Low Skew Clock Drivers and Their System Design Considerations

Prepared by
Chris Hanke
CMOS Design Engineering
Gary Tharalson
CMOS/TTL Product Planning Management

This application note addresses various system design issues to help ensure that Motorola's low skew clock drivers are used effectively in a system environment.

Low Skew Clock Drivers and Their System Design Considerations

ABSTRACT

Several varieties of clock drivers with 1ns or less skew from output-to-output are available from Motorola. Microprocessor-based systems are now running at 33MHz and beyond, and system clock distribution at these frequencies mandate the use of low skew clock drivers. Unfortunately, just plugging a high performance clock driver into a system does not guarantee trouble free operation. Only careful board layout and consideration of system noise issues can guarantee reliable clock distribution. This application note addresses these system design issues to help ensure that Motorola's low skew clock drivers are used effectively in a system environment.

INTRODUCTION

With frequencies regularly reaching 33MHz and approaching 40-50MHz in today's CISC and RISC microprocessor systems, well controlled and precise clock signals are required to maintain a synchronous system. Many microprocessors also require input clock duty cycles very close to 50%. These stringent timing requirements mandate the use of specially designed, low skew clock distribution circuits or 'clock drivers.' However, just plugging one of these parts into your board does not ensure a trouble free system. Careful system and board design techniques must be used in conjunction with a low skew clock driver to meet system timing requirements and provide clean clock signals.

Why are Low Skew Clock Drivers Necessary

An MPU system designer wants to utilize as much of a clock cycle as possible without adding unnecessary timing guardbands. Propagation delays of peripheral logic do not scale with frequency. Therefore, as the clock period decreases, the system designer has less time but the same logic delays to accomplish the function. How can he get more time? A viable option is to use a special clock source that minimizes clock 'uncertainty.'

A simple example illustrates this concept. At 33MHz, T_{cycle} = 30ns. An FCT240A, for example, has a High-Low uncertainty of the min/max spread of t_{PLH} to t_{PHL} of approximately 3.3ns. If 1.7ns of pin-to-pin skew due to the actual part and PCB trace delays is also considered, then only 25ns of the clock period is still available. The worst case t_P of clock-to-data valid on the 88200 M-Bus is 12ns, which leaves only 13ns to accomplish additional functions. In this case 17% of a cycle is required for clock distribution or clock 'uncertainty,' which is an unacceptable penalty from a system designer's point of view. At 50MHz this penalty becomes 25%. A maximum of 10% of the period allotted for clock distribution is an acceptable standard.

If multiple levels of clock distribution (one clock driver's output feeding the inputs of several other clock drivers) are necessary due to large clock fan-outs, the additional part-to-part skew variations add even more to the clock uncertainty. Standard logic has always been specified with a large (and conservative) delta between the minimum and maximum propagation delays. This delta creates the excessive amount of clock 'uncertainty' which the system designer has been forced to design into his system, even though it is not realistic. When system frequencies were below 16MHz this large clock penalty could be tolerated, but as the above example points out, not anymore. A clock driver's specs *guarantee* this min/max delta to be a specific, small value. To reduce the clock overhead to manageable levels, a clock driver with minimal variation (<5%) from a 50% duty cycle and guaranteed low output-to-output and part-to-part skew must be used.
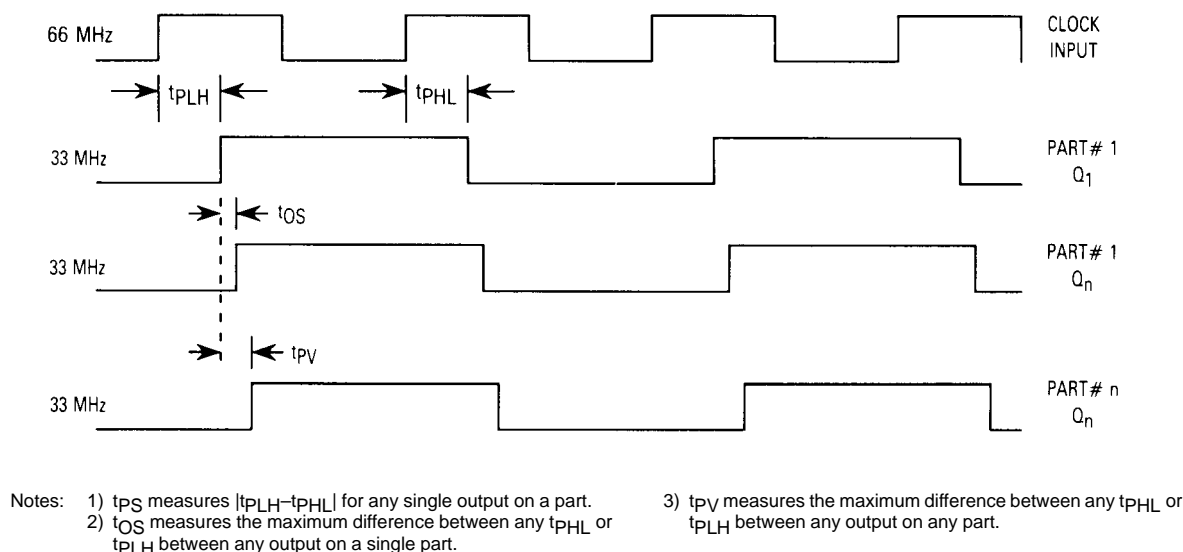
DEFINITIONS

A typical clock driver has a single input which is usually driven by a crystal oscillator. The clock driver can have any number of outputs which have a certain frequency relationship to the clock input. Clock driver skew is typically defined by three different specs. These specs are graphically illustrated in Figure 1.

The first spec, t_{OS}, measures the difference between the fastest and slowest propagation delays (any transition) between the outputs of a single part. This number must be 1ns or less for high-end systems.

The second, t_{PS}, measures the difference between the high-to-low and low-to-high transition for a single output (pin). This spec defines how close to a 50% duty cycle the outputs of the clock driver will be. For example, if this spec is 1ns (±0.5ns), at 33MHz the output duty cycle is 50% ±3.5%. A clock driver which only buffers the crystal input, creating a 1:1 input to output frequency relationship, can be a problem if a very tight tolerance to a 50% duty cycle is required. In this situation the output duty cycle is directly dependent on the input duty cycle, which is not well controlled in most crystal oscillators. The clock driver's outputs switching at half the input frequency (÷2) is a common relationship, which means that the outputs switch on only one edge of the oscillator, eliminating the output's dependence on the duty cycle of the input (crystal oscillator frequency is very stable).

The third spec, t_{PV}, measures the maximum propagation delay delta between any given pin on any part. This spec defines the part to part variation between any clock driver (of the same device type) which is ever shipped. This number reflects the process variation inherent in any technology. For CMOS, this spec is usually 3ns or less. High performance ECL technologies can bring this number down into the 1-2ns range. Another way to minimize the part-to-part variation is to use a phase-locked loop clock driver, which are just now becoming available.

Notes: 1) t_{PS} measures $|t_{PLH}-t_{PHL}|$ for any single output on a part.
2) t_{OS} measures the maximum difference between any t_{PHL} or t_{PLH} between any output on a single part.
3) t_{PV} measures the maximum difference between any t_{PHL} or t_{PLH} between any output on any part.

Figure 1. Timing Diagram Depicting Clock Skew Specs Within One Part and Between Any Two Parts

An important consideration when designing a clock driver into a system is that the skew specs described above are usually specified at a fixed, lumped capacitive load. In a real system environment the clock lines usually have various loads distributed over several inches of PCB trace which can contribute additional delay and sometimes act like transmission lines, so the system designer must use careful board layout techniques to minimize the total system skew. In other words, just plugging a low skew clock driver into a board will not solve all your timing problems.

DESIGN CONSIDERATIONS

Figure 2 is a scale replication of a section of an actual 88000 RISC system board layout. The section shown in the figure includes the MC88100 MPU and the MC88200 CMMU devices and the MC88914 CMOS clock driver. The only PCB traces shown are the clock output traces from the MC88914 to the various loads. For this clock driver the output-to-output skew (t_{OS}) is guaranteed to be less than 1ns at any given temperature, supply voltage, and fixed load up to 50 pF.

In calculating the total system skew, the difference in clock PCB trace length and loading must be taken into account. For an unloaded PCB trace, the signal delay per unit length, t_{PD}, is dependent only on the dielectric constant, e_r. of the board material. The characteristic impedance, Z_O, of the line is dependent upon e_r and the geometry of the trace. These relationships are depicted in Figure 3 for a microstrip line.[1] The formulas for t_{PD} and Z_O are slightly different for other types of strip lines, but for simplicity's sake all calculations in this article will assume a microstrip line.

The equations in Figure 3 are valid only for an unloaded trace; loading down a line will increase its delay and lower its impedance. The signal propagation delay (t_{PD}') and characteristic impedance (Z_O') due to a loaded trace are calculated by the following formulas:

$$t_{PD}' = t_{PD} \sqrt{1 + \frac{C_d}{C_O}}$$

$$Z_O' = \frac{Z_O}{\sqrt{1 + \frac{C_d}{C_O}}}$$

C_d is the distributed load capacitance per unit length, which is the total input capacitance of the receiving devices divided by the length of the trace. C_O is the intrinsic capacitance of the trace, which is defined as:
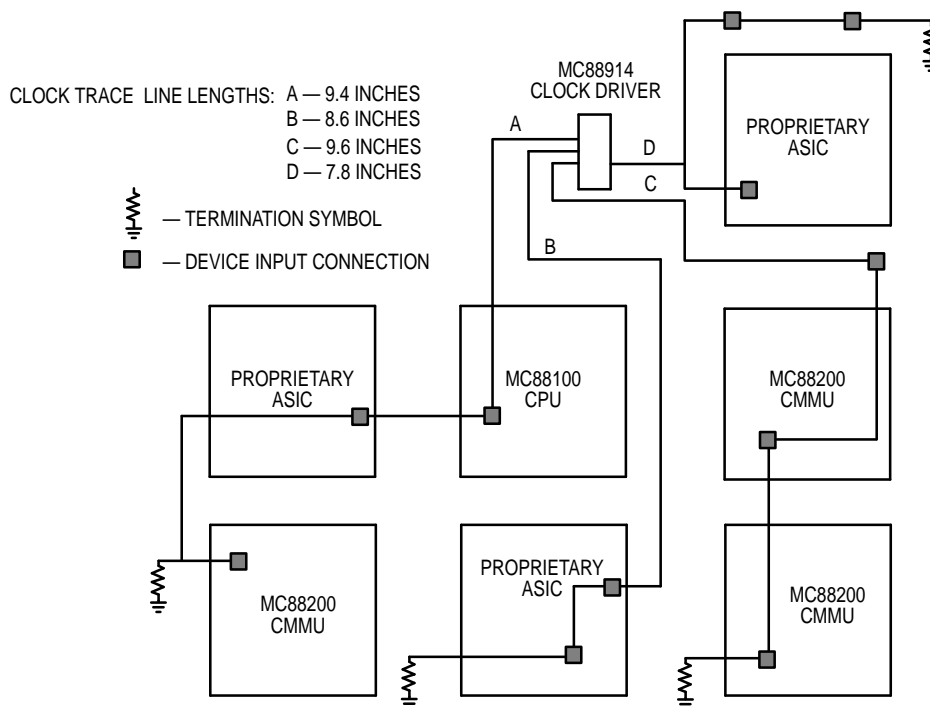
$$C_O = \frac{t_{PD}}{Z_O}$$

Assuming typical microstrip dimensions and characteristics as w = 0.01 in, t = 0.002in, h = 0.012in, and e_r = 4.7, the equations of Figure 3 yield Z_O = 69.4Ω and t_{PD} = 0.144ns/in C_O is then calculated as 2.075pF/in. If it is assumed that an MC88100 or 88200 clock input load is 15pF, and that two of these loads, in addition to a 7pF FAST TTL load, are distributed along a 9.6in clock trace,

$$C_d = (2 \times 15 + 7)pF/9.6in = pF/in.$$

The loaded trace propagation delay and characteristic impedance are then calculated as

$$t_{PD}' = 0.243ns/in \text{ and } Z_O' = 41\Omega.$$

Looking at trace C in Figure 2, the two MC88200's are approximately 3 inches apart. Using the calculated value of t_{PD}', the clock signal skew due to the trace is about 0.7ns. Since these two devices are on the same trace, this is the total clock skew between these devices. Upon careful inspection of all the clock traces, it can be seen that clock signal skew was accounted for and minimized on this board layout. The longest

CLOCK TRACE LINE LENGTHS: A — 9.4 INCHES
B — 8.6 INCHES
C — 9.6 INCHES
D — 7.8 INCHES

— TERMINATION SYMBOL

— DEVICE INPUT CONNECTION

MC88914
CLOCK DRIVER

PROPRIETARY ASIC

PROPRIETARY ASIC

MC88100 CPU

MC88200 CMMU

MC88200 CMMU

PROPRIETARY ASIC

MC88200 CMMU

Figure 2. Scale Representation of an Actual 88000 System PCB Layout
(Only sections of the board related to the clock driver outputs are shown.)

distance between any 88K devices on a single clock trace is about 4.5 inches, which translates to approximately 1.1ns of skew. The two 88K devices farthest away from the clock driver (traces a and c), are located at almost exactly the same distance along their respective traces, making the clock skew between them the 1ns guaranteed from output to output of the clock driver. This means that the worst case clock skew between any two devices on this board is approximately 2.1ns, which at 33MHz is 7% of the period. Without careful attention to matching the clock traces on the board, this number could easily exceed 3ns and the 10% cut-off point, even if a low skew clock driver is used.

MICROSTRIP LINE CROSS-SECTION

W

T

DIELECTRIC

H

GROUND OR POWER PLANE
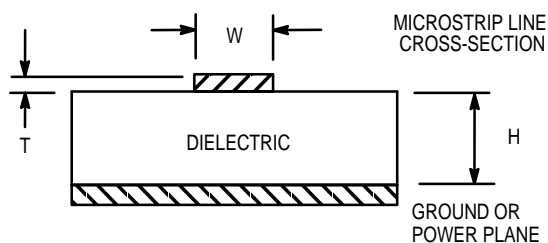
$$Z_O = \frac{87}{\sqrt{e_r + 1.41}} \ln\left(\frac{5.98h}{0.8w + t}\right)$$

$$t_{pd} = 1.017 \sqrt{0.475\, e_r + 0.67}\ \ ns/ft$$

WHERE:
  e_r = Relative Dielectric Constant of the Board Material
  w, h, t = Dimensions Indicated in a Microstrip Diagram

Figure 3. Formulas for the Characteristic Impedance and Propagation Delay of a Microstrip Line (Ref 1)

CLOCK SIGNAL TERMINATIONS

Transmission line effects occur when a large mismatch is present between the characteristic impedance of the line and the input or output impedances of the receiving or driving device. The basic guidelines used to determine if a PCB trace needs to be examined for transmission line effects is that if the smaller of the driving device's rise or fall time is less than three times the propagation delay of a switching wave through a trace, the transmission line effects will be present.[2] This relationship can be stated in equation form as:[3]

$$3 \times t_{PD}' \times \text{trace length} \leq t_{RISE} \text{ or } t_{FALL}$$

For the MC88914 CMOS clock driver described in this article, rise and fall times are typically 1.5ns or less (from 20% to 80% of V_{CC}). Analyzing the clock trace characteristics presented earlier for transmission line effects, 3 x 0.243ns/in x trace length ≤1ns (1ns is used as 'fastest' rise or fall time). Therefore the trace length must be less than 1.5 inches for the transmission line effects to be masked by the rise and fall times.

Figure 4 shows the clock signal waveform seen at the receiver end of an unterminated 0.5 inch trace and an unterminated 9 inch trace. These results were obtained using SPICE simulations. which may not be exact, but are adequate to predict trends and for comparison purposes. The 9 inch trace, which is well beyond the 1.5 inch limit where transmission line effects come into play, exhibits unacceptable switching characteristics caused by reflections going back and forth on the trace. Even the 0.5 inch line exhibits substantial overshoot and undershoot. Any unterminated line will exhibit some overshoot and undershoot at these edge rates.

Clock lines shorter than 1-1.5 inches are unrealistic on a practical board layout, therefore it is recommended that CMOS clock lines be terminated if the driver has 1-2ns edge rates. Termination, which is used to more closely match the line to the load or source impedances, has been a fact of life in the ECL world for many years (reference 1 is an excellent source for transmission line theory and practice in ECL systems), but CMOS and TTL devices have only recently reached the speeds and edge rates which require termination. CMOS outputs further complicate the issue by driving from rail to rail (5 V), with slew rates exceeding those of high performance ECL devices.

A. UNTERMINATED 0.5-INCH, 41Ω TRANSMISSION LINE

B. UNTERMINATED 9-INCH, 41Ω TRANSMISSION LINE
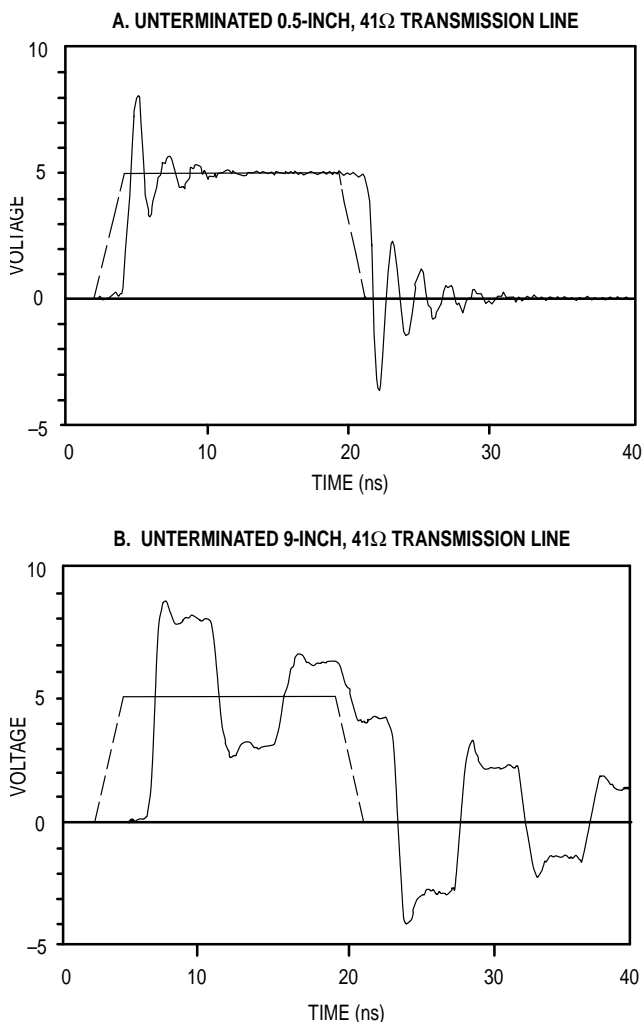
Figure 4. SPICE Simulation Results of 'Short' and 'Long' Transmission Lines. Simulations Were Run with Typical Parameters @ 25°C and V$_{CC}$ = 5.0V

Since clock lines are only driven from a single location, they lend themselves to termination more easily than bus lines which are commonly driven from multiple locations. Termination of bus lines with multiple drivers is a complicated manner which will not be addressed in this article. The most common types of termination in digital systems are shown in Figure 5. Since no single termination scheme is optimal in all cases, the tradeoffs involving the use of each will be discussed, and recommendations specific to clock drivers will

be made. Reference 2 is a comprehensive and practical treatment of transmission line theory and analysis of CMOS signals, and is recommended reading for those who want to gain a better understanding of transmission lines. Figure 6 shows SPICE simulated waveforms of the different termination schemes to be discussed. The driving device in the simulations was the MC88914 output buffer; in all simulations it drove a 9 inch 41Ω transmission line. The simulations were run using typical model parameters at 25°C and V$_{CC}$ = 5V.

Series termination, depicted in Figure 5b, is recommended if the load is lumped at the end of the trace and the output impedance of the driving device is less than the loaded characteristic impedance of the trace, or when a minimum number of components is required. The main problem with series termination occurs when the driving device has different output impedance values in the low and high states, which is a problem in TTL and some CMOS devices. A well designed CMOS clock driver should have nearly equal output impedances in the high and low states, avoiding this problem. An additional advantage is that series termination does not create a DC current path, thus the V$_{OL}$ and V$_{OH}$ levels are not degraded. The SPICE generated waveforms of series termination in Figure 6a show that series termination effectively masks the transmission line effects exhibited in Figure 4. If each clock output is driving only one device, series termination would be recommended, but this is not a realistic case in most systems, so series termination is not generally recommended for termination of clock lines.

Parallel termination utilizes a single resistor tied to ground or V$_{CC}$ whose value is equal to the characteristic impedance of the line. Its major disadvantage is the DC current path it creates when the driver is in the high state (if the resistor is tied to ground). This causes excessive power dissipation and V$_{OH}$ level degradation. Since a clock driver output is always switching, the DC current draw argument loses some credibility at higher frequencies because the AC switching current becomes a major component of the overall current. Therefore the main consideration in parallel termination is how much V$_{OH}$ degradation can be tolerated by the receiving devices. Figure 6b demonstrates that this termination technique is effective in minimizing the switching noise, but Thevenin termination has some advantages over parallel termination.

Thevenin termination utilizes one resistor tied to ground and a second tied to V$_{CC}$. An important consideration when using this type of termination is choosing the resistor values to avoid settling of the voltage between the high and low logic levels of the receiving device.2 TTL designers commonly use a 220/330 resistor value ratio, but CMOS is a little tricky because the switch point is at V$_{CC}$/2. With a 1:1 resistor ratio a failure at the driver output would cause the line to settle at 2.5V, causing system debug problems and also potential damage to the receiving devices.

In Thevenin termination, the parallel equivalent value of the two resistors should be equal to the characteristic impedance of the line. A DC path does exist in both the high and low states, but it is not as bad as parallel termination because the resistance in the Thevenin DC path is at least 2 times greater.

Figure 6c shows the termination waveforms, which exhibit characteristics similar to parallel termination, but with less V_{OH} degradation. The only real advantage of parallel over Thevenin is less resistors (1/2 as many) and less space taken up on the board by the resistors. If this is not a factor, Thevenin termination is recommended over parallel.

AC termination, shown in Figure 5e, normally utilizes a resistor and capacitor in series to ground. The capacitor blocks DC current flow, but allows the AC signal to flow to ground during switching. The RC time constant of the resistor and capacitor must be greater than twice the loaded line delay. AC termination is recommended because of its low power dissipation and also because of the availability of the resistor and capacitor in single-in-line packages (SIP). A pullup resistor to V_{CC} is sometimes added to set the DC level at a certain point because of the failure condition described in regards to Thevenin termination. As discussed earlier, the argument of lower DC current is less convincing at high frequencies. The AC terminated waveform walks out slightly toward the end of a high-to-low or low-to-high transition as seen in Figure 6d, making it slightly less desirable than Thevenin termination.

Thevenin and AC termination are the two recommended termination schemes for clock lines, but it depends on what frequency the clock is running at when making a decision between these types of termination. Although hard data is not provided to back this statement up, it is a safe assumption that at frequencies of 25 MHz and below AC is the best choice. If the system frequency could reach 40 MHz and beyond, Thevenin becomes the better choice.

ADDITIONAL CONSIDERATIONS WHEN TERMINATING CLOCK LINES

The results presented might imply that terminating the clock lines will completely solve noise problems, but termination can cause secondary problems with some logic devices. Termination acts to reduce the noise seen at the receiver, but that noise actually is seen as additional current and noise at the output of the driving device. If the internal and input logic on the source device is not sufficiently decoupled on chip from the high current outputs, internal threshold problems can occur. This phenomenon is commonly known as 'dynamic threshold.' It is usually evidenced by glitches appearing on the outputs of a fast, high current drive logic device as it switches high or low. This is most severe on 'ACT' devices which have high current and high slew rate CMOS outputs along with TTL inputs which have low noise immunity. This problem can be minimized by decoupling the internal ground and V_{CC} supplies on-chip and in the package. This decoupling is accomplished by having separate 'quiet' ground and V_{CC} pads on chip which supply the input circuitry's ground and V_{CC} references. These pads are then tied to extra 'quiet' ground and 'quiet' V_{CC} pins on the package, or to special 'split leads' which resemble a tuning fork and utilize the leadframe inductance to accomplish the decoupling. When choosing a clock source, make sure that the part has one of these decoupling schemes.

References

1. Blood, William R., *MECL System Design Handbook*, Motorola Inc., 1983.

2. Application Note AN1051, *Transmission Line Effects in PCB Applications*, Motorola Inc., 1990.

3. *Motorola FACT Data Book* DL138, Motorola Inc., 1990.

A. TRANSMISSION LINE WITH NO TERMINATION
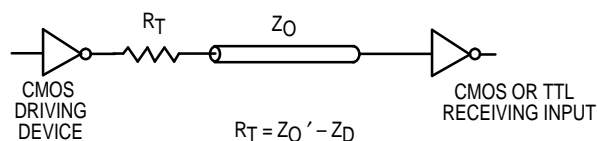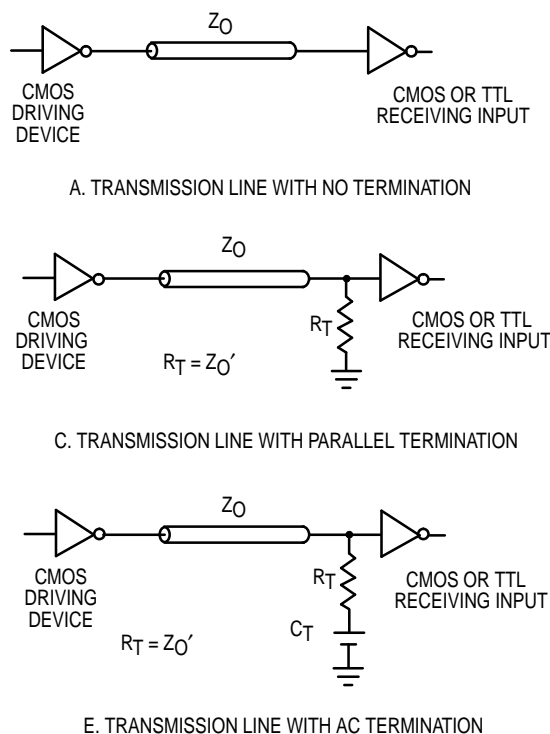
WHERE, Z_D = DRIVING DEVICE OUTPUT IMPEDANCE

B. TRANSMISSION LINE WITH SERIES TERMINATION

C. TRANSMISSION LINE WITH PARALLEL TERMINATION
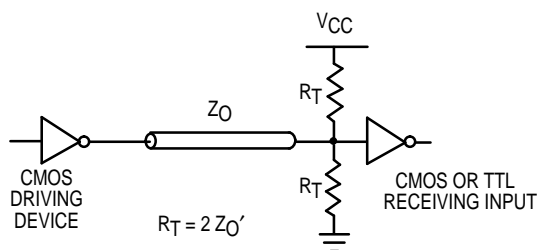
D. TRANSMISSION LINE WITH THEVENIN TERMINATION

E. TRANSMISSION LINE WITH AC TERMINATION

Figure 5. Schematic Representation of Common Termination Techniques

A. SERIES TERMINATION

B. PARALLEL TERMINATION

DRIVER OUTPUT WAVEFORM — — — — —
RECEIVER INPUT WAVEFORM ————————

C. THEVENIN TERMINATION
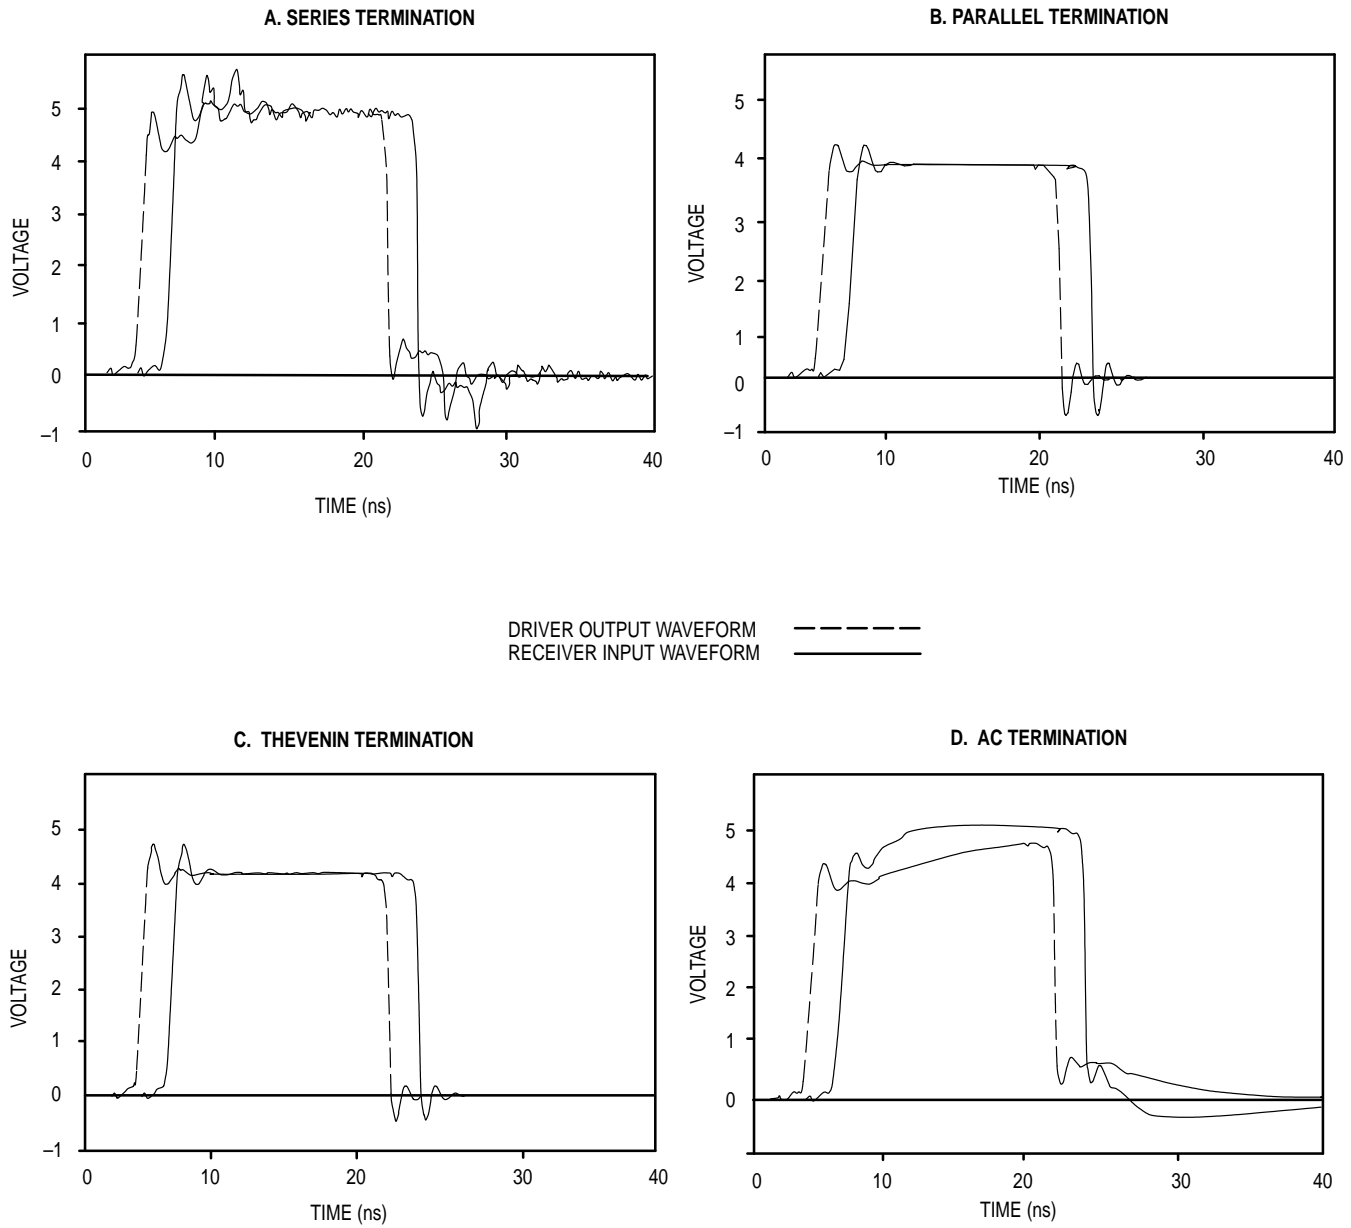
D. AC TERMINATION

Figure 6. SPICE Simulation Results for Various Terminations of a 9-Inch 41Ω Transmission Line. Simulations Were Run with Typical Model Parameters @ 25°C and V$_{CC}$ = 5.0V