
FAST BURST SRAM DESIGN CONSIDERATIONS

by Rickie TUTTLE

INTRODUCTION

This application note will discuss the definition of the fast Burst RAM (BRAM[®]) for high performance systems, and consider application issues to ensure proper design techniques. Design issues become more acute when changing clock rates from 12-25MHz to 33-66MHz and beyond. SGS-THOMSON has developed specialty processor specific cache Burst RAMs that have on-chip registers and counters to aid in synchronous cache designs. However, additional cautions must be taken in any high speed design to calculate noise margins, ground bounce, and PCB trace layout for transmission line affects. This application note will review the various issues for high speed design as applied to the high performance BRAM products. We will discuss many scenarios, and offer cost effective solutions that have been verified in a laboratory setting.

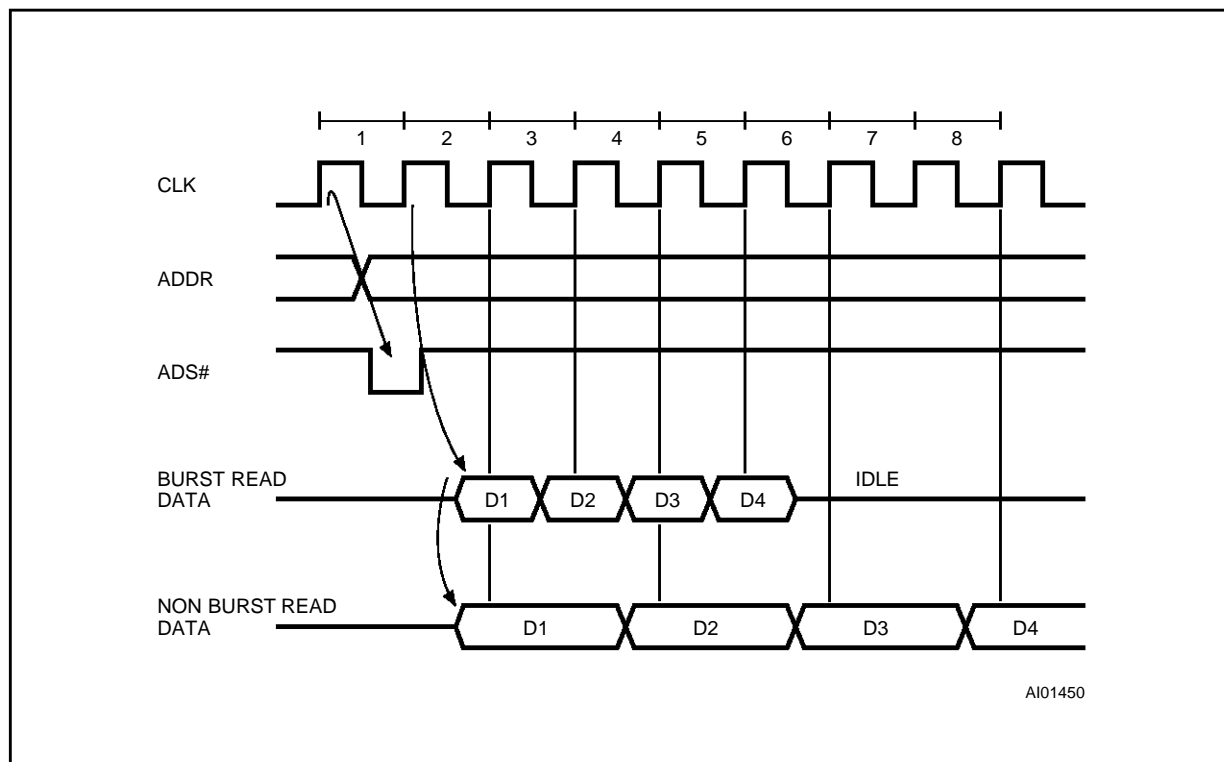
PROCESSOR SPECIFIC BURST SRAM

SGS-THOMSON Microelectronics has begun to develop a Processor specific SRAM family for cache applications. These devices are designed to support synchronous systems using the cache burst line fill algorithm present in the 680X0 and X86 family of processors. This includes the Pentium[™] and PowerPC[™] class of processors.

The present product offering includes the M63532P 32Kx32 BRAM for the X86 processor family, as well as RISC and PowerPC processors. These devices have clock to data valid access speed of 7-8ns, which meets system design requirements for clock rates as fast as 75MHz. Present microprocessors employ a burst cache line fill algorithm of either 128 bits (32-bit i486[™]) or 256 bits (64-bit Pentium). Each processor has a unique burst sequence requirement as four (4) bursts of data with optional byte parity. Therefore, the two devices incorporate a different 2-bit counter MODE along with address and data registers. The X86 family uses an interleave burst count algorithm, while the RISC and PowerPC use a linear burst count.

The BRAM is organized as 32Kx32 with control pins that connect directly to the processor or cache controller. For the 32-bit processors, the burst algorithm allows 16 bytes of data (128 bits) to be transferred in five (5) subsequent processor clock periods, whereas a non-burst sequence would take eight (8) clock periods (assuming zero wait). The 64-bit processors would transfer 32 bytes of data (256 bits) in the same (5) clock periods. This can be seen in Figure 1 where data and clock cycles for both burst and non-burst are shown. As a comparison, the burst sequence with a 66MHz clock rate supporting a 3-1-1-1 burst algorithm would allow data to be transferred at 178M bytes/second in a 32-bit data bus system, and 355M bytes/second in a 64-bit data bus system (this comparison assumes 3-1-1-1 clocking scheme for each cache burst line-fill). In a non-burst sequence, only 133M bytes/second can be transferred in the 32-bit bus system, and 266M bytes/second in the 64-bit bus system with the same 66MHz clock. Table 1 compares data flow rate with various burst sequences and inserted wait states. The advantage of the burst algorithm over a non-burst transfer is obvious for high performance systems. Only one (1) 32Kx32 cache Burst RAM is needed to design a direct mapped 128K byte cache for a 32-bit system, and two (2) devices for a 256K byte cache in a 64-bit system. A 512K byte cache solution can be provided using only four (4) 32Kx32 Burst SRAMs in a 64-bit data I/O system.

Figure 1. Burst vs. Non-Burst Cycles



KEY TIMING PARAMETERS

The BRAM registers the address at the rising clock edge where ADS# is asserted (asserted in this context means an active signal). After the BRAM has registered the base address, the three subsequent addresses are generated by the on-chip burst counter. The BRAM allows wait states to be inserted in the burst sequence by using the ADV# (Burst Advance) control pin. The on-chip burst counter is not affected, and the burst sequence will continue as expected after the wait state(s) is completed. It is up to the cache control design to abort a cache line-fill or begin a new one. A burst cycle in the BRAM is always initiated by asserting either ADSP# (Address Strobe Processor) or ADSC# (Address Strobe Cache Controller) at the rising edge of the clock. These signals allow the device to register a new base address, and set-up for the burst sequence. Chip select(s) must be asserted whenever ADSP# or ASC# is active at the rising edge of clock to remain selected.


The BRAM also allows immediate control of the data bus by featuring an asynchronous Output Enable (OE#). Although clock-to-data-out is the limiting access time, the fast Output Enable can provide data within 5-6ns after being asserted. This allows the OE# signal to be asserted late in the read cycle without access penalty. Using OE# as the output data bus control has no affect on the burst counter sequence. During wait-states the data bus can be driven to a high impedance state. Once the burst cycle continues, the data bus can become active for the duration of the burst without consequence to the burst algorithm. Detail timing examples and parameters are given in their respective product data sheets.

Table 1. Burst Sequence and Data Rate

Clock Rate	Cycle Time (ns)	Bata Bus Width	Burst Sequence	Mbytes / sec ⁽¹⁾
33 MHz	30	32 bit	2-1-1-1	106
66 MHz	15	32 bit	2-1-1-1	213
66 MHz	15	64 bit	2-1-1-1	426
66 MHz	15	32 bit	3-1-1-1	178
66 MHz	15	64 bit	3-1-1-1	355
66 MHz	15	32 bit	4-1-1-1	152
66 MHz	15	64 bit	4-1-1-1	304
66 MHz	15	32 bit	3-2-2-2	118
66 MHz	15	64 bit	3-2-2-2	237
66 MHz	15	32 bit	4-2-2-2	106
66 MHz	15	64 bit	4-2-2-2	213

Note: 1. Table presumes the noted burst sequence count for each burst cache line-fill.

Figure 2. Standard and Pipeline Burst



Note: 1. The ADS# signal applies to either ADSP# or ADSC# on the BRAM.

The rest of this application note will be given to the discussion of high speed design considerations using the BRAM product family from SGS-THOMSON Microelectronics. Specific attention will be given to system noise in relation to the control signals mentioned previously. Careful consideration must be given to signal integrity in order to avoid violation of set-up and hold times relative to the rising clock edge. General system design issues will be presented as well as specific issues regarding the cache BRAM device. This document is a brief summary of high speed system design issues, and some prior knowledge is presumed. Further information can be obtained by contacting SGS-THOMSON SRAM applications. Suggested reading can be found in references 6 and 7.

SYNCHRONOUS CIRCUITS

The BRAM is a synchronous CMOS device, with on-board registers and self-timed write circuits. Synchronous circuits can help relieve some system noise issues over asynchronous designs due to the settling time inherent between rising clock edges. For example, a high speed output switching after a rising clock edge can induce noise to a system, but will have time to settle before the next rising clock edge where data is sampled to an input. Although CMOS technology offers high speed and power preferred by many designers, care must be taken to incorporate a good clock distribution network. Clock distribution becomes much more of an issue at high frequencies because skew and slow rise times degrade operating margins ⁽¹⁾. Also, high performance logic families are affected by system noise, and generate system noise. The mixture of 3 volt and 5 volt technology has introduced voltage, current, and impedance matching issues to the system designer. Synchronous circuits can provide extra margin, but set-up and hold times relative to the rising clock edge must be maintained in all instances.

GROUND BOUNCE


Simply stated, ground bounce is a result of high speed device output switching where the load capacitance is discharged through the package inductance to ground. This discharge occurs when the device outputs switch from high-to-low causing an instantaneous capacitive current ($I=CdV/dt$) to flow through the ground lead. The discharge current causes a positive voltage pulse to appear across the package inductance of the ground lead as ($V=Ldi/dt$). The positive ground bounce pulse raises the internal ground of the system devices affected, making the inputs appear to get a negative pulse added to the input equal to the ground bounce amplitude (potential difference). Logic levels, specifically V_{IH} , can be degraded to an unknown logic state.

CMOS circuits with fast switching outputs and high drive capability can generate a large voltage transient between device ground and the ground plane. Additionally, CMOS outputs can pull to rail voltages, where V_{OH} is 4.5 to 5V. The higher the output voltage, and the greater the slew rate and current drive of the output transistor, the worse the transient voltage on the ground lead. This is the ground bounce phenomenon see (Figure 3). The worst case induced noise condition is when all device outputs switch from high-to-low simultaneously. Conversely, ground bounce can be observed in low-to-high transitions, but the magnitude is much smaller because of the absence of load current in the ground lead ⁽²⁾. Adding more devices with more switching outputs aggravates the issue. Ground bounce can be controlled, but it cannot be completely eliminated.

Remember that the voltage drop across the ground lead is the product of the rate of current change and the ground lead inductance. This positive ground bounce or overshoot is followed by a an undershoot voltage. Here are some methods to minimize ground bounce noise:

1. Take steps to reduce parasitic inductance between the package and ground and V_{CC} . This includes using a PCB with power and ground planes, or a good power distribution network.
2. Whenever possible, design synchronous circuits to allow maximum signal settling time.

Figure 3. Ground Bounce Voltage



3. Use various techniques to slow switching transition edges to deter the current rate of change. PCB traces exhibit parasitic resistance, inductance, and capacitance that depend on trace length, topology, and material. A common technique is series dampening resistors at the driver's output in the range of 20 to 40 ohms connected to the receiver's input. Series resistors also limit signal overshoot and undershoot.
4. Use separate device ground lead traces for signal and data drivers when possible ⁽³⁾.
5. Select packages that exhibit less inductance. Surface mount packages have less pin inductance than DIP devices.

Ground bounce can be controlled by connecting zener diodes and filter capacitors between power and ground to control over voltage transients and filter AC noise. The zener diode must be able to sink the amount of current that is dissipated in the circuit. Connect the cathode of the zener diode to V_{CC} , and the anode to ground. Over-voltages on V_{CC} will cause the zener diode to break down and clamp the power supply voltage to the specified zener voltage. In extreme cases, positive voltage ground bounce transients can cause negative V_{CC} spikes that may forward bias the diode. This condition allows the zener diode to act as a protection device ⁽⁴⁾.


Finally, ground bounce is dampened by decoupling capacitors across each active device between power and ground. It is common practice to use $0.001\mu\text{F}$ to $0.1\mu\text{F}$ for each device to provide high frequency AC decoupling or filtering. Capacitance can be calculated as $Q = CV$, and $C = (Idt/dv)$. These decoupling capacitors also provide the instantaneous switching current required by the output drivers of the device. More decoupling capacitance is required for the high speed BRAM because the output slew rate for the 32Kx9 BRAM is 2.0V/ns . The BRAM offers numerous power and ground pins in a 44-pin PLCC. Therefore, it is suggested that decoupling capacitors be placed across two to four pairs of V_{CC} -to-GND pins for each device. A total decoupling capacitance of $0.4\mu\text{F}$ to $0.8\mu\text{F}$ is suggested. The actual decoupling capacitance needed for the BRAM may vary depending on the application. Additional capacitive isolation is obtained by placing a $1\mu\text{F}$ to $10\mu\text{F}$ across the power plane inputs to the board or module. Usually this is an electrolytic type capacitor used to filter lower frequencies that are generated off the board (such as the 60-Hz line frequency). The smaller capacitors on the active devices are generally an encapsulated ceramic type used to filter high frequency harmonics generated on the board ⁽⁵⁾.

PCB DESIGN ISSUES

The printed circuit board is commonly regarded as a low impedance medium that interconnects electrical signals in a specific path from one point in a circuit to another. This is true where signals have slow edge rates (rise and fall times). However, PCB can no longer be considered an ideal interconnect for high-speed circuit design. PCB traces subjected to fast edge rates will exhibit impedance characteristics where the designer must account for the effects of distributed capacitance, inductance, and propagation delay. Fast edge rates will cause the PCB trace to appear as a transmission line that can cause signal distortion, overshoots, undershoots, and crosstalk between adjacent lines. Here are several tips to keep in mind regarding PCB design:

1. Use power and ground planes or a grid network. A low impedance ground is necessary since any noise on the ground line may be coupled into signal lines. A power grid network can be used to reduce the capacitance between the isolated power and ground busses.
2. Avoid sharp corners in trace layout, and keep vias to a minimum due to transmission line effects.
3. Crosstalk between signal lines can be avoided by separating signal layers by the power planes. If a double layer board is used, then ground traces should be placed between high frequency parallel signal lines with fast edge rates.
4. Use bypass or decoupling capacitors on each active device in the range of 0.001F to 0.1µF, and a low frequency capacitor (1µF to 10µF) between the power inputs of the board or module. (Remember additional decoupling capacitance is suggested for the active high speed cache BRAM devices.)

Figure 4. PCB Trace Diagram



5. High speed system design requires short traces for signal lines. This will minimize overshoot and undershoot ringing, and simplify timing considerations. Ringing is due to the intrinsic inductance and capacitance of the trace itself, as well as the lumped capacitance at the end of the line. Intrinsic inductance and capacitance are reduced by shortening the signal lines.
6. Longer trace lengths should be terminated. A common terminating method is a series dampening resistor placed close to the output of the driver ⁽⁶⁾. This is an impedance matching technique to be discussed later.

Most PCB boards employ microstrip, stripline, or a combination of both. Microstrip construction on a double-sided printed circuit board with power and ground nets can suffice for a low-to-medium performance, and a low density board. Higher performance PCB will incorporate strip-line construction where power planes will isolate signal layers to provide a higher quality power and ground as well as premium signal integrity. A microstrip line (see Figure 4) is a strip conductor that represents a signal line on a PCB separated from a ground plane by a dielectric. Proper control of the line's thickness, width, and distance from the ground plane allows the designer to calculate the characteristic impedance of the line. A strip line usually consists of a copper line centered in a dielectric between two conducting planes (see Figure 4). The characteristic impedance of the line can be calculated knowing the line's thickness, width, material dielectric constant, and distance between power planes ⁽⁷⁾.

It is important to be able to calculate the characteristic impedance of the PCB trace, because the trace acts as a load to the driving circuit. In addition, the transmissive effect introduces a finite signal delay from source (driving end) to destination (receiving end). The impedance of the output driver is very low (5-10 ohms) compared to the impedance and load of the PCB trace. The equivalent ideal circuit for a transmission line is represented by distributed L+R and C (inductance, resistance, and capacitance: LRC) as shown in Figure 5A. You can minimize undershoot and overshoot voltage transients (ringing) by matching the impedance of the driving device to the line's characteristic impedance. As an example, and for the purpose of this discussion, we ignore the series resistance where R=0, and the parallel resistance (G) is infinite to make a "lossless" ideal case transmission line. The following equations can be derived from these assumptions, and are considered valid in theory study:

$$L = Z_0^2 C_0 \text{ therefore,}$$

$$(1) Z_0 = (L_0 / C_0)^{1/2} \text{ defined in ohms (ideal transmission line)}$$

$$(2) \text{tpd}_0 = (L_0 C_0)^{1/2} \text{ per unit length (nanoseconds)}$$

where Z_0 is the trace impedance, tpd_0 is the propagation delay time per unit length, L_0 is the inductance of the trace per unit trace length, and C_0 is the capacitance per unit length.




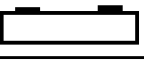
For a microstrip PCB trace separated by a dielectric (15 mils glass filled epoxy) on one side, and free air on the other 3 sides as a rectangular trace which is 10 mils wide, 1.5 mils thick, using a typical $C_0 = 2\text{pF}/\text{inch}$ and $L_0 = 10\text{nH}/\text{inch}$, then:

$$Z_0 = (10\text{nH} / 2\text{pF})^{1/2} = 70\Omega, \text{ and}$$

$$\text{tpd}_0 = (10\text{nH} \times 2\text{pF})^{1/2} = 0.14\text{ns} \text{ each unit length.}$$

Note that Z_0 is independent of the length of the trace. Table 2 gives examples of various transmission lines, geometry layout, and typical parameters ⁽⁸⁾.

Table 2. Transmission Line Impedances

TYPE	GEOMETRY TYPE	T_0 ns/in.	Z_0 Ω
WIRE OVER GROUND		0.14	70-170
MICROSTRIP LINES		0.15	30-150
STRIP LINE		0.19	15-100
PCB TRACES		0.16	50-200

AI01454

TRANSMISSION LINES AND TERMINATION METHODS

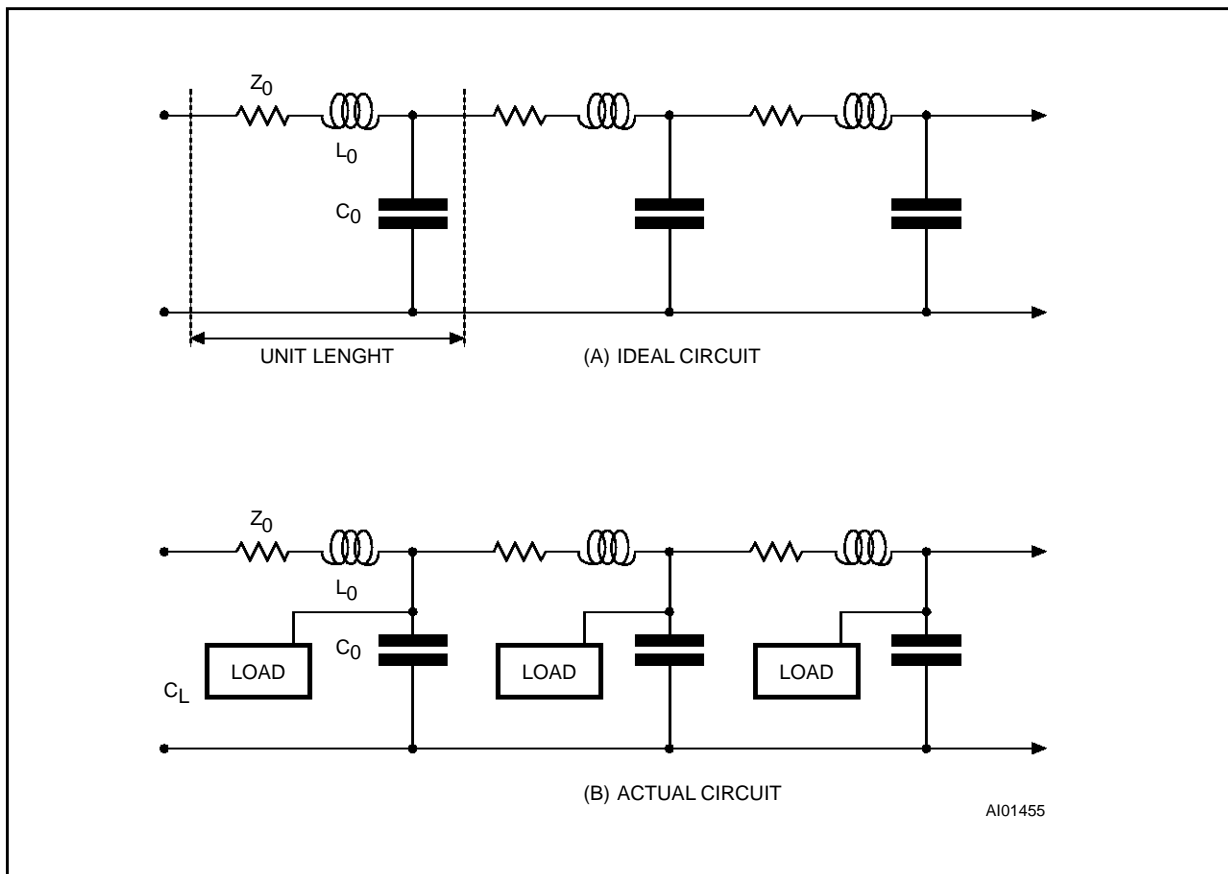
Transmission Line Summary

In general, a printed circuit board (PCB) connection or trace appears as a transmission line in high speed designs due to high frequencies and fast edge rates of signal drivers. In practice, transmission lines on PCBs are designed to be as lossless (non-resistive) as possible. This simplifies the mathematics required for analysis, and allows assumptions regarding characteristic line impedance, intrinsic capacitance, and inductance. The ideal case would be for all signals between ICs to travel over constant impedance transmission lines that are terminated in their characteristic impedances at the load ⁽⁹⁾. However, this is not the case since transmission line effects depend upon the frequency of the applied signal and the trace load. Signal traces usually have multiple loads as shown in Figure 5B. Treat a PCB trace as a transmission line if:

$$Tr \leq (2L)(TL) ,$$

where L is the length of the trace, Tr is the rise or fall time of the driver, and TL is the loaded transmission line propagation delay per unit length. Maximum trace length can be calculated as: $L = (Tr / 2TL)$ ⁽¹⁰⁾.

Figure 5. Transmission Line Equivalentts



Signal noise known as reflections and ringing are caused by impedance mismatch due to non-uniform transmission line impedances. These voltage transients can compromise dynamic system noise margins and cause unwanted circuit switching. The amplitude of the reflection is determined by the impedance mismatch between the line and the load, and is a factor of the driver's rise/fall time in relation to the trace length and propagation delay⁽¹¹⁾. There are two mismatch coefficients to consider: 1) source to load, and 2) source to line. If the impedance mismatch is large enough, either a positive or negative voltage can reflect back from the load to the source, where the voltage either adds or subtracts from the original voltage. Conversely, a mismatch between the line and source can cause a voltage reflection which can reflect back to the load causing an invalid input bias. There will always be some signal overshoot and undershoot, but the objective is for the designer to accurately predict these effects for a given design.

The ideal case for calculating transmission line effects is when the characteristic impedance appears as pure resistance, where $Z_0 = (L/C)^{1/2}$, and $R_S = Z_0 = R_L$. Z_0 is the impedance, and R_S is the source resistance and R_L is the load resistance. The maximum energy transfer from source to load occurs in this instance with no consequent reflection voltage transients. Half of the energy is dissipated in the source resistance (R_S), and half is dissipated in the load resistance (R_L) where the line or trace is lossless. Extra energy is available at the load when the load resistor is larger than the line's characteristic impedance, and a voltage reflection is reflected back to the source. This is called the underdamped condition, because the load does not use the full energy available. One of the transmission theory rules is that a positive voltage reflected at the load toward the source results in a corresponding negative current. The total voltage or current in the line is the algebraic sum of all incident and reflected waves. Therefore, a positive voltage reflection results in a negative current reflection and vice versa. If the load resistance is smaller than the line impedance, the load will attempt to use more energy than is available. Obviously the load cannot use what does not exist, so a reflection occurs back to the source to send more energy. This condition is known as overdamped. From a system design viewpoint, the most conservative approach is the overdamped termination condition such that no energy is reflected back to the source⁽¹²⁾.

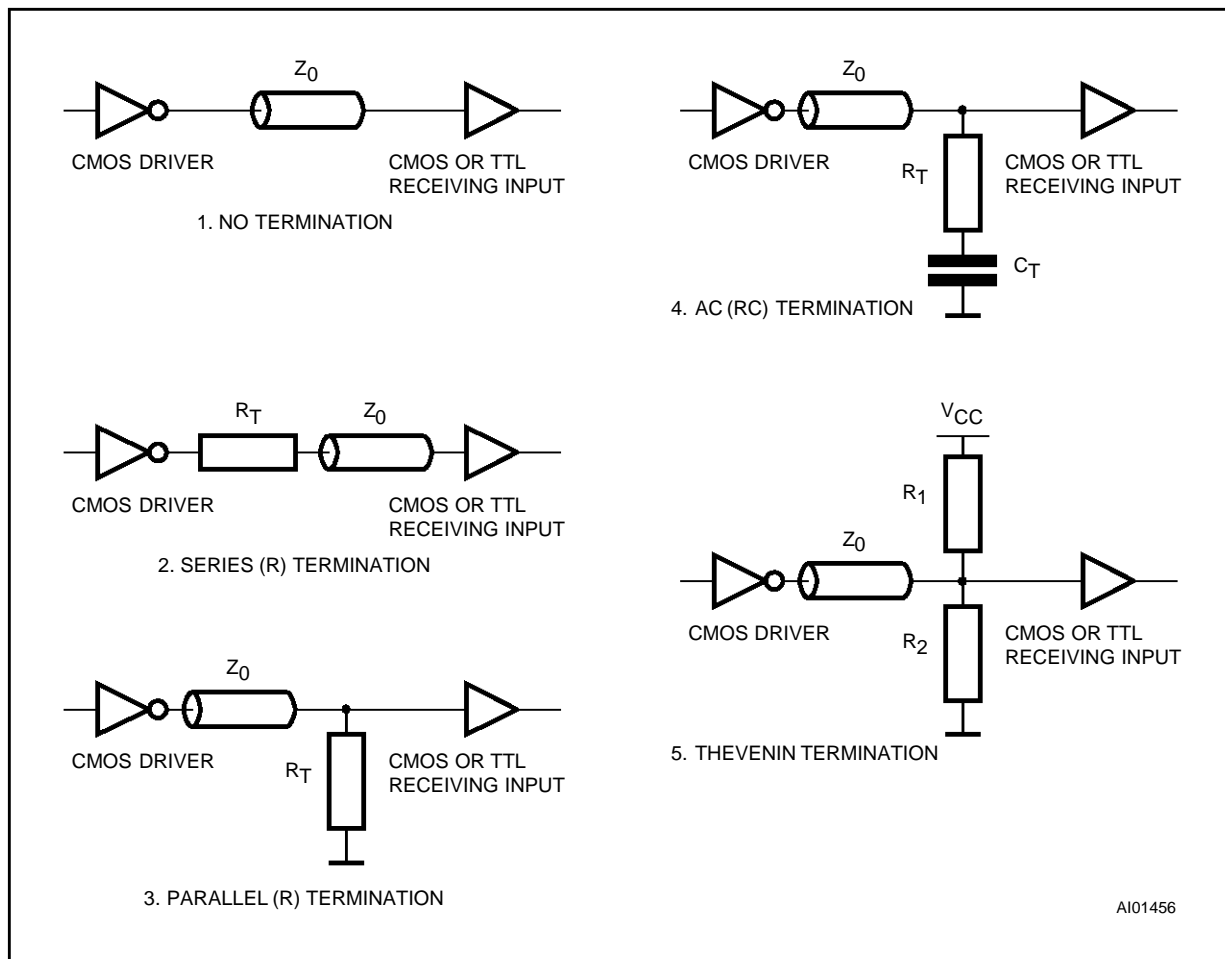
Transmission Line Termination Methods

There are several methods of termination available to the designer (see Figure 6)⁽¹³⁾. Each method will have to be reviewed and studied as it applies to a unique design. Termination advantages and efficiency, design preferences, board space, cost, time-to-market, and device types are factors that must be considered.

Series termination is recommended if the output impedance of the driving device is less than the loaded characteristic impedance of the trace. A single resistor is connected close to the output driver node between the node and the PCB trace. The resistance should be chosen such that the total source impedance is close to the characteristic impedance of the line. In a 50 ohm impedance environment, the resistor value is typically 20-40 ohms. One benefit of series termination is that it does not create a DC current path where V_{OL} and V_{OH} levels could be degraded. Series termination is not suggested for clock distribution lines because they are usually driving more than one load device.

Parallel termination is a single resistor connected to V_{CC} or ground at the end of the trace, observing no distortion in the signal. The resistor value should be equal to the characteristic impedance of the line. For a 50 ohm impedance line, a 50 ohm resistor to power or ground would be required. This type of termination creates a DC current path and requires more power dissipation. This can cause excessive degradation to V_{OL} or V_{OH} depending if the resistor is connected to power or ground. Since the AC switching current is the major component of the overall current, the DC current can be ignored. With high frequencies, the parallel method allows easy and inexpensive transmission line termination, but must review the source drive current ability and possible V_{OH} degradation.

Figure 6. Termination Line Methods



Thevenin termination is where two resistors are joined to make a resistor divider network at the end of the transmission line. One resistor is tied to V_{CC} and the other resistor is tied to GND. The transmission line is connected to the junction of the resistors. The parallel equivalent value should be equal to the characteristic impedance of the line.

The designer must choose resistor values to avoid settling in an invalid logic state while the line is not being driven. This method requires more discrete resistors than the parallel termination, but the DC component is less since the resistor values are at least two times greater. It is suggested that for this type of termination when using the cache BRAM device, that control signals should settle to a logic high (V_{IH} or greater) when not being driven. This is probably true for most CMOS devices. Table 3 gives examples of resistor values for line impedance when using the Thevenin termination method ⁽¹⁴⁾.

$$R_{Thevenin} = (R1R2)/(R1+R2) = Z_0$$

$$V_{Thevenin} = (R2/R1+R2)V_{CC}$$

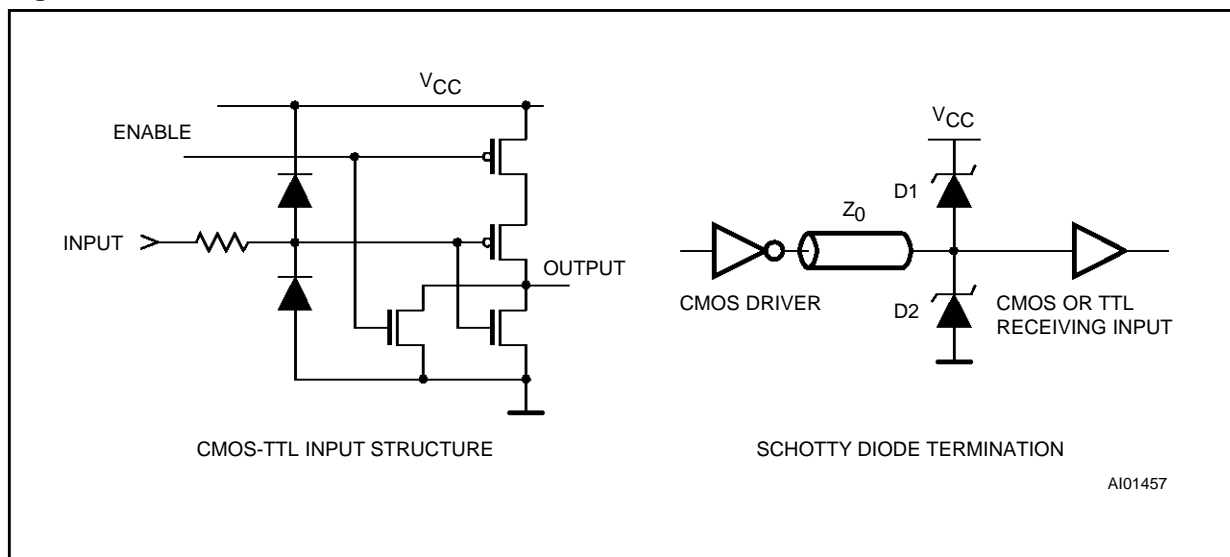
AC termination uses an RC network series combination connected between the transmission line signal and ground. The capacitor blocks DC current but allows the AC signal to flow to ground during high frequency switching. The RC time constant should be greater than twice the loaded line delay. It should be noted that the capacitance chosen must be able to supply or absorb the energy contained in a positive or negative going glitch, and that it must be small enough to avoid delaying the signal's rise and fall time to a given design limit⁽¹⁵⁾. AC termination is recommended where low power dissipation is paramount. Sometimes a pull-up resistor to V_{CC} is required to avoid the failure condition of an unknown logic state. RC networks are now available in SIP (single-in-line) packages, making this method a little easier for manufacturing.

The final termination method we will discuss is the **Schottky diode termination**. Many bipolar logic families and BiCMOS as well as CMOS logic families have inherent clamping diodes on the inputs and outputs (see Figure 7). PCB traces connected to such devices have an automatic silicon diode termination method incorporated such that voltage reflections may be suppressed by the clamping action of the diodes. Otherwise, use a discrete method of termination by connecting Schottky diodes to the transmission line. A pair is required such that they appear stacked across V_{CC} and ground (Figure 7). The junction of the anode and cathode of the two diodes is connected to the PCB trace. The diode switching time must be at least as fast as the signal rise time, and leads should be as short as possible. This method of termination can significantly reduce signal undershoot and overshoot. The Schottky diode low forward voltage (V_f = 0.3V to 0.45V) clamps the input signal to (V_{CC}+V_f) if D1 is activated, and V_f below ground if D2 is activated.

Table 3. Thevenin Resistor Termination

Z ₀ Ohms	R1 Ohms	R2 ohms	V _{TH} , V _{CC} = 5V
50	81	130	3.08
70	113	182	3.09
75	121	190	3.05
80	130	208	3.07
90	140	234	3.12
100	162	260	3.08
120	194	312	3.08
150	243	390	3.08

Figure 7. Diode Termination



There are several advantages using diode termination:

1. Impedance matched lines are no longer a major factor.
2. A precise controlled transmission line environment is not required.
3. Diode clamping action reduces undershoot and overshoot.
4. Diodes can replace termination resistors or RC networks.
5. If ringing is a problem on a trace, the diodes can be added when the system is debugged.

Diodes typically cost more than resistors, but make many allowances for the designer. The CMOS input structures of the 32Kx9 BRAM from SGS-THOMSON employ a single silicon clamping diode to V_{CC} . This type of termination for the 32Kx9 BRAM would require only one Schottky diode to ground. As stated previously, termination methods and preferences must be studied for each application.

BRAM DESIGN HINTS

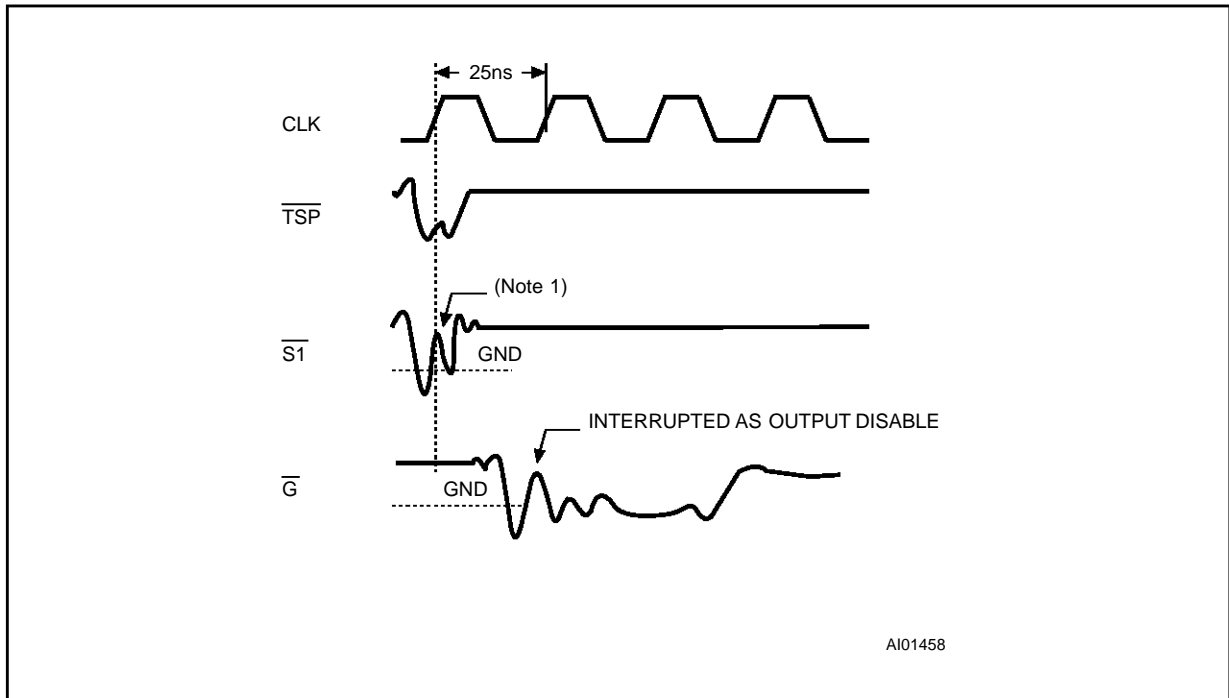
This application note has been directed to provide helpful hints and reminders for high performance system design. In an effort to provide actual data for design consideration, we have studied two different 256K byte cache modules for the Power Macintosh 8100 series. The processor is a MPC601-80MHz with a 40MHz external clock. Using (8) 32Kx9 BRAMs (M62486A) per module, we found it necessary to use termination techniques on certain control signals to ensure proper operation of the cache module. In a lab setting, using the Power Macintosh 8100 series system, the cache modules were plugged into the Burndy ELF card edge connector. The self power-up memory test passed as well as memory diagnostic software when using the above mentioned termination methods to dampen transmission line effects.

We found it necessary to use parallel line termination on Chip Select (S1#), Output Enable (G#), and Write Enable (W#) on the cache modules obtained from various vendors. Schottky diode, AC (RC) network, and parallel R terminations were successful. The AC network consisted of a 50ohm resistor as RL with a 47pF CL. The 50 ohm parallel resistor and Schottky diode termination methods were employed in a separate experiment. Figures 8 and 9 show the difference between terminated and un-terminated signal lines (these are not actual pictures, but serve only as examples to visually depict circuit differences). Figure 10 is a block diagram example showing the placement of required line terminations for the 32Kx9 BRAM in a 256K byte cache application. Depending on the physical layout of the board or module, design sophistication, and cost constraints, one of the mentioned transmission termination methods must be used for proper system operation.

SUMMARY

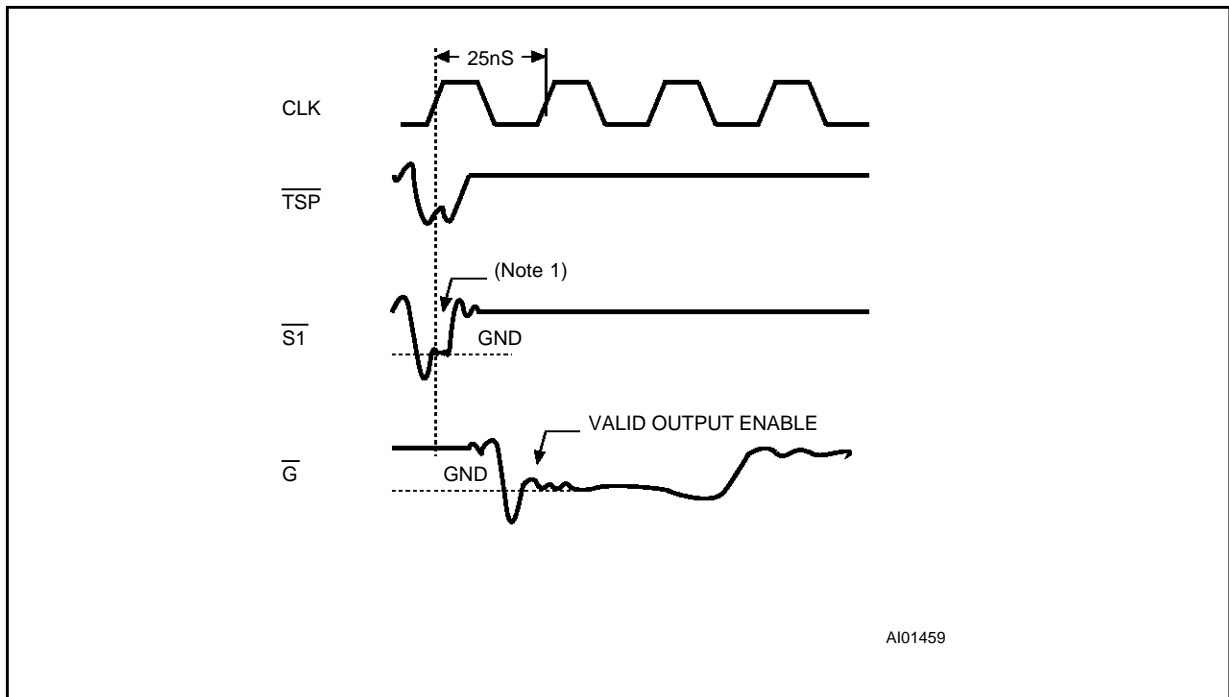
This application note has reviewed general high speed design issues relative to the processor specific cache BRAM from SGS-THOMSON. These issues are encountered in all high speed designs, especially where clock rates exceed 33MHz. Ground bounce noise, PCB design considerations, and transmission line effects and methods of termination have been discussed. Low cost termination techniques have been given and verified for designs using SGS-THOMSON high speed 32Kx9 BRAM in a 256K byte cache application. The same design issues should be noted and evaluated for any application using fast processor specific SRAMs such as the M62486A, M63532P and future high performance cache SRAMs.

Figure 8. Non-terminated Control Lines



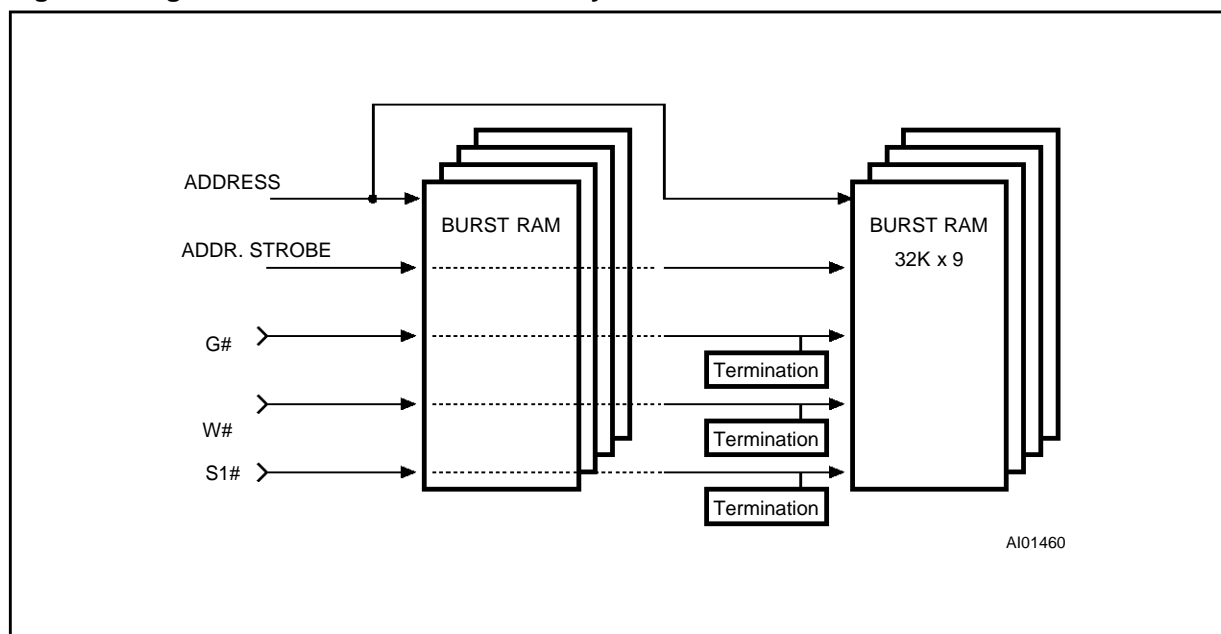
Note: 1. Causes erratic behaviour due to set-up time violation at rising clock edge.

Figure 9. Terminated Control Lines



Note: 1. Ringing dampened by termination of transmission line corrects error at rising clock edge.

Figure 11. Signal Lines Termination for 256K byte Cache



REFERENCES

1. "Tips for High-Speed Logic Design", Cypress Semiconductor, *Applications Handbook*, (Aug. 1991), p.1-29.
2. Suren Kodical, "Simultaneous Switching Noise", *1993 IDT High-Speed CMOS Logic Design Guide*, (1993), 42.
3. Reference 1, p. 1-30.
4. "Protection, Decoupling, and Filtering", Cypress Semiconductor, *Applications Handbook*, (Aug. 1991), p. 1-35.
5. *High-Speed-Board Design Techniques*, Advanced Micro Devices, (Publication # 16356, 1992), 3-4.
6. Blood, William R., *MECL System Design Handbook*, Motorola, Inc., (1983), 26-27.
7. "System Design Considerations When Using Cypress CMOS Circuits", Cypress Semiconductor, *Applications Handbook*, (Aug. 1991), p.1-13.
8. Suren Kodical, "Characteristics of PCB Traces", *1993 IDT High-Speed CMOS Logic Design Guide*, (1993), 50.
9. Reference 7, p.1-2.
10. Reference 8, 51.
11. Reference 7, p1-4.
12. Reference 7, p.1-5.
13. "Low Clock Skew Drivers and Their System Design Considerations", Motorola, Inc., *Timing Solutions*, (1993 Rev.3), 10.
14. Reference 6, 59.
15. Reference 7, p. 1-17.

Information furnished is believed to be accurate and reliable. However, SGS-THOMSON Microelectronics assumes no responsibility for the consequences of use of such information nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of SGS-THOMSON Microelectronics. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. SGS-THOMSON Microelectronics products are not authorized for use as critical components in life support devices or systems without express written approval of SGS-THOMSON Microelectronics.

© 1996 SGS-THOMSON Microelectronics - All Rights Reserved

® BRAM is a registered trademark of SGS-THOMSON Microelectronics

i486 and Pentium are trademarks of INTEL Corporation
PowerPC is a trademark of IBM Corporation

SGS-THOMSON Microelectronics GROUP OF COMPANIES
Australia - Brazil - Canada - China - France - Germany - Hong Kong - Italy - Japan - Korea - Malaysia - Malta - Morocco - The Netherlands -
Singapore - Spain - Sweden - Switzerland - Taiwan - Thailand - United Kingdom - U.S.A.

