Application Note 101
75 MHz Board Design

APPLICATION NOTE 101

75 MHz Board Design

Summary

At high bus frequencies, motherboard design becomes geometrically more difficult since a multiplicity of design factors interact to defeat the integrity of the design. Circuit traces become transmission lines, coupling occurs between traces, and power supplies require additional decoupling. Shorter lead length increases circuit density and produces greater thermal concerns. Clock generator location and clock-line layout becomes critical. High speed design can be made easier and more reliable using circuit simulation and other sophisticated design techniques and by considering the recommendations given in this application note.

This document applies to all Cyrix CPU unless otherwise specified.

Introduction

In recent years, the Cyrix x86 family of processors has gained increased performance due, in part, to higher bus speeds. This chapter is intended to help motherboard designers properly design a motherboard to accommodate a Cyrix 6x86MX operating at 75 MHz.

At this high frequency of operation, more attention must be paid to board characteristics such as:

- Trace to trace capacitance
- Transmission line effects
- Power supply decoupling
- Thermal constraints
- Clock distribution

In the following paragraphs these characteristics and methods for the reducition of these effects are discussed.

Trace to Trace Capacitance

The impedance of a capacitor gets smaller at higher frequencies (neglecting effective series inductance and resistance). Specifically, the impedance from one board trace to another gets smaller as board frequencies go up. In the equation below the impedance (Z) is calculated at a specific frequency (f), and for a typical capacitance between two board traces (C').

$$Z = \frac{1}{j\omega C} = \frac{1}{2j\pi f C'}$$

At 20 MHz with 5 pf capacitance between board traces, there is an interconnective impedance of:

$$Z = \frac{1}{j\omega C} = \frac{1}{2j\pi \times 20MHz \times 5pf} = 1591\Omega$$

Since the characteristic impedance of the lines is about 150 ohms the line impedance is negligible, creating a small amount of cross coupling. These coupling is further reduced when the driving source impedance is included.

At 75MHz, the effect of intertrace impedance becomes more pronounced:

$$Z = \frac{1}{j\omega C} = \frac{1}{2j\pi \times 75MHz \times 5pf} = 424\Omega$$

At still higher frequencies, the trace to trace impedance gets closer to the characteristic impedance of the source and transmission line and intertrace crosstalk becomes even more pronounced. For example, at 100MHz the impedance is 318 ohms, and this is only twice the characteristic impedance of the line.

In higher speed designs, signal integrity analysis should be performed to predict the effects of cross connection interference. The primary observable effect will be digital noise appearing at random. Setup and hold times are particularly succeptible and this may lead to unstable operation.

Expedient layout, such as long parallel bus runs between devices, can cause difficulty in troubleshooting noise and cross-talk problems. These problems can be largely prevented by employing short traces and minimizing parallel bus runs. If two lines must run a fairly long distance, two different trace paths are preferred. If device location prevents this approach, then the use of vias and jump-overs, every centimeter or two, is recommended.

Although there will be many more cross couplings to a given line due to line shuffling, the cross couplings will be to different lines and no one given line will bear the bulk of all the interference.

Boards with lower dielectric coefficients are recommended, as they have less lead-to-lead capacitance.

Capacitance can also be reduced by increasing the spacing between conductors, and increasing the interplane board material thickness.

Transmission Line Effects and Reflection

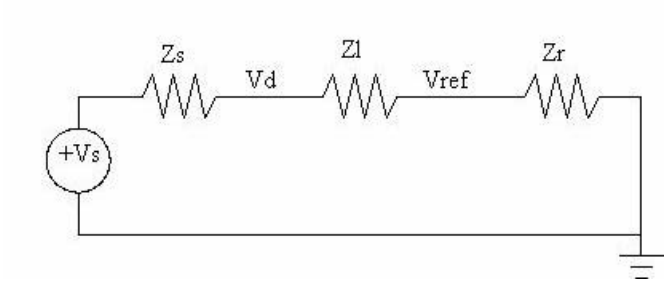

Electrical Model of a Transmission Line

A model of a transmission line is shown above. When a wavefront of source voltage Vs. is launched into a transmission line, the wavefront voltage (Vd) becomes:

$$Vd = Vs \times \frac{Zs}{Zs + Zl} \qquad \text{[Equation 1]}$$

where:

Vd is the drive voltage at the pin of the driver (that is after Zs)
Zs is the source impedance
Zl is the characteristic impedance of the transmission line.

This smaller wave then propagates down the line to the end of the transmission line where it encounters Zr, the receiver impedance. However, when a wavefront hits an impedance, it divides down in amplitude, then reflects, doubling the voltage difference after being divided by its impedance mismatch. The reflection amplitude is given by:

$$Vref = (Vw - Vi) \times 2 \times \frac{Zt}{Zt + Zo} \qquad \text{[Equation 2]}$$

where:

Vw is the incident wavefront voltage
Vi is the initial voltage at the reflection point
Zt is the impedance the wavefront hits
Zo is the characteristic impedance of the line the wave was following.

For example, if a 1 V wave on a line impedance of 100 ohms hits a 200 ohm impedance, and the 200 ohm impedance already has 1/2 volt on it, the apparent voltage of the reflected wave will be:

$$Va = \frac{1}{2}V + \left(1 - \frac{1}{2}\right)V \times 2 \times \frac{200}{200 + 100} = 1.166 \text{ V}$$

This reflected wave then propagates back, and the procedure repeats on the transmission side, reflecting back and forth until the steady state line voltage is reached. Depending upon how mismatched the source impedance, line impedance, and receiver impedance are, and the propagation time on the line, the reflections can take a long time to settle out.

When the source voltage has a rise time associated with it, the reflections that bounce back and forth are not distinctly identifiable, since the source wave may still be rising during the reflections. This will cause the reflected waves to take on a more rounded shape, which gives rise to the phenomenon commonly known as overshoot, undershoot, and ringing.

These characteristics should not be considered as noise, since they are an inherent part of the signal propagation. However, the ringing can cause erroneous operation on a motherboard. Additionally, the ringing can couple to other circuits through capacitive paths between traces.

The key to preventing these problems is to provide proper termination for the transmission lines. A properly terminated transmission line will not reflect any signal, and no overshoot, undershoot or ringing will occur.

Other line characteristics that affect a clean signal transmission are turning radii of wires and extraneous vias in a line. Ninety degree turns and vias in a path create discontinuities that result in transmission line impedance mismatch. These discontinuities cause signal reflections. To minimize these effects, use as few vias as possible, and use 45 degree turns instead of 90 degree turns.

Transmission Line Termination

Several common ways to reduce reflections on a transmission line include the use of diode clamping, series dampening resistors, balanced and unbalanced termination resistors, and resistive-capacitive termination.

Ideally, if the driver source, the transmission line and receiver impedances were all the same, the signal would propagate to the end cleanly with no reflections. This is rarely the case, so a proper termination technique should be employed. Note that on a PCI bus, reflections are a required part of the bus specification; therefore, so the following techniques should only be exercised on signals where reflections and ringing need to be kept to a minimum.

Diode Clamping is a common technique where a pair of normally reverse biased diodes are wired from the end of a transmission line to either Vcc or ground. When a wave hits the end of the line, it starts to reflect off the high impedance and tries to double in voltage; however, the fast diode starts to conduct and holds the overshoot or undershoot to a one diode voltage drop above or below the associated power rail. Thought of another way, the impedance of the diode, being variable according to its forward bias, will somewhat closely match itself to the transmission line impedance, thereby limiting the reflection to a diode drop. This method does not prevent the reflection, but limits the overshoot on the first reflection. Successive reflections will not be suppressed since they won't fall above the diode's bias requirement. The diode clamping technique is shown in the figure on the next page.

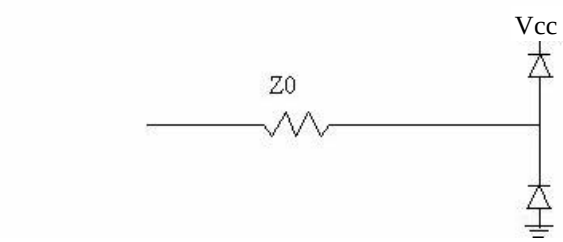

Diode Clamp Termination

Series resistive dampening is another technique used to limit ringing and reflections. By putting a small valued resistor (10 to 30 ohms) in series with the drivers, as shown below, the rise time of the signal at the transmission line is increased due to the time required to charge the line capacitance. This allows the reflection to be reduced in amplitude; however, this method also allows the source impedance of the driver to be more closely matched to the characteristic impedance of the transmission line. This does not address the reflection at the high impedance end of the line. Additionally, series resistors can limit the drive levels on a heavily loaded line to the point where a logical high may not have the noise immunity headroom for the circuit being driven. This method is shown in the figure on the next page.

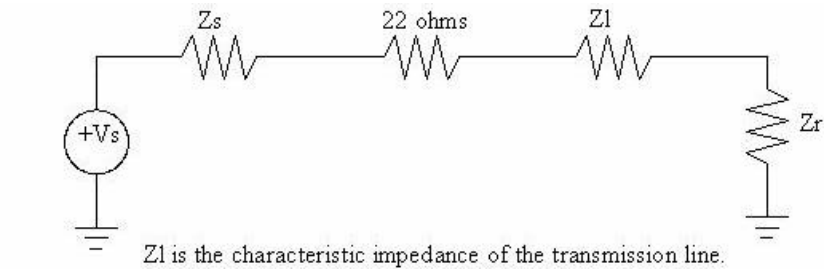

Series Resistor Termination

Another method of terminating a transmission line is the use of unbalanced termination resistors. This technique involves placing resistors at the end of the transmission line to Vcc or more preferably the ground bus. By selecting a resistor value that comes close to the characteristic impedance of the line, the reflected waveform voltage will tend to cancel out with the incident waveform voltage per Equation 2, and no apparent reflection will occur.

The unbalanced termination resistor (Zt) is selected so that when it is connected in parallel across the receiver (load) impedance (Zr), the resulting impedance is equal to the transmission line impedance. Selecting a termination resistor based on the transmission line impedance will often not only create a mismatch, but will heavily load down the line driver, possibly holding the steady state level of the signal below the required high level voltage. It also has the effect of forcing a normally tri-stated bus into a logical one or zero state, and increasing the current requirements during steady state signaling. The unbalanced resistor termination technique is shown below.

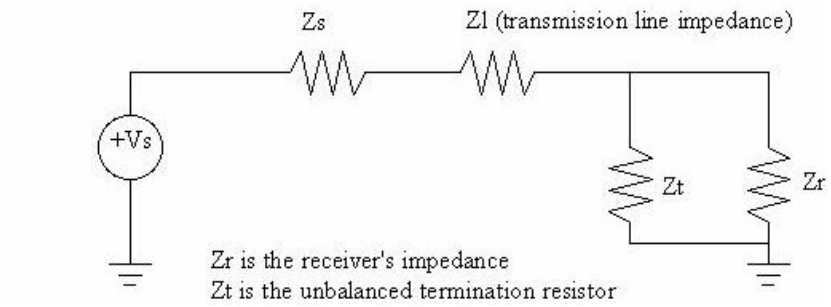

Unbalanced Resistor Termination

Balanced Termination Resistors

This technique is like the single resistive technique, except that by using two resistors, each twice the value of the line impedance and wiring the bus to both Vcc and to ground in a “T” network, the drive low current and the drive high current required to pull the bus is reduced by a factor of 2. It also allows the bus to “float” to the midpoint between the voltage rails. It does share the disadvantage of making the drivers work harder in steady state, but to a lesser extent.

RC Series Termination

This technique involves attaching an RC series network to the end of the transmission line. The RC series network is designed to provide an impedance that is matched to the characteristic impedance of the transmission line, but at a frequency determined by the rise and fall times of the incident wave. This will suppress any reflections while not wasting the DC drive capability of the source drivers. This method has the disadvantage that the line being terminated should be well categorized as to it's characteristic impedance. The RC phase diagram is shown on the next page.

The value of C should be $\frac{3 \times Tr}{Zo}$ where Tr is the rise time of the incident signal, and Zo is the characteristic impedance of the transmission line. This is derived from the following: At the primary frequency of $\frac{1}{2 \times Tr}$, you want a very low impedance path provided by C, and the rest of the characteristic termination impedance to be taken up by R. This is to allow the AC component to be absorbed by the termination, while allowing the DC portion of the wave to remain unloaded. If Zo is the transmission line impedance, and a factor of 10 lower AC impedance is chosen for the termination capacitor, then

$$|Zc| = \frac{1}{2\pi fC} = \frac{Zo}{10} \qquad \text{[Equation 3]}$$

Substituting $\frac{1}{2 \times Tr}$ for f, and solving for C we get:

$$C = \frac{10Tr}{Zo\pi} \approx \frac{3Tr}{Zo}$$

To choose an appropriate Series Resistor, we need to compute the total impedance of the RC network, and solve for the resistance value.

$$|Zo| = \sqrt{Zr^2 + |Zc|^2} \qquad \text{[Equation 4]}$$

The magnitude of Zo is equal to the square root of the sum of the squares of Zr (impedance of the resistor) and Zc (the impedance of the capacitor) as shown below.

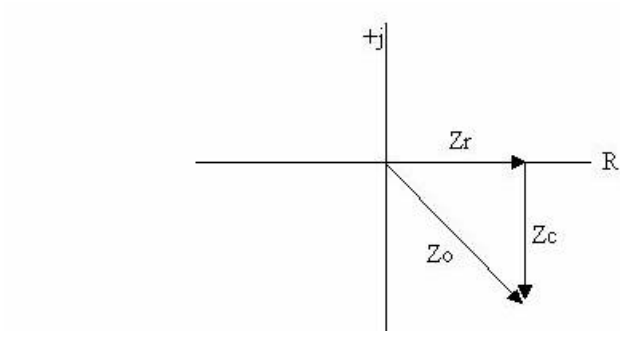

Phasor Diagram for Series RC Termination

For a 100 ohm characteristic line impedance and 4 ns rise time, C = 120 pF, Zc is found by:

$$Zc = \frac{1}{j\omega C} = \frac{-j}{\omega C}$$

$$Zc = -\frac{1}{2\pi \times 125MHz \times 120pf} = -10.6j$$

Solving Equation 4 for R, we get

$$R = \sqrt{Zo^2 - Zc^2} = \sqrt{100^2 - 10.6^2} = 99\Omega$$

So the complete RC termination circuit for a characteristic line impedance of 100 ohms with a signal rise time of 4 ns is shown in the figure below.

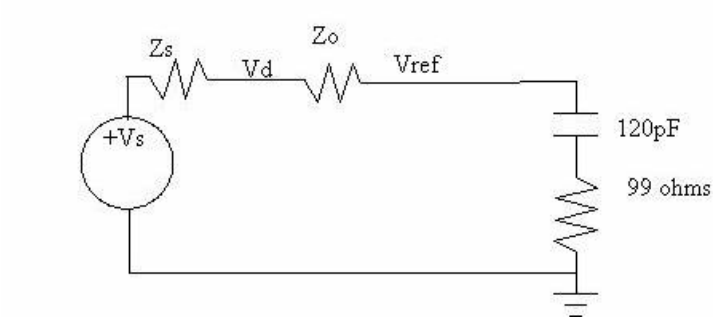

Series RC Termination

Of course, if the receiving element has a termination impedance that is within a factor of ten of the characteristic line impedance, its impedance should be considered in parallel with the termination network, and the network should be recalculated appropriately. This termination network will generally work well over a 60% to

150% range of the desired characteristic impedance. For example, if the network is designed for 100 ohms, it will work well from 60 ohms to 150 ohms line impedance.

Power Supply Decoupling

At higher bus frequencies, internal circuits and bus lines switch more often. Additionally, faster parts of a given technology often have faster signal rise and fall times. The more frequent and faster transitions dump more electrical energy through board capacitance and internal CMOS gate capacitance. Since all real life capacitors are lossy, this creates the need for a fast response from the power supply to keep the voltage supplying the circuits at a steady state.

At 75 MHz and beyond, almost all of the current being used by a CMOS circuit is due to switching the internal capacitances on and off very quickly, and the value of the current becomes more proportional to the operating frequency.

To supply a steady power supply voltage to the quickly switching circuits, the power supply must provide very low impedance to the high speed frequencies involved. This is accomplished using an adequate quantity and value of low ESR (Effective Series Resistance) capacitors.

Decoupling Capacitors

There are usually three values of decoupling capacitors used on motherboards, each addressing a different frequency spectrum required of the power supply. The large, bulk decoupling capacitors, which usually have a high ESR, are placed on a board to compensate for voltage regulator recovery delays and long leads from the power supply. They provide the ability to detune and decouple the low frequency changes in the current. One such example is the entry and exit from suspend mode where motherboard current can swing from near zero to 5 A or more. These large capacitors are selected to hold the power bus reasonably stable during the surges until the regulator can stabilize at the new current draw.

The next smaller valued capacitors, usually in the 0.1 pF range of value, is used to help decouple some of the higher frequency components that the self-inductance of the larger, bulk capacitors cannot deliver in microsecond times. These are usually chip capacitors with about 1 nH of self inductance. This type of capacitor usually has a self resonant frequency given by:

$$fr = \frac{1}{2\pi\sqrt{LC}} \qquad \text{[Equation 5]}$$

$$fr = \frac{1}{2\pi\sqrt{1 \times 10^{-9} \times 100 \times 10^{-9}}} = 16MHz$$

Above these frequencies, the 0.1 pF capacitors are not as effective at shunting out the higher frequencies. Additionally, the inductance of the copper lines or planes from the capacitors to the actual device pins adds a high frequency impedance to the capacitor and further reduces the capacitor's ability to decouple higher frequencies.

At higher bus speeds, however, even these chip capacitors have too much self inductance in series with their capacitance to meet the demands of high speed, high current switching circuitry. A third, smaller value of capacitor, specifically selected for high frequency decoupling is required. These are usually 0.01 µF capacitors with very low ESR (about 500 pH) whose self-resonance is given by Equation 5 or:

$$fr = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{500 \times 10^{-12} \times 10 \times 10^{-9}}} = 71MHz$$

In the following example, the high frequency portion of the power grid of a motherboard is shown.

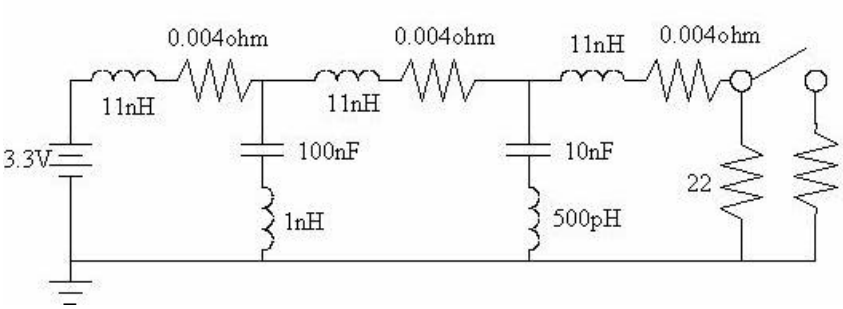

High Frequency Model of Power Grid

In this circuit example, the switch selects between a 150mA and a 300 mA load, emulating the nominal current switching of 3 power pin pairs of a Cyrix 6x86MX processor during normal operating conditions. The 11nH inductances in series with the 4 milliohm resistors approximates the voltage and ground plane interconnects between decoupling capacitors and the Cyrix 6x86MX processor power pins. The inductances in series with the capacitors approximate the decoupling capacitors, complete with their series ESR. The 3.3V battery emulates the voltage regulator and it's bulk filtering capacitors. This circuit, when simulated in SPICE, results in voltage spikes, as measured at the switch simulated pin, of 1.5 volts, and approximately 1 nanosecond in duration. Fifteen capacitor pairs will provide sufficient decoupling of the high frequency noise induced by the switching logic in Cyrix 6x86MX processor. In general, use one pair per 300mA of load current unless the wiring to the capacitors is excessively long.

Low Frequency Decoupling Techniques

Low Frequency decoupling, also known as bulk decoupling, provides energy storage capacity for the processor during prolonged transients during which the power source is unable to meet due to its transient response time. Power sources vary widely in their response to transients from the processor. Typically, the response time of most power sources range from 1 to 100ms.

Power plane and power trace inductance to the processor must be minimized. This inductance impedes the response time of the power source, resulting in more intense voltage spikes during processor switching transients. By placing the decoupling capacitors as close to the processor or load as possible, the inherent inductance in the supply leads is compensated for, resulting in a smoother power supply voltage.

Another very important parameter to be concerned with in choosing capacitors for bulk decoupling is Effective Series Resistance (ESR). High ESR capacitors will drop too much voltage across the capacitor at high currents, adding to the voltage transients. Precautions should be made to select bulk decoupling capacitors which have low ESR values.

The formula for calculating the required bulk capacitance including ESR is as follows:

$$C = I\frac{dt}{dV} - I \times ESR \qquad \text{[Equation 6]}$$

where:
I = Worst case change in current for the supply
dt = Response time of the power source to the change in current
dV = Maximum allowable change in voltage for the load
ESR = ESR of all capacitors used in decoupling the power source.

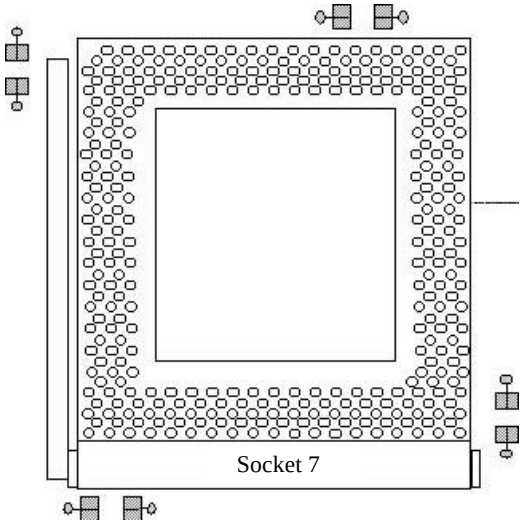

Bulk Decoupling Capacitors Placement

High Frequency Decoupling Techniques

The high frequency decoupling technique is required to smooth transients that are too fast to be handled by bulk decoupling capacitors. Typically, ceramic high frequency decoupling capacitors are used as they have low inductance.

These capacitors should be placed as close to the power pins of the processor as physically possible. It is also critical that the traces from the capacitors go directly to the power and ground planes with leads as short as possible. This will prevent any board inductances from voiding the decoupling effectiveness.

Cyrix recommends that a motherboard designer use a minimum of 30 high frequency decoupling capacitors. Surface mount capacitors that are evenly distributed close to the processor work well. Suggested values are fifteen 0.1 µF capacitors and fifteen 0.01 µF capacitors. These capacitors are in addition to the low frequency decoupling capacitors.

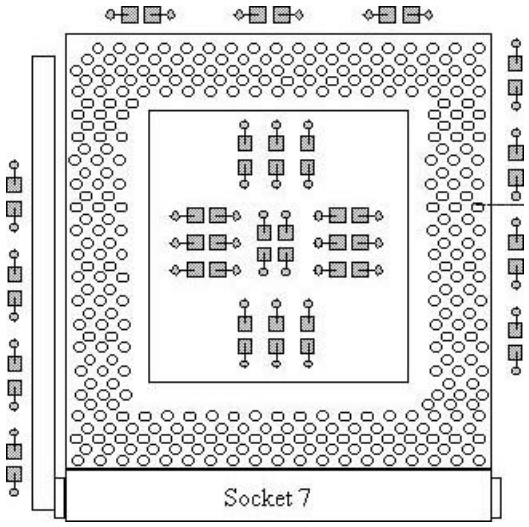

Placement suggestions for high frequency decoupling capacitors for Socket 7.

Thermal Concerns

As the frequency increases on a board, the transistors switch more frequently and thus dissipate more power. In order to design a board to run at higher frequencies, the intersignal capacitance must be reduced by making signal lines as short as possible.

As the frequency is increased, the thermal problem is aggravated by the shorter line lengths recommended in previous paragraphs. The more densely populated board components block the airflow and to a lesser degree absorb each others radiate heat.

The overall capability of a motherboard to dissipate heat is controlled by many variables including case air flow, case volume, ambient temperature, number and position of any plug-in cards, as well as the number and thicknesses of signal and power planes.

Assuming natural convection cooling in an operating environment temperature of 35° C, an average junction temperature of devices not exceeding 100° C, a minimum of two signal layers of 1 oz/square foot thickness, 2 power planes of 3 oz/square foot thickness, the ideal maximum power density to put on a motherboard is 0.04 watts per square centimeter, evenly distributed over the board. If adequate airflow is included in a design, this dissipation recommendation can be made higher. Refer to the *Cyrix Thermal Considerations Application Note 105* for more information.

Clock Distribution

As clock frequency goes up, clock-skew tolerance through the motherboard drops dramatically. Clock skew results from using different clock drivers and the existence of different clock path lengths. Since the chipsets, CPU, and peripheral components all use the same clock, it is critical that the clock signal be correctly distributed in a way that minimizes skew from chip-to-chip.

It is best to keep clock skew, clock reflections and noise to a minimum to ensure that the clock waveform shape stays as close to ideal as possible. Another approach in clock distribution design is to introduce delay purposely in the appropriate clock lines. Circuits dependent on setup or hold times are often the most affected by clock skew.

Since all components must be in different physical positions, distance to each component from the clock driver will vary. By planning a clock distribution network early in the layout process, clock skew can be minimized.

Ideally, a good design should start with equal clock flight times to every chip using the clock signal. There are three distinct distribution methods for a clock line: random spawning, daisy chaining, and radial distribution. Each has their merits and deficiencies and are discussed below. Random spawning distribution is essentially randomly attaching a branching path to the main clock path wherever it is most convenient. A random-spawning clock distribution is shown in the figure below.

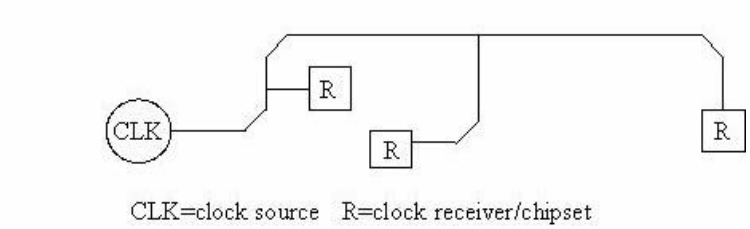

Random Spawning Clock Distribution

In this example, the clock is distributed from one point, and receiving components pick off the clock signal from the net as needed. This method is the most convenient for routing a clock signal, however it results in so many impedance mismatches along the route, unterminated ends, and mismatched flight times, that it will cause timing skew errors, reflections, and noise pickup from nearby lines. This distribution technique is only valid for circuits with high noise margins, and clock rise times in excess of the longest flight time. This distribution technique is a poor choice for high speed motherboard design and *should be avoided*.

Daisy Chaining, as shown in the figure on the next page, is a more dependable version of the random spawning method of clock signal distribution.

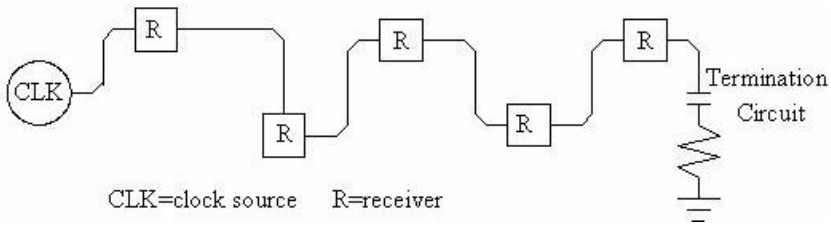

Daisy Chaining Clock Distribution

In this example, a single transmission line goes from source to receiver to receiver, without creating any unnecessary “stubs lines” along the way. When the clock launches a waveform, it will propagate all the way to the termination circuit, which will prevent reflections and maintain the original clock waveform along the entire circuit. Often, a small valued resistor of 10 to 22 ohms is placed immediately after the clock source. This has several benefits. First, it helps to match the impedance of the clock source to the clock transmission line, thereby reducing signal launch ringing and reflections.

Secondly, the series resistance, coupled with the distributed capacitance of the transmission line will slow down rise and fall times of the wave by filtering out the high frequency harmonics of the clock line. This will have the effect of “rounding the edges” of the clock signal. This will reduce electromagnetically radiated noise from the motherboard, and also reduce the amount of noise the clock line introduces into nearby PCB conductors. Higher frequency harmonics couple to other signal lines much easier than the primary frequency of the clock.

Thirdly, if reflections are produced off the end of the transmission line due to a mismatch between the characteristic line impedance and the termination circuit impedance, the series resistor will help to dissipate the resultant ringing. Using the daisy chain method of clock distribution will help minimize impedance mismatches along the transmission line by eliminating stub-outs. If a branch is required in a daisy chain distribution, it should be kept to no longer than 2 centimeters in length.

One flaw in the daisy chain distribution technique is that each receiver gets clocked at a slightly different time since the wave takes a finite time to propagate to the end of the transmission line. If each chip's setup and hold times have sufficient slack in them relative to all the other drivers and receivers in the other chips using the clock, then the daisy chain solution will work. If cycle timings are so close that the total flight time approaches the slack between chipsets, then setup or hold time violations are imminent.

Radial Clock Distribution is shown in the figure below. It addresses most of the deficiencies of the above distribution techniques.

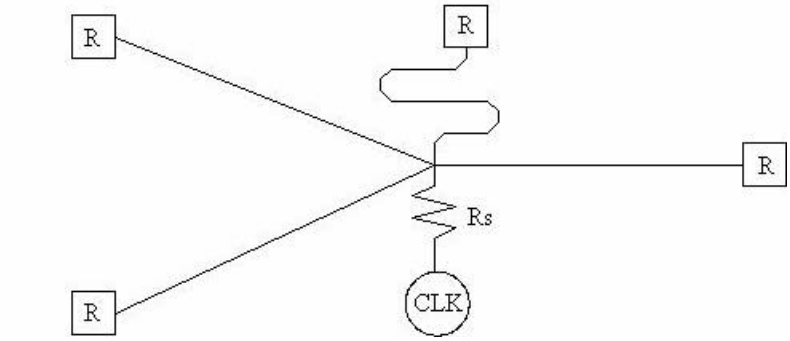

Radial Clock Distribution

In this distribution scheme, the clock generator is centralized, and distribution is made from one common point. Additionally, if any physical distance from the clock to the receivers are significantly shorter, a path must be chosen to that receiver to compensate for the shorter flight times. The goal in the radial distribution technique is to provide the same wave flight time to each receiving element in the circuit. That way, clock skew is minimized and setup and hold times are equalized, resulting in a more robust design.

A small source resistor is also sometimes included in the layout for the same reasons given in the daisy chain method. Additionally, an RC termination circuit should be added to each receiving element to keep the reflections and ringing to a minimum. The RC termination is preferred since it does not induce any DC loading to the clock source. A disadvantage to this method is that the clock layout needs to be carefully considered and the remaining layout will become more difficult.

Regardless of which distribution method is used, it is important that reflections and ringing be minimized in any design. Reflections will appear as periodic noise on the clock line and will cause noisy coupling to nearby circuits, setup/hold time violations, data corruption, and radiated electromagnetic field emissions. The safest way to eliminate this is by properly terminating all clock transmission lines. Another good design practice in clock distribution is to surround the clock line by two ground traces along its flight as shown in the figure on the next page.

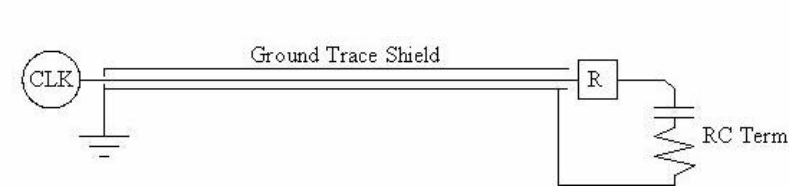

Clock Shielding

In this example, a pair of ground lines run adjacent to the clock line, and the termination circuit completes the transmission line. It is important to ground the shield at the source and not at the receiving end of the line. If this looks like a coaxial solution, you're on the right track. At high frequencies, a shield around the clock line will help prevent transmission of clock harmonics to surrounding circuitry, as well as prevent surrounding circuitry from coupling their signals into the clock line. Ideally, the clock and shield should be located next to the ground plane, and there should be no signals on adjacent planes in the vicinity of the clock line.

There are a number of different ways of distributing the clock line in a board design. Through careful planning and by keeping the impedances of the line and termination consistent one can minimize:

- Clock skew for different loads
- Reflections
- Ringing
- Clock waveform harmonics
- Coupling of signals and clocks.

Additional General Design Considerations

There are a number of other design considerations to account for in a high speed design. Any turns taken by a high speed signal should be made smoothly, or at 45 degree angles and never at 90 degree angles, except for via (plane to plane connection) holes. Additionally, every via hole for a signal introduces two 90 degree turns which will appear as slight impedance discontinuities. Excessive use of vias or turns on a signal will increase reflective ringing on the line.

When designing a motherboard for use with the Cyrix 6x86MX processor and Intel® Pentium™ processors, it is very important to compare the timing specifications of the two parts, megahertz for megahertz, and design for the worst case of each AC signal timing. Most of the signal valid, setup, and hold times are compatible, but a few differences exist, and can vary from revision to revision of the parts. For example, if a certain signal has a setup requirement of 3 ns on a Pentium processor, and 2 ns on an Cyrix 6x86MX processor, then the worst case of 3 ns should be used in design and simulation. If a clock to valid time for a signal is 5 ns for an Cyrix 6x86MX processor and 3 ns for a Pentium processor, then the 5 ns time should be used for design. Generally, use the longest setup times, the longest output launch valid times, and the shortest hold times. Be sure to consider all required frequency variations of parts when choosing the correct design timings.

Since design timing errors are difficult to prevent in most high speed designs, it is very useful to include debug tools in the first phases of a design. For example, putting a zero ohm jumper in line with the control signals of the CPU allows easier probing of the signals and the ability to insert a delay or resistive element if needed. Placing test points in a first pass design also allows control signals and data busses to be probed easier without introducing excessive bus capacitance. As a design is debugged, small test points and debugging aids can be extracted from the circuit to reduce fabrication costs.

10/20/97
C:\DATAOEM.CUR\!M2MXAP\101ap.fm5

Revision 1.1 Corrected typos
Revision 1.0 Name Change, Final Release