SECTION 1
System Basics

Power Supply Considerations

The following text gives a brief description of the requirements and recommendations for treatment of power supplies in an ECLinPS* system design. A more thorough narration on the general subject of power supplies can be found in the Motorola System Design Handbook.

V_{CC} Supply

As with all previous ECL families the ECLinPS logic family is designed to operate with negative power supplies; in other words with V_{CC} connected to ground. However, ECLinPS circuits will work fine with positive power supplies as long as special care has been taken to ensure a stable, quiet V_{CC} supply. For more detailed information about using positive supplies and ECL, designers are encouraged to refer to Application Note AN1406 on page NO TAG. The output voltage levels for a positive supply system can be determined by simply subtracting the absolute value of the standard negative output levels from the desired V_{CC}.

To provide as small an AC impedance as possible, and minimize power bus IR drops, the V_{CC} supply should have a dedicated power plane. By providing a full ground plane in the system the designer ensures that the current return path for the signal traveling down a transmission line does not encounter any major obstructions. It is imperative that the noise and voltage drops be as small as possible on the V_{CC} plane as the internal switching references and output levels are all derived off of the V_{CC} power rail. Thus, any perturbations on this rail could adversely affect the noise margins of a system.

V_{EE} Supply

To take advantage of increased logic density and temperature compensated outputs, many designers are building array options with both, temperature compensated output levels and a – 5.2V V_{EE} supply. To alleviate any problems with interfacing these arrays to ECLinPS 100E devices, Motorola has specified the operation of 100E devices to include the standard 10H V_{EE} voltage range. Moreover, because of the superior voltage compensation of the bias network, this guarantee comes without any changes in the DC or AC specification limits. With the availability of both 10H and 100K compatible devices in the ECLinPS family, there is generally no need to run 10E devices at 100K voltage levels. If, however, this is desired, the 10E devices will function at 100E V_{EE} levels with, at worst, a small degradation in AC performance for a few devices due to soft saturation of the current source device.

Although both the 10E and 100E devices can tolerate variations in the V_{EE} supply without any detrimental effects, it is recommended that the V_{EE} supply also have a dedicated powerplane. If this is not a feasible constraint, care should be taken so that the IR drops of the V_{EE} bus do not create a V_{EE} voltage outside of the specification range. To provide the switching currents resulting from stray capacitances and asymmetric loading, the V_{EE} power supply in an ECL system needs to be bypassed. It is recommended that the V_{EE} supply be bypassed at each device with an RF quality 0.01µF capacitor to ground. In addition, the supply should also be bypassed to ground with a 1.0µF – 10µF capacitor at the power inputs to a board. If a separate output termination plane is used the V_{EE} supply will be of a static nature as the output switching current will return to ground via the V_{TT} supply, thus, the bypassing of every device may be on the conservative side. If the design is going to include a liberal use of serial or Thevenin equivalent termination schemes, a properly bypassed V_{EE} plane is essential.

V_{TT} Supply

The output edge rates of the ECLinPS family necessitate an almost exclusive use of controlled impedance transmission lines for system interconnect (the details of this claim will be discussed in a latter section). Thus, unless Thevenin equivalent termination schemes are going to be used, a V_{TT} supply is a must in ECLinPS designs. The choice of using only Thevenin equivalent termination schemes to save a power supply should not be made lightly as the Thevenin scheme consumes up to ten times more power than the equivalent parallel termination to a – 2.0V V_{TT} supply.

As was the case for the V_{EE} supply, a dedicated power plane, liberally bypassed as described above, should be used for the V_{TT} supply. In designs which rely heavily on parallel termination schemes the V_{TT} supply will be responsible for returning the switching current of the outputs to ground, therefore, a low AC impedance is a must. For bypassing, many SIP resistor packs have bypass capacitors integrated in their design to supply the necessary bypassing of the supply. The use of SIP resistors will be discussed more thoroughly in a later chapter.

Handling of Unused Inputs and Outputs

Unused Inputs

All ECLinPS devices have internal 50kΩ – 75kΩ pulldown resistors connected to V_{EE}. As a result, an input which is left open will be pulled to V_{EE} and, thus, set at a logic LOW. These internal pulldowns provide more than enough noise margin to keep the input from turning on if noise is coupled to the input, therefore, there is no need to tie the inputs to V_{EE} external to the package. In addition, by shorting the inputs to V_{EE} external to the package, one removes the current limiting effect of the pulldown resistor and, under extreme V_{EE}

*** Any reference to ECLinPS in this section include the ECLinPS Lite and Low Voltage ECLinPS families.**

conditions, the input transistor could be permanently damaged. If there are concerns about leaving sensitive inputs, such as clocks, open, they should be tied low via an unused output or a quiet connection to V_{TT}.

Unless otherwise noted on the data sheets, the outputs to differential input devices will go to a defined state when the inputs are left open. This is accomplished via an internal clamp. Note that this clamp will only take over if the voltage at the inputs fall below ≈ – 2.5V. Therefore, if equal voltages of greater than – 2.5V are placed on the inputs, the outputs will attain an undefined midswing state.

Unlike saturating logic families, the inputs to an ECLinPS, or any ECL device, cannot be tied directly to V_{CC} to implement a logic HIGH level. Tying inputs to V_{CC} will saturate the input transistor and the AC and DC performance will be seriously impaired. A logic HIGH on an ECLinPS input should be tied to a level no higher than 600mV below the V_{CC} rail and, more typically, no higher than the specified V_{IH}max limit. A resistor or diode tree can be used to generate a logic HIGH level or, more commonly, an output of an unused gate can be used.

Unused Outputs

The handling of unused outputs is guided by two criteria: power dissipation and noise generation. For single ended output devices it is highly recommended to leave unused outputs unterminated as there are no benefits in the alternative scheme. This not only saves the power associated with the output, but also reduces the noise on the V_{CC} line by reducing the current being switched through the inductance of the V_{CC} pins. For the counters and shift registers of the family, the count and shift frequencies will be maximized if the parallel outputs are left unterminated. Of course, for applications where these parallel outputs are needed this is not a viable alternative.

For the differential outputs, on the other hand, things are a little less cut and dry. If either of the outputs of a complimentary output pair is being used, both outputs of the pair should be terminated. This termination scheme minimizes the current being switched through the V_{CC} pin and, thus, minimizes the noise generated on V_{CC}. If, however, neither of the outputs of a complimentary pair are being used it makes most sense to leave these unterminated to save power. Note that the E111 device has special termination rules; these rules are outlined on the data sheet for the device.

Minimizing Simultaneous Switching Noise

A common occurrence among ECL families is the generation of crosstalk and other noise phenomena during simultaneous switching situations. Although the noise generated in ECL systems is minor compared to other technologies, there are methods to even further minimize the problem.

Figure 1.1 below illustrates the two output scenarios of an ECL device: differential outputs and single-ended outputs. During switching, the current in the output device will change by ≈17mA when loaded in the normal 50Ω to – 2.0V load. With differential outputs, as one output switches from a low to a high state the other switches from a high to a low state simultaneously, thus, the resultant current change through the V_{CCO} connection is zero. The current simply switches between the two outputs. However, for the single-ended output, the current change flows through the V_{CCO} connection of the output device. This current change through the V_{CCO} pin of the package causes a voltage spike due to the inductance of the pin.

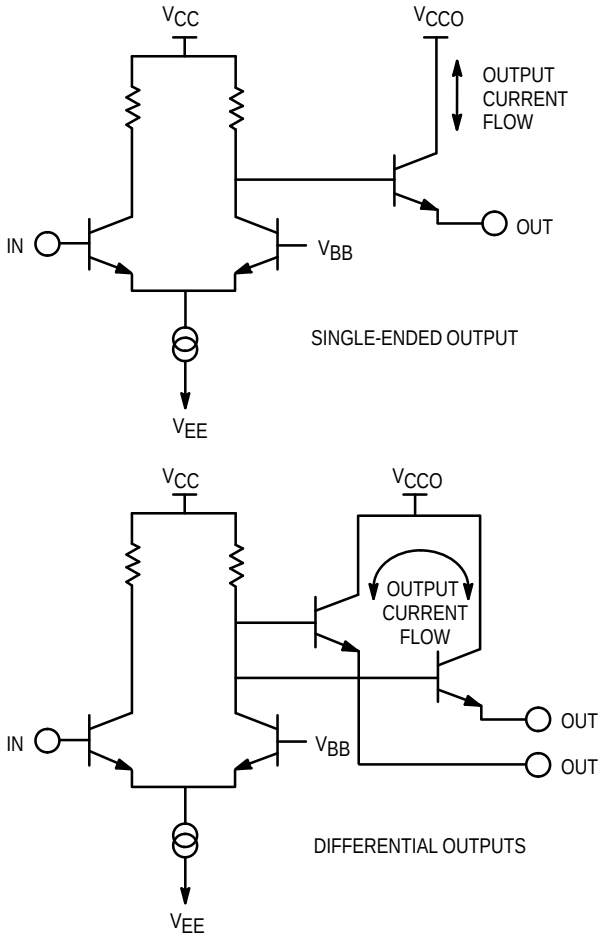

Figure 1.1. ECL Output Structures

Traditionally, manufacturers of ECL products have attempted to combat this problem by providing a separate V_{CC} pin for the output device (V_{CCO}, V_{CCA} etc.) and the internal circuitry. By doing this the noise generated on the V_{CCO} of the output devices would see a high impedance internal to the chip and not couple onto the the V_{CC} line which controls the output and internal bias levels. Unfortunately, in practice the noise generated on the V_{CCO} would couple into the chip V_{CC} through the collector base capacitance of the output device, thus, a large portion of the noise seen on the V_{CCO} line would also be seen on the V_{CC} line.

For the ECLinPS family and its associated edge speeds, it was decided that multiple V_{CCO} pins would be necessary to minimize the inductance and the associated noise generation. A design rule was established so that there would be no more than three single-ended outputs per V_{CCO} pin. Initially, the V_{CC} and V_{CCO} pins were kept isolated from

one another. However, it was discovered that in certain applications the parasitics of the package and the output device would combine to produce an instability which resulted in the outputs going into an oscillatory state. To alleviate this oscillation problem, it was necessary to make the V_{CC} and V_{CCO} metal common internal to the package. Subsequent evaluation showed that because of the liberal use of V_{CCO} pins, the noise generated is equal to or less than that of previous ECL families.

To further reduce the noise generated there are some things that can be done at the system level. As mentioned above, there should be adequate bypassing of the V_{CC} line and the guidelines for the handling of unused outputs should be followed. In addition, for wide single-ended output devices, an increase in the characteristic impedance of the transmission line interconnect will result in a smaller time rate of change of current; thus, reducing the voltage glitch caused by the inductance of the package. This noise improvement should, of course, be weighed against the potential slowing of the higher impedance trace to optimize the performance of the entire system. In addition, the connection between the device V_{CC} pins and the ground plane should be as small as possible to minimize the inductance of the V_{CC} line. Note that a device mounted in a socket will exhibit a larger amount of V_{CC} noise due to the added inductance of the socket pins.

Effects of Capacitive Loads

The issue of AC parametric shifts with load capacitance is a common concern especially with designers coming from the TTL and CMOS worlds. For ECLinPS type edge speeds, wire interconnect starts acting like transmission lines for lengths greater than 1/2″. Therefore, for the majority of cases in ECLinPS designs, the load on an output is seen by the transmission line and not the output of the driving device. The effects of load capacitance on transmission lines will be discussed in detail in the next section.

If the load is close to the driving output (<1/2″), the resulting degradation will be 15–25ps/pF for both propagation delays and edge rates. In general, a capacitive load on an emitter follower has a greater impact on the falling edge than the rising edge. Therefore, the upper end of the range given above represents the effect on fall times and the associated propagation delays, while the lower end represents the effect on the rising output parameters.

For ECLinPS devices, the capacitive load produced by an input ranges from 1.2pF to 2.0pF. The majority (≈95%) of this capacitance is contributed by the package with very little added by the internal input circuitry. For this reason the range is generally a result of the difference between a corner and a center pin for the PLCC package. A good typical capacitance value for a center pin is 1.4pf and for a corner pin 1.7pf. The capacitances for the other pins can be deduced through a linear interpolation.

Wired-OR Connections

The use of wired-OR connections in ECL designs is a popular way to reduce total part count and optimize the speed performance of a system. The limitation of OR-tying ECL outputs has always been a combination of increased delay per OR-tie and the negative going disturbance seen at the output when one output switches from a high to a low while the rest of the outputs remain high. For high speed devices the latter problem is the primary limitation due to the increased sensitivity to this phenomena with decreasing output transition times. The following paragraph will attempt to describe the wire-OR glitch phenomena from a physical perspective.

Figure 1.2 illustrates a typical wire-OR situation. For simplicity, the discussion will deal with only two outputs; however, the argument could easily be expanded to include any number of outputs. If both the A and the B outputs start in the high state they will both supply equal amounts of current to the load. If the B output then transitions from a high to a low the line at the emitter of B will see a sudden decrease in the line voltage. This negative going transition on the line will continue downward at the natural fall time of the output until the A output responds to the voltage change and supplies the needed current to the load. This lag in the time it takes for A to correct the load current and return the line to a quiescent high level is comprised of three elements: the natural response time of the A output, the delay associated with the trace length between the two outputs and the time it takes for a signal to propagate through the package. The trace delay can be effectively forced to zero by OR-tying adjacent pins. The resulting situation can then be considered “best case”. In this best case situation, if the delay through the package is not a significant portion of the transition time of the output, the resulting negative going glitch will be relatively small (≈100mV). A disturbance of this size will not propagate through a system. As the trace length between OR-tied outputs increases, the magnitude of the negative going disturbance will increase. Older ECL families specified the maximum delay allowed between OR-tied outputs to prevent the creation of a glitch which would propagate through a system.

As this glitch phenomena is a physical limitation, due to decreased edge rates, ECLinPS devices are susceptible to the problem to an even greater degree than previous slower ECL families. The package delay of even the 28-lead PLCC

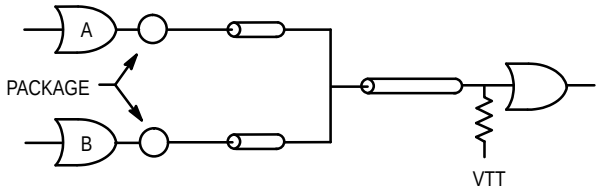

Figure 1.2. Typical Wire-OR Configuration

is a significant portion of the transition times for an ECLinPS device. Therefore, even in the best case situation described above, one can expect an ≈200mV glitch on the OR-tied line. A glitch of this magnitude will not propagate through the system but it is significantly worse than the best case situation of earlier ECL families. In fact, as long as the distance between OR-tied outputs is kept to less than 1/2″ the resulting line disturbance will not be sufficient to propagate through most systems.

With this in mind, the following recommendations are offered for OR-tying in ECLinPS designs. First, OR-tying of clock lines should be avoided as even in the best case situation the disturbance on the line is significant and could cause false clocking in some situations. In addition, wire ORed outputs should be from the same package and preferably should be adjacent pins. Non-adjacent outputs should be within 1/2″ of each other with the load resistor connection situated near the midpoint of the trace (Figure 1.2). By following these guidelines, the practice of wire-ORing ECL outputs can be expanded to the ECLinPS family without encountering problems in the system.

A detailed discussion of wire-OR connections in the ECLinPS world of performance is beyond the scope of this text. For this reason a separate application note has been written which deals with this situation in a much more thorough manner. Anyone planning to use wire-OR connections in their ECLinPS design is encouraged to contact a Motorola representative to obtain this application note.

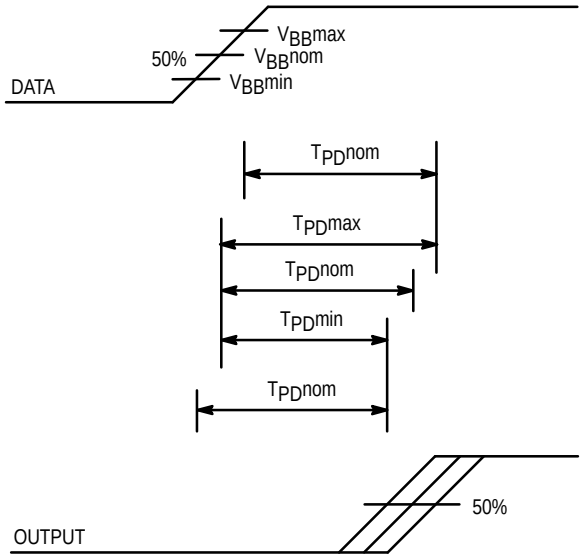

Figure 1.3. Delay vs Switching Reference Offset

Clock Distribution

Clock skew is a major contributor to the upper limit of operation of a high speed system; therefore, any reduction in this parameter will enhance the overall performance of a system. Through the ECLinPS family and new offerings in the 10H family, Motorola is providing devices uniquely designed to meet the demands of low skew clock distribution.

By far the largest contributor to system skew is the variation between different process lots of a given device. This variation is what defines the total delay window specified in the data sheets. This window can be minimized if the devices are fully differential due to the output level defined thresholds which ensure a “centered” input swing. The propagation delay windows of single-ended ECL and other logic technologies, are intimately tied to variations in the input thresholds. As illustrated in Figure 1.3 although the delays, when measured from the threshold of the input to the 50% point of the output, are equal; when measured from the specified 50% point of the input to the 50% point of the output, the delays will vary with any shift in the switching reference. Obviously, the magnitude of the delay difference is also proportional to the edge rate of the input. In addition to increasing the size of the delay windows, this reference shift will cause the duty cycle of the output of a device to be different than that of the input. Unfortunately, these thresholds are perhaps the most difficult aspects of a logic device to control. As a result, for the ultimate in low skew performance differential ECL devices are a must. A quick perusal of the ECLinPS databook will reveal a relatively large number of totally differential devices which will lend themselves nicely to very low skew applications such as clock distribution.

In addition to these generic differential devices there are several devices which were designed exclusively for clock distribution systems. With past ECL families designers were forced to build clock distribution trees with devices which were compromises at best. The ECLinPS family, however, was built around the E111 clock distribution chip; a fully differential 1:9 fanout device which boasts within part skews as well as part- to-part skews unequaled in today’s market.

Additionally, to further deskew clock lines the E195 programmable delay chip is available. Although static delay lines can remove built-in path length difference skew, they cannot compensate for variations in the delays of the devices in the clock path. The E195 allows the user to delay, a signal over a 2ns range in ≈20ps steps. Through the use of this device, the designer can match skews between clocks to 20ps.

Although these two devices satisfy the needs for many ECL designers, they do overlook the needs of a special subset; the designer who mixes ECL and TTL technologies. When translating between ECL and TTL, much of the skew performance gained through the E111 is lost when passed through the translator and distributed in TTL. To solve this problem, a new set of translators has been introduced in the MECL 10H family. The H641 and H643 receive a differential ECL input and fan out nine TTL outputs with a guaranteed unparalleled skew between the TTL outputs. The H640 and H642 take differential ECL inputs and generate low skew TTL outputs which are ideal for driving clocks in 68030 and 68040 microprocessor systems. By using the ECL aspects of the E111 to distribute clock lines across the backplane to TTL cards and receiving and translating these signals with the H640, H641, H642 or H643, a TTL clock distribution tree can be designed with a performance level unheard of with past logic families.

Through the development of a library of differential devices, specialized low-skew distribution chips and high-resolution programmable delay chips, Motorola has serviced the need for low-skew clock distribution designs. These offerings open the door for even higher performance next generation machines. For more information on clock distribution, designers are encouraged to read Application Note AN1405 on page NO TAG.

Metastability Behavior

The metastability behavior and measurement of a flip-flop is a complicated subject and necessitates much more time than is available in this forum for a thorough explanation. As a result, the following description is of an overview nature. Anyone interested in a more thorough narration on the subject is encouraged to read Application Note AN1504 on page NO TAG, which contains a more detailed discussion on the subject.

In many designs, occasions arise where an asynchronous signal needs to be synchronized to the system clock. Generally, this task is accomplished with the use of a single or series of D flip-flops as pictured in Figure 1.4. Because the data signal and the clock signal are asynchronous, the system designer cannot guarantee that the setup and hold specifications for the device will be met. This in and of itself would not cause a problem if it was not for the metastable behavior of a D flip-flop. The metastable behavior of a flip-flop is described by the outputs of a device attaining a nondefined logic level or, perhaps, going into an oscillatory state when the data and the clock inputs to the flip-flop switch simultaneously. It has been shown that this metastable behavior occurs across technology boundaries as well as across performance levels within a technology.

For ECL the characteristic of a flip-flop in a metastable state is a device whose outputs are in a non-defined state near the midpoint of a normal output swing. The output will return randomly to one of the two defined states some time later (Figure 1.5). The two parameters of importance when discussing metastability are the metastablity window; the window in time for which a transition on both the data and the

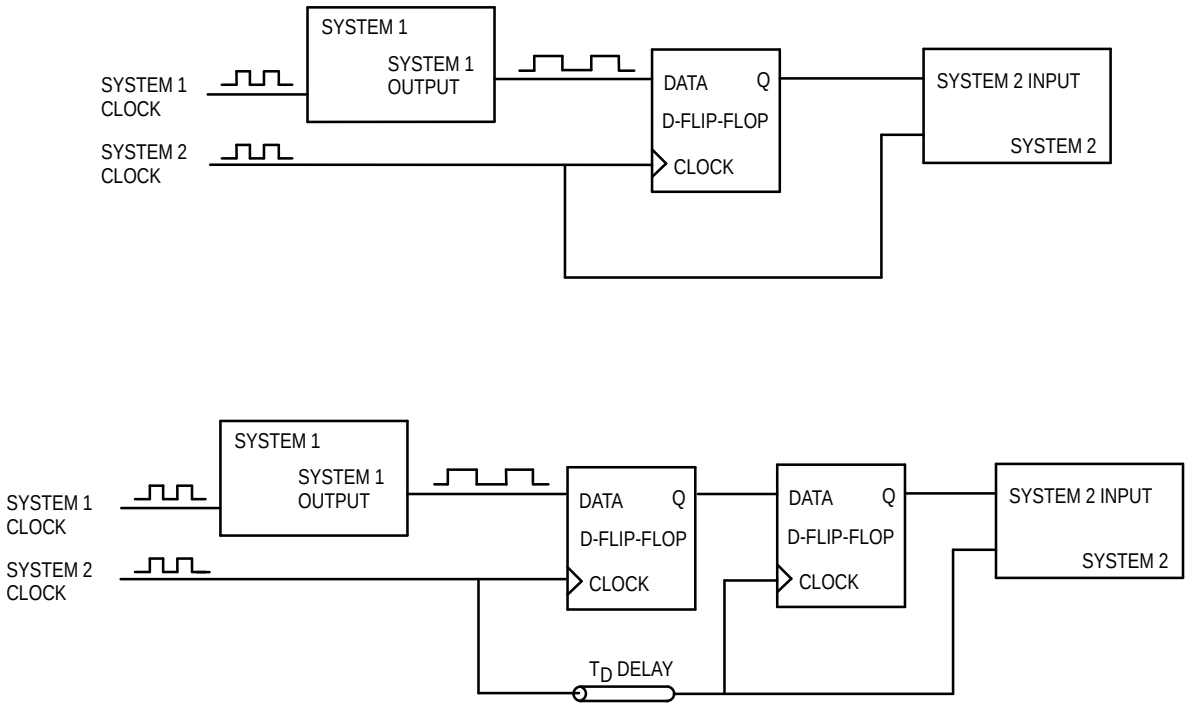

Figure 1.4. Clock Synchronization Schemes

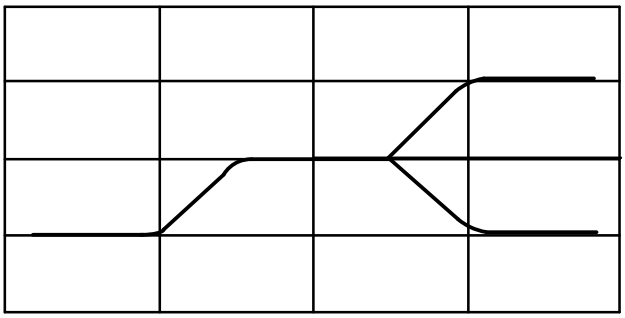

Figure 1.5. Metastable Behavior of an ECL Flip Flop

clock will cause a metastable output, and the settling time; the time it takes for a metastable output to return to a defined state. For the single flip-flop design of Figure 1.3, the data being fed into system 2 will be in an undefined state and, thus, unusable if the synchronizing flip-flop enters a metastable state. Because of this, a more popular design incorporates multiple flip-flop chains with cascading data inputs and clock inputs which are delayed with respect to each other. This redundancy of flip- flops helps to reduce the probability that the data entering system 2 will be at an undefined level which could wreak havoc on the logic of that system. This reduction in probability relies on the fact that even if the preceding flip-flop goes metastable, it will settle to a defined state prior to the clocking of the following flip-flop. Obviously, once the first flip-flop goes metastable there is an even chance that it will settle in the wrong state and, thus, information will be lost. However, there are error detection and correction methods to circumvent this problem. The larger the flip-flop chain the lower the probability of metastable data being fed into system 2.

Unfortunately for ECLinPS, levels of performance, both the window width and the settling time, are difficult or impossible to measure directly. The metastable window for an ECLinPS flip-flop is assuredly less than 5.0ps and most likely less than 1.0ps based on SPICE level simulation results. In either case, with today's measuring equipment, it would be impossible to measure this window width directly. Although it is feasible to measure the settling time for a given occurrence, this parameter is not fixed but, rather, is of a variable length which makes it impossible to provide an absolute guarantee.

The challenge then becomes, how to characterize metastability behavior given the above circumstances. The standard method in the industry is to use Stoll's[1] equation, combined with the standard MTBF equation, to develop the following relationship:

$$MTBF = 1 / (2 * f_C * f_D * T_P * 10^{-(t/\tau)})$$

where:

	f_C:	Clock Frequency
	f_D:	Data Frequency
	T_p:	FF Propagation Delay
	t:	Time Delay Between FF Clocks
	τ:	FF Resolution Time Constant

Note that the clock frequency, data frequency and time delay between flip-flops are user-defined parameters, thus it is up to Motorola to provide only the propagation delays and the resolution time constants for the ECLinPS flip-flops.

The propagation delays are, obviously, already defined leaving only the resolution time constant yet to be determined. An evaluation fixture was fabricated and several ECLinPS flip-flops were evaluated for resolution time constants.The results of the evaluation showed that the time constant was somewhat dependent on the part type as all the flip-flops in the ECLinPS family do not use the same general design. The time constants range from 125–225 ps depending on the part type.

As an example, for a system with a 100MHz clock and 75MHz data rate, the required delay between clock edges of a cascaded flip-flop chain for the E151 register, assuming a τ of 200ps, would be:

$$MTBF = 1 / (2 * 100MHz * 75MHz * 800ps * 10^{-t/200ps})$$

solving for an MTBF of 10 years yields:

t = 3.1ns, therefore:

$$T_D = \Delta t + T_P = 3.9ns$$

So, for an MTBF of 10 years for the above situation the second flip-flop should be clocked 3.9ns after the first. Similar results can be found by applying the equation to different data and clock rates as well as different acceptable MTBF rates.

[1] Stoll, P. "How to Avoid Synchronization Problems," VLSI Design, November/December 1982. pp. 56–59.