

AN1405

ECL Clock Distribution Techniques

Prepared by: Todd Pearson
 ECL Applications Engineering

Abstract

This application note provides information on system design using ECL logic technologies for reducing system clock skew over the alternative CMOS and TTL technologies.

INTRODUCTION

The ever increasing performance requirements of today's systems has placed an even greater emphasis on the design of low skew clock generation and distribution networks. Clock skew, the difference in time between "simultaneous" clock transitions within a system, is a major component of the constraints which form the upper bound for the system clock frequency. Reductions in system clock skew allow designers to increase the performance of their designs without having to resort to more complicated architectures or more costly, faster logic. ECL logic technologies offer a number of advantages for reducing system clock skew over the alternative CMOS and TTL technologies.

SKEW DEFINITIONS

The skew introduced by logic devices can be divided into three parts: duty cycle skew, output-to-output skew and part-to-part skew. Depending on the specific application, each of the three components can be of equal or overriding importance.

Duty Cycle Skew

The duty cycle skew is a measure of the difference between the T_{PLH} and T_{PHL} propagation delays (Figure 1). Because differences in T_{PLH} and T_{PHL} will result in pulse width distortion the duty cycle skew is sometimes referred to as pulse skew. Duty cycle skew is important in applications where timing operations occur on both edges or when the duty cycle of the clock signal is critical. The later is a common requirement when driving the clock inputs of advanced microprocessors.




Figure 1. Duty Cycle Skew

Output-to-Output Skew

Output-to-output skew is defined as the difference between the propagation delays of all the outputs of a device. A key constraint on this measurement is the requirement that the output transitions are identical, therefore if the skew between all edges produced by a device is important the

output-to-output skew would need to be added to the duty cycle skew to get the total system skew. Typically the output-to-output skew will be smaller than the duty cycle skew for TTL and CMOS devices. Because of the near zero duty cycle skew of a differential ECL device the output-to-output skew will generally be larger. The output-to-output skew is important in systems where either a single device can provide all of the necessary clocks or for the first level device of a nested clock distribution tree. In these two situations the only parameter of importance will be the relative position of each output with respect to the other outputs on that die. Since these outputs will all see the same environmental and process conditions the skew will be significantly less than the propagation delay windows specified in the standard device data sheet.

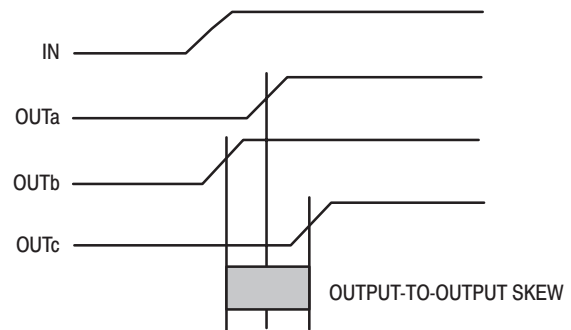


Figure 2. Output-to-Output Skew

Part-to-Part Skew

The part-to-part skew specification is by far the most difficult performance aspect of a device to minimize. Because the part-to-part skew is dependent on both process variations and variations in the environment the resultant specification is significantly larger than for the other two components of skew. Many times a vendor will provide subsets of part-to-part skew specifications based on non-varying environmental conditions. Care should be taken in reading data sheets to fully understand the conditions under which the specified limits are guaranteed. If the part-to-part skew is specified and is different than the specified propagation delay window for the device one can be assured there are constraints on the part-to-part skew specification.

Power supply and temperature variations are major contributors to variations in propagation delays of silicon devices. Constraints on these two parameters are commonly seen in part-to-part skew specifications. Although there are situations where the power supply variations could be ignored, it is difficult for this author to perceive of a realistic system whose devices are all under identical thermal conditions. Hot spots on boards or cabinets, interruption in air flow and variations in IC density of a board all lead to thermal gradients within a system. These thermal gradients will guarantee that

devices in various parts of the system are under different junction temperature conditions. Although it is unlikely that a designer will need the entire commercial temperature range, a portion of this range will need to be considered. Therefore, a part-to-part skew specified for a single temperature is of little use, especially if the temperature coefficient of the propagation delay is relatively large.

For designs whose clock distribution networks lie on a single board which utilizes power and ground planes an assumption of non-varying power supplies would be a valid assumption and a specification limit for a single power supply would be valuable. If, however, various pieces of the total distribution tree will be on different boards within a system there is a very real possibility that each device will see different power supply levels. In this case a specification limit for a fixed V_{CC} will be inadequate for the design of the system. Ideally the data sheets for clock distribution devices should include information which will allow designers to tailor the skew specifications of the device to their application environment.

SYSTEM ADVANTAGES OF ECL

Skew Reductions

ECL devices provide superior performance in all three areas of skew over their TTL or CMOS competitors. A skew reducing mechanism common to all skew parameters is the faster propagation delays of ECL devices. Since, to some extent, all skew represent a percentage of the typical delays faster delays will usually mean smaller skews. ECL devices, especially clock distribution devices, can be operated in either single-ended or differential modes. To minimize the skew of these devices the differential mode of operation should be used, however even in the single-ended mode the skew performance will be significantly better than for CMOS or TTL drivers.




Figure 3. V_{BB} Induced Duty Cycle Skew

ECL output buffers inherently show very little difference between T_{PLH} and T_{PHL} delays. What differences one does see are due mainly to switching reference levels which are not ideally centered in the input swing (see Figure 3). For worst case switching reference levels the pulse skew of an ECL device will still be less than 300 ps. If the ECL device is used differentially the variation in the switching reference will not impact the duty cycle skew as it is not used. In this case the pulse skew will be less than 50 ps and can generally be ignored in all but the highest performance designs. The problem of generating clocks which are capable of meeting the duty cycle requirements of the most advanced microprocessors, would be a trivial task if differential ECL compatible clock inputs were used. TTL and CMOS clock drivers on the other hand have

inherent differences between the T_{PLH} and T_{PHL} delays in addition to the problems with non-centered switching thresholds. In devices specifically designed to minimize this parameter it generally cannot be guaranteed to anything less than 1ns.

The major contributors to output-to-output skew is IC layout and package choice. Differences in internal paths and paths through the package generally can be minimized regardless of the silicon technology utilized at the die level, therefore ECL devices offer less of an advantage in this area than for other skew parameters. CMOS and TTL output performance is tied closely to the power supply levels and the stability of the power busses within the chip. Clock distribution trees by definition always switch simultaneously, thus creating significant disturbances on the internal power busses. To alleviate this problem multiple power and ground pins are utilized on TTL and CMOS clock distribution devices. However even with this strategy TTL and CMOS clock distribution devices are limited to 500 ps – 700 ps output-to-output skew guarantees. With differential ECL outputs very little if any noise is generated and coupled onto the internal power supplies. This coupled with the faster propagation delays of the output buffers produces output-to-output skews on ECL clock chips as low as 50 ps.

Two aspects of ECL clock devices will lead to significantly smaller part-to-part skews than their CMOS and TTL competitors: faster propagation delays and delay insensitivity to environmental variations. Variations in propagation delays with process are typically going to be based on a percentage of the typical delay of the device. Assuming this percentage is going to be approximately equivalent between ECL, TTL and CMOS processes, the faster the device the smaller the delay variations. Because state-of-the-art ECL devices are at least 5 times faster than TTL and CMOS devices, the expected delay variation would be one fifth those of CMOS and TTL devices without even considering environmental dependencies.

The propagation delays of an ECL device are insensitive to variations in power supply while CMOS and TTL device propagation delays vary significantly with changes in this parameter. Across temperature the percentage variation for all technologies is comparable, however, again the faster propagation delays of ECL will reduce the magnitude of the variation. Figure 4 on the following page represents normalized propagation delay versus temperature and power supply for the three technologies.

Low Impedance Line Driving

The clock requirements of today's systems necessitate an almost exclusive use of controlled impedance interconnect. In the past this requirement was unique to the performance levels associated with ECL technologies, and in fact precluded its use in all but the highest performance systems. However the high performance CMOS and TTL clock distribution chips now require care in the design and layout of PC boards to optimize their performance, with this criteria established the migration from these technologies to ECL is simplified. In fact, the difficulties involved in designing with these "slower" technologies in a controlled impedance environment may even enhance the potential of using ECL devices as they are ideally suited to the task.

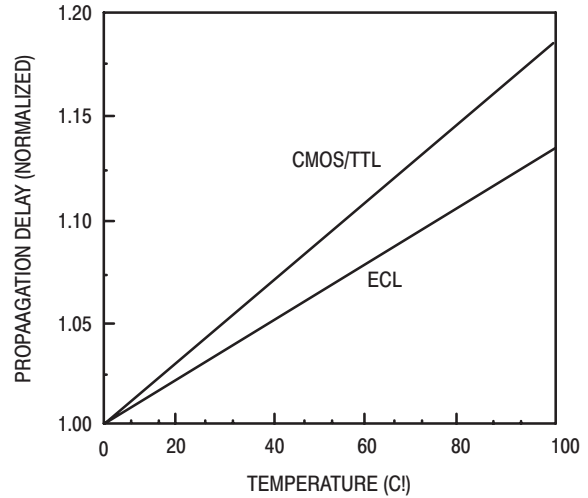
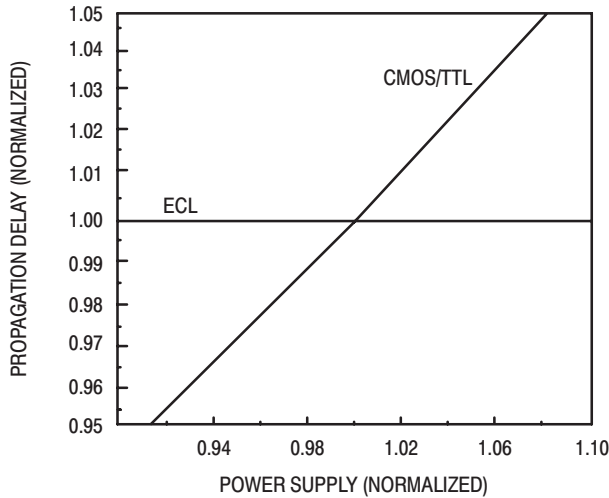


Figure 4. T_{PD} vs Environmental Condition Comparison

The low impedance outputs and high impedance inputs of an ECL device are ideal for driving 50 Ω to 130 Ω controlled impedance transmission lines. The specified driving impedance of ECL is 50 Ω , however this value is used only for convenience sake due to the 50 Ω impedance of most commonly used measurement equipment. Utilizing higher impedance lines will reduce the power dissipated by the termination resistors and thus should be considered in power sensitive designs. The major drawback of higher impedance lines (delays more dependent on capacitive loading) may not be an issue in the point to point interconnect scheme generally used in low skew clock distribution designs.

Differential Interconnect

The device skew minimization aspects of differential ECL have already been discussed however there are other system level advantages that should be mentioned. Whenever clock lines are distributed over long distances the losses in the line and the variations in power supply upset the ideal relationship between input voltages and switching thresholds. Because differential interconnect “carries” the switching threshold information from the source to the load the relationship between the two is less likely to be changed. In addition for long lines the smaller swings of an ECL device produce much lower levels of cross-talk between adjacent lines and minimizes EMI radiation from the PC board.

There is a cost associated with fully differential ECL, more pins for equivalent functions and more interconnect to be laid on a typically already crowded PC board. The first issue is really a non-issue for clock distribution devices. The output-to-output and duty cycle skew are very much dependent on quiet internal power supplies. Therefore the pins sacrificed for the complimentary outputs would otherwise have to be used as power supply pins, thus functionality is actually gained for an equivalent pin count as the inversion function is also available on a differential device. The presence of the inverted signal could be invaluable for a design which clocks both off the positive and negative edges. Figure 5 shows a method of obtaining very low skew (<50 ps) 180° shifted two phase clocks.

It is true that differential interconnect requires more signals to be routed on the PC board. Fortunately with the wide data and address buses of today’s designs the clock lines represent a small fraction of the total interconnect. The final choice as to whether or not to use differential interconnect lies in the level of skew performance necessary for the design. It should be noted that although single-ended ECL provides less attractive skew performance than differential ECL, it does provide significantly better performance than equivalent CMOS and TTL functions.




Figure 5. 180° Shifted Two Phase Clocks

USING ECL WITH POSITIVE SUPPLIES

It is hard to argue with the clock distribution advantages of ECL presented thus far, but it may be argued that except for all ECL designs it is too costly to include ECL devices in the distribution tree. This claim is based on the assumption that at least two extra power supplies are required; the negative V_{EE} supply and the negative V_{TT} termination voltage. Fortunately both these assumptions are false. PECL (Positive ECL) is an acronym which describes using ECL devices with a positive rather than negative power supply. It is important to understand

that all ECL devices are also PECL devices. By using ECL devices as PECL devices on a +5 volt supply and incorporating termination techniques which do not require a separate termination voltage (series termination, thevenin equivalent) ECL can be incorporated in a CMOS or TTL design with no added cost.

The reason for the choice of negative power supplies as standard for ECL is due to the fact that all of the output levels and internal switching bias levels are referenced to the V_{CC} rail. It is generally easier to keep the grounds quieter and equal potential throughout a system than it is with a power supply. Because the DC parameters are referenced to the V_{CC} rail any disturbances or voltage drops seen on V_{CC} will translate 1:1 to the output and internal reference levels. For this reason when communicating with PECL between two boards it is recommended that only differential interconnect be used. By using differential interconnect V_{CC} variations within the specified range will not in any way affect the performance of the device.

Finally mentioning ECL to a CMOS designer invariably conjures up visions of space heaters as their perception of ECL is high power. Although it is true that the static power of ECL is higher than for CMOS the dynamic power differences between the technologies narrows as the frequency increases. As can be seen in Figure 6 at frequencies as low as 20 MHz the per gate power of ECL is actually less than for CMOS. Since clock distribution devices are never static it does not make sense to compare the power dissipation of the two technologies in a static environment.

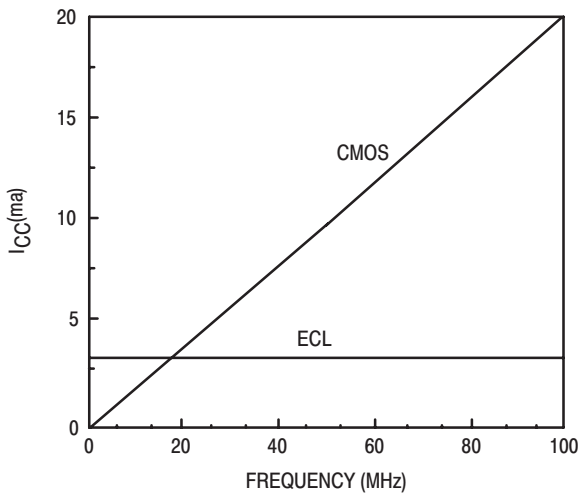


Figure 6. ICC/Gate vs Frequency Comparison

MIXED SIGNAL CLOCK DISTRIBUTION

ECL Clock Distribution Networks

Clock distribution in a ECL system is a relatively trivial matter. Figure 7 illustrates a two level clock distribution tree which produces nine differential ECL clocks on six different cards. The ECLinPS E211 device gives the flexibility of disabling each of the cards individually. In addition the synchronous registered enables will disable the device only when the clock is already

in the LOW state, thus avoiding the problem of generating runt pulses when an asynchronous disable is used. The device also provides a muxed clock input for incorporating a high speed system clock and a lower speed test or scan clock within the same distribution tree. The ECLinPS E111 device is used to receive the signals from the backplane and distribute it on the card. The worst case skew between all 54 clocks in this situation would be 275 ps assuming that all the loads and signal traces are equalized.




Figure 7. ECL Clock Distribution Tree

Mixed Technology Distribution Networks

Building clock networks in TTL and CMOS systems can be a little more complicated as there are more alternatives available. For simple one level distribution trees fanout devices like the MECL 10H645 1:9 TTL to TTL fanout tree can be used. However as the number of levels of fanout increases the addition of ECL devices in an other wise TTL or CMOS system becomes attractive. In Figure 8 on the next page an E111 device is combined with a MECL H641 device to produce 81 TTL level clocks. Analyzing the skew between the 81 clocks yields a worst case skew, allowing for the full temperature and V_{CC} range variation, of 1.25 ns. Under ideal situations, no variation in temperature or V_{CC} supply, the skew would be only 750 ps. When compared with distribution trees utilizing only TTL or CMOS technologies these numbers represent ~50% improvement, more if the environmental conditions vary to any degree. For a 50 MHz clock the total skew between the 81 TTL clocks is less than 6.5% of the clock period, thus providing the designer extra margin for layout induced skew to meet the overall skew budget of the design.

Many designers have already realized the benefits of ECL clock distribution trees and thus are implementing them in their designs. Furthermore where they have the capability, i.e. ASICs, they are building their VLSI circuits with ECL compatible clock inputs. Unfortunately other standard VLSI circuits such as microprocessors, microprocessor support chips and memory still cling to TTL or CMOS clock inputs. As a result many systems need both ECL and TTL clocks within the same system. Unlike the situation outlined in Figure 8 the ECL levels are not merely intermediate signals but rather are driving the

clock inputs of the logic. As a result the ECL edges need to be matched with the TTL edges as pictured in Figure 9.




Figure 8. ECL to TTL Clock Distribution




Figure 9. Mixed ECL and TTL Distribution

An ECL clock driver will be significantly faster than a TTL or CMOS equivalent function. Therefore to de-skew the ECL and TTL signals of Figure 9 a delay needs to be added to the input of the ECL device. Because a dynamic delay adjust would not lend itself to most production machines a static delay would be used. The value of the delay element would be a best guess estimate of the differences in the two propagation delays. It is highly unlikely that the temperature coefficients of the propagation delays of the ECL devices, TTL devices and delay devices would be equal. Although these problems will add skew to the system, the resultant total skew of the distribution network will be less than if no ECL chips were used.

PLL Based Clock Drivers

A potential solution for the problem outlined in Figure 9 is in the use of phase locked loop based clock distribution chips. Because these devices feedback an output and lock it to a reference clock input the delay differences between the various technology output buffers will be eliminated. One might believe that with all of the euphoria surrounding the performance of PLL based clock distribution devices that the need for any ECL in the distribution tree will be eliminated. However when analyzed further the opposite appears to be the case.

For a single board design with a one level distribution system there obviously is no need for ECL. When, however, a multiple board system is required where nested levels of devices are needed ECL once again becomes useful. One major aspect of part-to-part skew for PLL based clock chips often overlooked is the dependence on the skew of the various reference clocks being locked to. As can be seen in Figure 10 the specified part-to-part skew of the device would necessarily need to be added to the reference clock skew to get the overall skew of the clock tree. From the arguments presented earlier this skew will be minimized if the reference clock is distributed in ECL. It has not been shown as of yet where a PLL based ECL clock distribution chip can provide the skew performance of the simple fanout buffer. From a system standpoint the buffer type circuits are much easier to design with and thus given equivalent performance would represent the best alternative. The extra features provided by PLL based chips could all be realized if they were used in only the final stage of the distribution tree.

The MPC973 is a PLL based clock driver which features differential PECL reference clock inputs. When combined with the very low skew MC10E111 fanout buffer, very low skew clock trees can be realized for multiprocessor MPP designs. There will be a family of devices featuring various technology compatible inputs and outputs to allow for the building of precisely aligned clock trees based on either ECL, TTL, CMOS or differential GTL (or a mixture of all four) compatible levels.

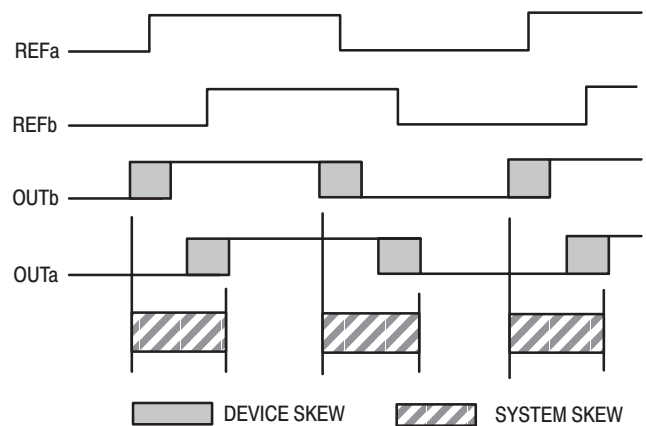



Figure 10. System Skew For PLL Clock Distribution

Conclusion

The best way to maximize the performance of any synchronous system is to spend the entire clock period performing value added operations. Obviously any portion of the clock period spent idle due to clock skew limits the potential performance of the system. Using ECL technology devices in clock distribution networks will minimize all aspects of skew and thus maximize the performance of a system. Unfortunately the VLSI world is not yet ECL clock based so that the benefits

of a totally ECL based distribution tree cannot be realized for many systems. However there are methods of incorporating ECL into the intermediate levels of the tree to significantly reduce the overall skew. In addition the system designers can utilize their new found knowledge to incorporate ECL compatible clocks on those VLSI chips of which they have control while at the same time pressuring other VLSI vendors in doing the same so that future designs can enjoy fully the advantages of distributing clocks with ECL.

NOTES

Motorola reserves the right to make changes without further notice to any products herein. Motorola makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Motorola assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Motorola data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Motorola does not convey any license under its patent rights nor the rights of others. Motorola products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Motorola product could create a situation where personal injury or death may occur. Should Buyer purchase or use Motorola products for any such unintended or unauthorized application, Buyer shall indemnify and hold Motorola and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Motorola was negligent regarding the design or manufacture of the part. Motorola and  are registered trademarks of Motorola, Inc. Motorola, Inc. is an Equal Opportunity/Affirmative Action Employer.

MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. © Motorola, Inc. 2001.

How to reach us:

USA/EUROPE/Locations Not Listed: Motorola Literature Distribution; P.O. Box 5405, Denver, Colorado 80217. 1-303-675-2140 or 1-800-441-2447

JAPAN: Motorola Japan Ltd.; SPS, Technical Information Center, 3-20-1, Minami-Azabu. Minato-ku, Tokyo 106-8573 Japan. 81-3-3440-3569

ASIA/PACIFIC: Motorola Semiconductors H.K. Ltd.; Silicon Harbour Centre, 2 Dai King Street, Tai Po Industrial Estate, Tai Po, N.T., Hong Kong. 852-26668334

Technical Information Center: 1-800-521-6274

HOME PAGE: <http://www.motorola.com/semiconductors/>



MOTOROLA



AN1405/D