



Planning for High Speed XC9500XL Designs

XAPP115 September 28, 1998 (Version 1.0)

Application Note

Introduction

CPLD design has advanced significantly beyond that of fast PAL design. Today's CPLDs must operate in systems that include microprocessors, memories, I/O devices, buses, multiple power supplies and multiple frequency clocks. The actual logic design is frequently minor with respect to the electrical issues that must be dealt with during debug.

Discovering electrical problems at debug is too late. The printed circuit board has been built and may have to be significantly changed to debug. The best approach is to avoid the problem. By anticipating common problems, designs can be substantially "bullet-proofed" before debug. This means planning for options at the outset is the best solution. Like NASA space flights, a thorough, but practical checklist is one aspect of planning for success. This application note provides a framework for checklisting a design early, to eliminate problems.

Overview

To address the issues, they must first be identified. Noise, ground bounce, signal coupling, ringing and reflections rise quickly to the top of the list. CPLDs are frequently used in the heart of complex designs creating vital control signals like read/write strobes, bus and chip enables, etc. The CPLDs must provide signals to the rest of the system that are crisp, noise free and strong.

Signal Integrity Issues

Noise

Noise is created in a system at many places. The most common sources are power distribution or timing signal distribution (i.e., clocks). Power supply noise is dealt with in many ways, but capacitive decoupling is the most common. Recognizing that capacitors are not purely capacitive, but also include inductance and resistance means they have both low and high frequency responses. For unregulated boards (i.e. VCC brought in from outside), the standard procedure is to include low frequency decoupling capaci-

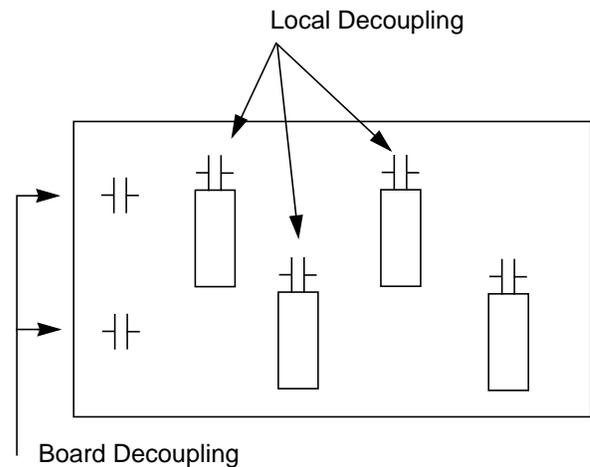


Figure 1: Board V_{CC} Decoupling

tors at the board power entry sites. (See Figure 1.) Then, each VCC pin of the CPLD should have high frequency capacitance attached physically close to the CPLD and the nearest ground. When discussing clock speeds near 200 MHz, it is vital that separate VCC and GND planes are included in the design.

Ground Bounce

Simultaneous switching signals can cause high current to be drawn into a chip, exiting at the chip ground leads. Due to inductance in the leads themselves, a voltage develops ($V=Ldi/dt$) that adds onto the chip ground path, briefly raising it. This is called bouncing. There is also a smaller effect related to the VCC pins on a chip. Figure 2 shows the standard ground bounce measurement technique, where all inputs save one transition, with the exception signal being driven low. Ground bounce is observed at the output of the statically driven gate, if it exists.

There are two ways to deal with this phenomenon: live with it or reduce it. If the rising ground does not cause signal integrity issues, it may be ignored. If it does create an issue, the signal must be reduced to the point where it doesn't.

Living With It

Signals jump around all the time. As long as the target signal does not disrupt other signals with its behavior, it may be ignored. If that signal is attached to a clock input of another function, or its asynchronous set/reset, it may cause problems. Likewise, if the signal is delivered to another flip flop's D input and settles before clock arrival, it may be ignored. However, if it causes setup time to be violated, it must be dealt with to assure stable operation.

Reducing It

Ground bounce can be reduced below the receiving gate's switching threshold. One way is to introduce some signal skew among the various offending gates. This can be done by setting some of the outputs to different slew rates than the rest (i.e., split up the simultaneous switching gates). Other ways exist to introduce time delay for some of the signals. Alternately, speeding up some paths may help reduce the total number of lines switching at once. With XC9500XL CPLDs, this may be handled by setting software optimization switches (design constraints).

Advanced recognition of sensitive signals is also useful. If it is known that a sensitive signal is prone to disastrous ground bounce response, that signal may be located very near (i.e., adjacent to) a physical ground. If this is not possible, leaving unused pins nearby will be useful for the Xilinx CPLD Fitter software to introduce user programmable ground pins. When programmed, these provide extra ground sites that work best when externally attached to the PCB ground.

Although not preferred, introducing additional impedance directly attached to the rising pin may be done. Attaching a pull-down resistor or a lag capacitor to the pin can reduce bounce, but may affect other signals that must switch that pin. Each case will require a different resistor or capacitor, depending on signal strength, pulse width, etc.

Signal Coupling

Coupling of electrical signals typically occurs on PCB traces where long lines run parallel for a significant length. The coupling mechanism is most likely edge to edge capacitance. The best way to manage this, is to prevent it. Examining the PCB signal file or the physical board can best tell this one. Carefully separating sensitive/critical signals for special handling may be needed. This is not specifically a CPLD issue.

Ringings (Over/Undershoot)

Ringings is ever present. With today's fast edge rates, short PCB traces (only a few inches) can create appreciable ringing. As with ground bounce, the choice of living with ringing or reducing it must be addressed. If the desired signal resolves to a clean binary value when sampled by other circuitry, it can be ignored. However, overshoot or undershoot

can affect other circuits, taking their signal levels "out of spec" with respect to power rails. This should be handled. One caveat is that the offending signal also is difficult to observe without introducing scope probe connection issues. After being certain that inappropriate ringing is indeed present, it should be decreased. The classic method is inserting series resistance.

Reflections

Reflections result when attaching a signal source to an interconnect with widely differing impedances. To reduce reflections, the standard approach is to insert compensation impedance. In general, this problem is tough to solve exactly, because the media has a constant impedance but the signal source has a variable one. It is important to remember that some signaling standards accept imbalance and rely on reflected wave switching for their basic signal approach (i.e., PCI).

Standard solutions to minimize reflection are based on termination, where series and parallel resistive termination are the most common. Other methods include distributed impedance and diode/resistance combinations. This document will focus on PCB terminations, with the stripline as the primary trace model. Standard termination practices are well known for the stripline model.

Expression 1 relates the length of a rising pulse edge with the distance it travels in a medium.

$$\text{Exp. 1: } L = Tr/D$$

where: L = length of rising edge,

Tr = rise time (10-90%), in picoseconds

D = delay in ps/in.(property of trace materials)

Tr is the single most telling parameter for how a signal will behave on a printed circuit board. A rule of thumb on whether reflections will occur is given by Expression 2:

$$\text{Exp. 2 1: } XL > Tr / 2Tpd$$

where: XL = length of trace

Tr = rise time of signal

Tpd = travel time from source to end (see Table 1)

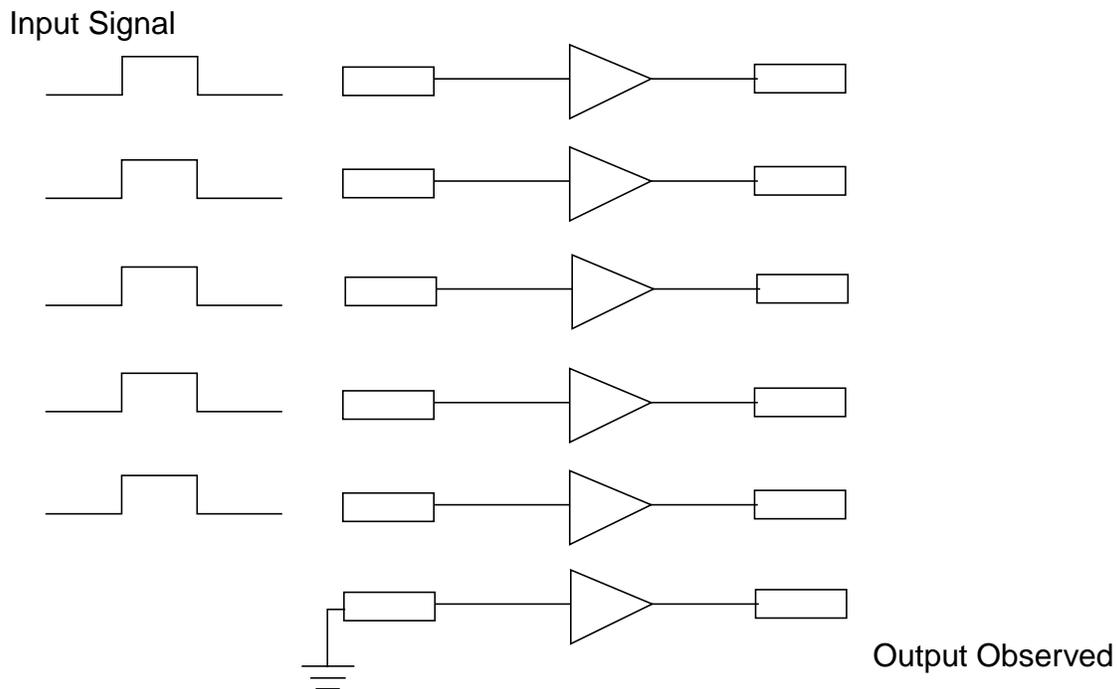


Figure 2: Ground Bounce Measurement

Table 1: Common Media with Delays and Dielectric

Medium	Delay (ps/in.)	Dielectric Constant
Air	85	1.0
Coax Cable (75% velocity)	113	1.8
Coax Cable (66% velocity)	129	2.3
FR4 PCB, outer trace	140-180	2.8-4.5
FR4 PCB, inner trace	180	4.5
Alumina PCB, inner trace	240-270	8-10

Constants

If Expression 2 is true, the signal will probably reflect, needing termination. More on this later.

This completes our basic list of signal integrity issues. In most cases, prevention can be accomplished by introducing external impedances or internal CPLD settings. To handle these problems, it is critical that the chosen CPLD supplies abundant features to manage the various problems. Xilinx XC9500XL CPLDs were designed to deliver the key properties needed for high signal integrity.

XC9500XL Electrical Capabilities

Let's look closer at the XC9500XL CPLD products. First, the XC9500XL family has CMOS output drivers. The supply structure splits an internal VCC (V_{CCINT}) and an I/O VCC (V_{CCIO}). This is convenient for interfacing with several signaling standards. Table 2 summarizes voltage family compliance for the XC9500XL family with respect to several EIA Standards.

Table 2: Voltage Family Compatibility of XC9500XL Family.

	5V CMOS	5V TTL	3.3V LVCMOS	3.3V LVTTTL	2.5V Normal
V _{IL}	√	√	√	√	√
V _{IH}	√	√	√	√	√
V _{OL}	√	√	√	√	√
V _{OH}	3.3V	3.3V	√	√	√

Quick Look at XC9500XL I/O Structure

Figure 3 gives a simplified look at the XC9500XL output structure. ESD protection is present with Figure 3, but is not shown. Because there is a P-channel pullup transistor attached to VCCIO, the output pin will be driven to full rail VCCIO, which will be either 3.3 or 2.5V.

When the output is driven high, the impedance to VCCIO will be approximately 20 ohms. When the output is driven low, the impedance to ground will be approximately 10 ohms. This gives some idea of range for inserting series resistance, to reduce ringing. Table 3 summarizes these values, and are derived from simulation. Note that both VCCIO conditions are distinguished in Table 3.

Table 3: XC9500 Impedance Switching Characteristics

Output Condition	Output R @ 3.3V	Output R @ 2.5V
Output High ¹	15 ohms	20 ohms
Output Low ¹	8 ohms	10 ohms

Note 1: Values from simulation

Because the resistance in the high and low direction are different, it will be impossible to match an external connection impedance to the impedance of an output pin. A trade-off must be made, if needed.

Of more importance, for termination, are the edge rates of the output switching signal. Note that we are typically interested in the largest value of an edge rate for worst case termination calculations. Edge rates for the XC9500XL are summarized in Table 4. Taking the rise time as 10-90% of the final value crosses a range of 2.64 volts (i.e., 2.97-0.33 = 2.64). Table 5 states the same information in terms of rise and fall times.

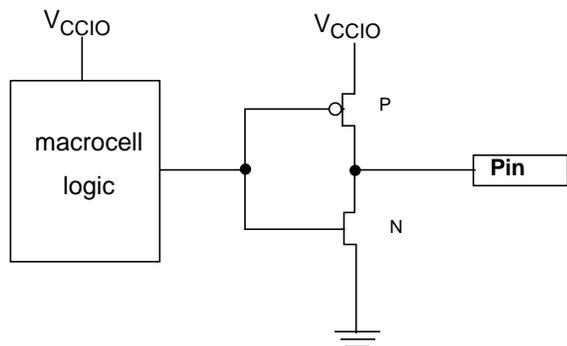


Figure 3: XC9500XL Output Structure

Table 4: Simulated XC9500XL Edge Rates

Slew Rate Setting	Worst Case High to Low	Worst Case Low to High
Slow	515 mV/ns	527mV/ns
Fast	882 mV/ns	846 mV/ns

Table 5: Simulated XC9500XL Fall/Rise Times (10% to 90%)

Slew Rate Setting	Worst Case Tf	Worst Case Tr
Slow	5.13 ns	5 ns
Fast	2.99 ns	3.12 ns

Worst case XC9500XL edge rates (from simulation) are 882 mV/nsec with fast slew rate and 527 mV/nsec with slow slew rate selected. Both edge rates are fast, but that is the price for high speed.

With that basic information, we can address the list of issues earlier defined.

Signal Integrity Solutions

Noise

Reducing noise is best done by adding capacitance. Making direct contact to the printed circuit board with minimum stub length is vital. All connections should be between the respective power planes. A good placement for decoupling capacitors is on the reverse side of the board from the component. Burying signals that do not need external exposure also cuts down on cross-talk and output switching noise in a CPLD. Although it has been common to just recommend putting 10 μF capacitors across VCC and GND points where power enters the board, 0.1μF caps at all CPLD VCCs and be done with it, things are a bit more complex, today. See references 1 and 5.

We should treat each decoupling/bypassing site independently, based on specific needs. We'll outline that process here. Both the board level decoupling and VCCINT bypassing will be calculated by similar considerations. VCCIO bypassing are found in a slightly different way. Remember, VCCINT and VCCIO are not necessarily the same values. Gardner and Johnson outline a five-step approach that makes practical sense.

Board Level Decoupling (C_{decouple})

In summary:

1. Estimate the largest ΔI the board will encounter.
2. Determine the largest ΔV the system is specified to tolerate (worst case VCC delta for any chips on the board set the limit)
3. Calculate $X_{max} = \Delta V / \Delta I$
4. Need to know power supply output series inductance (L_{psw})
5. $C_{decouple} = L_{psw} / (X_{max1})^2$ (i.e., greater than or equal to)

Typical digital board level decoupling capacitors range 10-1000 μF . Note that step 1 is typically estimated by summing the charge time for all other decoupling/bypassing capacitors on the board, simultaneously charging over a short period of time (say, 5 nsec.)

V_{CCINT} Bypassing ($C_{1bypass}$)

In this situation, we are noting that the internal (core) CPLD current is one component of the CPLD switching behavior. The output component is handled separately. Again, using the above approach:

1. Estimate the maximum change in current.

Basic approach: From the power estimation equation for an XC9500XL part (see Appendix), this is design dependent. Maximum current is based on maximum frequency of operation, number macrocells used and the split between high and low power macrocells. Minimum current is found by evaluating the equation near D.C. (say 100 Hz). Evenly distribute the current among the number of V_{CCINT} pins on the chosen chip. This will calculate a bypass value per V_{CCINT} pin
2. The voltage variation for the CPLD is specified between 3.6V down to 3.0V (i.e. specified range limit). This is 0.6V.
3. $X_{max2} = \Delta V / \Delta I$
4. Need the series inductance of the decoupling capacitor. This is supplied by the capacitor manufacturer and is usually called the lead inductance. For greater accuracy, the CPLD lead inductance and PCB stub inductance should be added to this value, which will be called L_C .
5. $C_{1bypass} = L_C / (X_{max2})^2$ (again, greater than or equal to)

V_{CCINT} decoupling capacitors typically range from 0.1 μF to 2 μF . Best results are found with low resistance, low inductance ceramic capacitors.

V_{CCIO} Bypassing ($C_{2bypass}$)

Here, the situation is slightly different. The capacitive load output edge rates come to play.

1. To arrive at a ΔI , we need to know how many loads will be simultaneously switching (worst case) and what their specific load capacitances are. Using Ohm's Law:

$I = Cx/dt$ we can arrive at I. C comes from the board PCB traces. $dV = 3.6V$ for 3.3V VCCIO (worst case) or $dV = 2.75 V$ VCCIO (worst case). $dt =$ the worst case edge rate (from Table 4). Again, the actual design comes into play, because knowledge of how many signals can possibly switch at once sets the limit. For XC9500XL parts, the design report file gives the number of outputs tied to a given clock, which will be the number of potentially switching outputs for that clock. Be sure to tally all output current components into a total.

2. Assume that you can tolerate ΔV volts from the VCC and remain within specification (here, you must know the quality of your power supply regulation). The VCCIO decoupling capacitor is then taken as:

$C_{2bypass} = I \Delta t / \Delta V$ where Δt is the worst case output edge rate.

Typical values for $C_{2bypass}$ typically range from 0.1 to 10 μF . Again, low resistance and low inductance ceramic capacitors provide best results.

Ground Bounce

Ground bounce can first be assessed by examining the design report file and determining how many signals may potentially switch at the same time. It is hard to know how many inputs might simultaneously switch, but flip flops attached to a common clock line are likely internal signals to simultaneously switch. If more than 50 signals attach to the same clock, adding user programmable grounds (a software switch) is a good idea. The corresponding pins can deliver better grounding if they are also tied externally to the PCB ground plane.

Coupling

Coupling is not an XC9500XL issue, but rather a property of the printed circuit board traces. Visual inspection of a plot and comparing to a design database helps determine which signals track each other (i.e., run parallel). One trick here is to run appropriate tracking signals on the opposite sides of the board, isolated by the VCC and GND planes.

Ringing

Ringling may be present on even short traces. Reserving a series resistance location for values between 20 and 50 ohms is probably a good idea. That way, if ringing is a problem, the series impedance can be adjusted to lower the effect.

Reflections

As suggested earlier, reflections can be planned for. The two classic termination types are series (source) termination and parallel (endpoint) termination. To be effective, source termination must be physically located near the source. Howard Johnson suggests that the resistor be

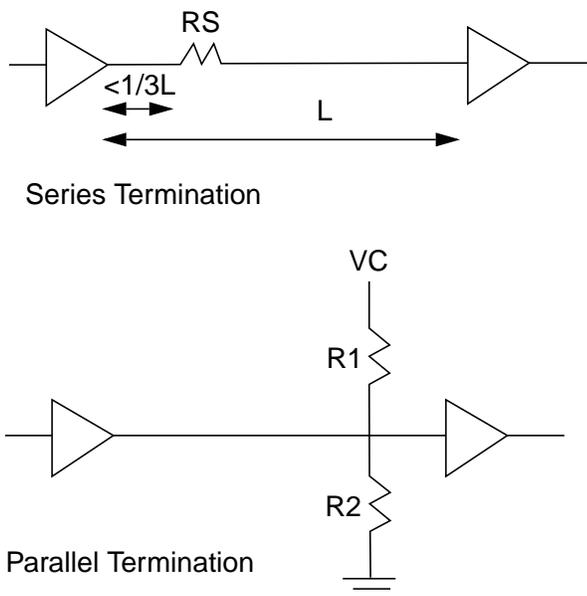


Figure 4: Single Ended Series/Parallel Termination

stubbed away from the source no more than one-third the trace length. Similarly, it is vital that parallel resistance is located physically near the destination. Figure 4 shows the situation for both styles of termination.

For series termination, it is important to know Z_O of the line and the source impedance. Then select a series resistance value so that $R_{source} + R_{series} = Z_O$. As noted in the figure, R_{series} should be located near the driver for greatest effect. Locating R_{series} further than $1/3$ the length of the trace from the driver is ineffective.

For the parallel termination case, it is important to select the parallel combination of $R1$ and $R2$ to have a combined resistance equal in magnitude to Z_O .

Figure 5 shows a common case where two mutually exclusive bidirectional drivers can drive the same line. In the case, R_{s1} plus the left driver source impedance to sum to equal Z_O for the line. R_{s2} and the right driver source impedance should also sum to match Z_O of the line. Another version of bidirectional termination is similar to that shown as parallel termination in Figure 4, but on both ends. In that case, the resistor sets must deliver a parallel match to Z_O .

Series termination requires fewer components, dissipates less power, but typically delays signal arrival by a predict-

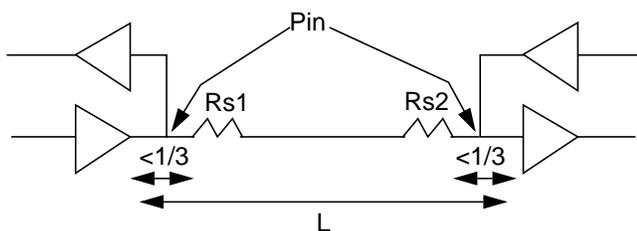


Figure 5: Bidirectional Series Termination

able RC time constant. Parallel termination is faster, requires more components and is always drawing power because the top transistor is attached to the VCC level of the switching drivers.

Table 6 summarizes impedance parameters for common media. Note that the values are provided indicate that the Z_O may vary depending on type and manufacture.

Table 6: Common Characteristic Impedances and Incremental Time Delays

Type	Z_O Ohms	T_O ns/inch
Coax	50-125	0.13
Wire over Ground	70-170	0.14
Microstrip Lines	30-150	0.15
Stripline	15-100	0.19
PC Board Traces	50-200	0.16

Printed Circuit Board Layout and Debug Checklists

Do these things before or during first PCB Layout:

1. Always use VCC/GND planes
2. Decouple and bypass all VCC (internal and I/O) pins to the nearest ground site
3. Minimize long runs
4. Minimize trace tracking (i.e., parallel signal traces).

Do these things after first PCB layout, but prior to final first PCB file creation:

1. Insert line termination where needed. Decision should be based on level of mismatch between PCB Z_O and Z_{source} .
2. Series is recommended for output only signals.
3. Bidirectional lines should have series at each end, located near the source to account for impedance mismatch and minimize over/undershoot.

Debug Checklist

This list is a set of things to do after PCB Layout if signals still need work.

Software adjustments:

1. eliminate unneeded outputs (signals that may remain buried).
2. place noncritical signals in slow slew rate mode
3. assign unused pins as User Programmable Ground
4. assign low power to macrocells not requiring high speed
5. introduce skew among large groups of simultaneously switching signals

Board adjustments:

6. modify termination
7. reduce VCCIO (if possible)

Final Comments

This application note covers most of the practical considerations that can slow down debug. Advanced planning is clearly the most cost effective and time effective approach for best results. Just as a good “pinlocking” architecture can minimize PCB changes for your logic, advanced signal planning can minimize PCB changes for electrical reasons. Minimizing the number of PCB changes to complete a design on schedule is the goal, and one of the most impressive features of programmable logic.

References

1. High Speed Digital Design A Handbook of Black Magic, Prentice-Hall, 1993, H.W.Johnson, M.Graham
2. Effective Use of Line Termination in High Speed Logic
Stanley Hronik, IDT
3. How close is close enough? H.W. Johnson, EDN, pg 24, 4/9/98
4. Bypass Capacitor Selection for High Speed Designs, TN-00-06, 1998, Micron Technology, Inc.

Appendix - Basic Core Current Estimation

Power estimation for CPLDs is complex. A formula that can be used to give a ballpark estimate for XC9500XL parts is given as:

$$P_{TOTAL} = P_{INT} + P_{IO} = I_{CCINT} \times V_{CCINT} + P_{IO}$$

Separating internal and I/O power here is convenient because XC9500XL CPLDs also separate the corresponding power pins. P_{IO} is a strong function of the load capacitance driven, so it is handled by $I = CdV/dt$. I_{CCINT} is another situation, that reflects the actual design considered and the internal switching speeds. An estimation expression for I_{CCINT} (taken from simulation) is:

$$I_{CCINT}(mA) = MC_{HP}(0.175 \times PT_{HP} + 0.345) + MC_{LP}(0.052 \times PT_{LP} + 0.272) + 0.04(MC_{HP} + MC_{LP}) \times f_{MAX} \times MC_{TOG}$$

where:

- MC_{HP} = # macrocells used in high speed mode
- MC_{LP} = #macrocells used in low power mode
- PT_{HP} = average p-terms used per high speed macrocell
- PT_{LP} = average p-terms used over low power macrocell
- f_{MAX} = max clocking frequency in the device
- MC_{TOG} = % macrocells toggling on each clock

To get a credible estimate, the macrocell and product term counts should be extracted from the design file report (.rpt) and evaluated at f_{MAX} to get the high value. A minimum value is obtained by setting f_{MAX} to zero, which gives an estimate of standby I_{CCINT}. The difference is what can change (i.e., ΔI) and what must be accounted for in the bypass equations.