

BOARD LEVEL SIMULATION SPECIALISTS

ICD Stackup Planner - offers engineers/PCB designers unprecedented simulation speed, ease of use and accuracy at an affordable price

- 2D (BEM) field solver precision
- Characteristic impedance, edge-coupled & broadside-coupled differential impedance
- Unique field solver computation of multiple differential technologies per stackup
- Heads-up impedance plots of signal and dielectric layers
- User defined dielectric materials library - over 16,250 materials up to 100GHz

ICD PDN Planner - analyze multiple power supplies to maintain low impedance over entire frequency range dramatically improving product performance

- Fast AC impedance analysis with plane resonance
- Definition of plane size/shape, dielectric constant & plane separation for each on-board power supply
- Extraction of plane data from the integrated Stackup Planner
- Definition of voltage regulator, bypass/decoupling capacitors, mounting loop inductance
- Frequency range up to 100GHz
- Extensive Capacitor Library – over 5,250 capacitors derived from SPICE models

Critical Placement

by Barry Olney

IN-CIRCUIT DESIGN PTY LTD

SUMMARY: *Controlling the placement of devices limits maximum trace length, reduces flight time delay and skew, and assists in compliance to timing specifications.*

Today's high-speed digital products use high frequencies and fast rise times that demand careful attention to PCB layout to ensure that the system works not only in the prototype phase, but also in mass production, and under all possible operating conditions.

An insidious problem with high-speed boards, in fact, is that their failure mode may simply be flakiness or intermittent behavior across multiple manufacturing runs. In these cases, the proper layout of the PCB can mean the difference between a reliable product and a board that may perform intermittently. A single PCB design error can bring down part or all of a board. Given all the possible mistakes that can slip through even the most dutiful design process, it is a wonder that we ever get the design to work properly. Maybe some of the advice that follows will help ensure you get it right first time.

Mechanical Constraints

Interactive mechanical and electronics co-design of PCBs within a true concurrent flow is becoming common these days, as inevitably every board has to fit into a box or enclosure of some sort.

As the first step, the mechanical designer uses the MCAD system to define the shape of the PCB in the mechanical product in "three space," setting mechanical con-

straints and placing components such as connectors that have 3D positioning requirements. After completion of the mechanical product design, the mechanical designer extracts the PCB shape and height restrictions from the design, and exports them to the PCB database. Once the PCB designer has completed component placement, the layout can then be transferred back to the MCAD system for a final 3D clearance check.

Electrical isolation and electrostatic requirements may also have to be adhered to. These may include cut-outs, keep-outs, vertical clearances, mounting holes and thermal heatsinks that need to be defined.


Thermal Analysis

Power dissipation is an important issue in present-day PCB design. To the degree that smaller dies enable smaller packages, this can result in higher concentration of heat.

In addition to the issue of reliability, excess heat will also negatively affect electrical performance. The working temperature of an IC should be kept below the maximum allowable, worst-case limit. In general, the temperatures of junction and ambient are 125°C and 55°C, respectively. Consideration should be given to thermal conduction and convection in the operating environment, as well as radiation.

Flight Time and Skew

Flight time delay and skew are key pillars in high-speed PCB design. One of the driving factors for flight time and skew performance






Figure 1: DDR2 processor and memory placement requirements.

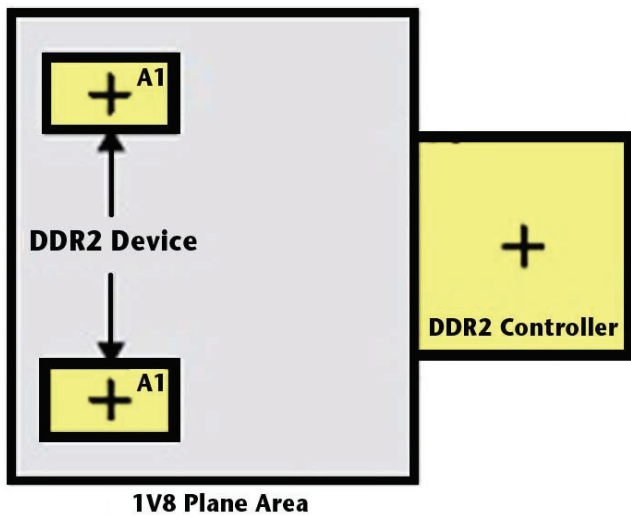


Figure 2: 1.8V plane and keep-out area.

is the placement of components. Maximum placement refers to the placement in which the distances between the devices are the maximum distance permitted. Controlling the maximum placement of devices, combined with the assumption that good general design practices are adhered to, limits maximum trace delay to roughly the longest Manhattan distance of the signals contained in a specific clock domain. Why the longest Manhattan distance?

This is due to skew matching requirements:

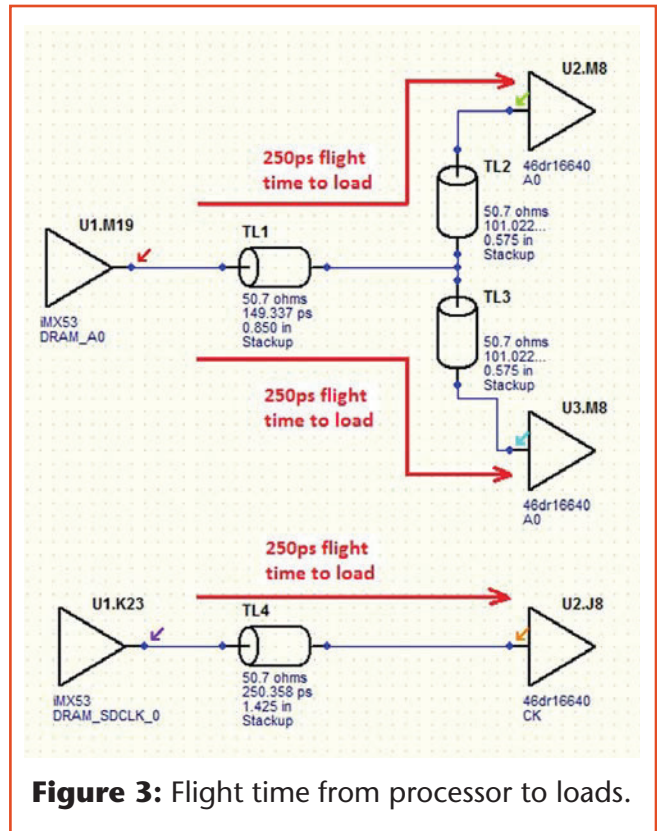


Figure 3: Flight time from processor to loads.

All of the shorter nets in a clock domain must be lengthened to skew match to the longest run length. Therefore, flight time and skew—for an entire clock domain—are governed by the maximum placement, along with the routing rules that constrain the matching of the trace lengths.

In the classic high-speed design flow, timing specifications simulation results are compared to determine placement and routing constraints. Given a length constraint, a designer can control signal integrity by controlling the PCB trace topology of the various parts of an interface. Included in this topology are any terminations.

Figure 1 shows the required placement of the processor and DDR2 memory chips. The purpose of the placement guide is to limit the maximum trace lengths and to allow for routing and via space—which can be a challenge. (Note that this placement does not restrict whether these devices are placed on the top or bottom of the board.)

The region of the PCB used for the DDR2 circuitry must be isolated from other signals.

The DDR2 keep-out region is defined for this purpose and is shown in Figure 2. The 1.8V power plane should cover this entire region and non-DDR2 signals should be kept out of this region.

The main requirement for DDRx timing is that the clock must be slightly longer than the address, control and command traces, and that the strobe must be longer than the data and data masks. When looking at the total length of signals in the EDA tool, one can easily be mistaken in calculating the total distance and hence delay required—totally annihilating the timing.

Figure 3 illustrates the relationship between an address signal (A0) routed in a T-section topology to the address pins on the two DDR2 loads. The total length in the EDA tool is given as $0.850 + 0.575 + 0.575 = 2000$ mils. So, one would assume that the clock must be at least 2000 mils to be longer than the address trace A0. This, of course, is incorrect, as we only need to worry about the longest point-to-point length.

If we now look at the flight time from the iMX53 processor to A0 pin of the first memory chip, the total flight time to the pin is 250ps, which is identical to the total flight time from the processor to the clock pin of each memory chip. So, the total length should be $0.850 + 0.575 = 1425$ mils.

The same applies for the data to strobe—which acts as the clock for the data signals. The strobe must be longer than any of the data or data mask signals for each lane.

Figure 4 illustrates the relationship between the clock and address signals. If a designer was to inadvertently use 2000 mils for the clock

length then the clock would be delayed by an additional 100ps, putting it way out of spec and causing a system-timing failure.

Having simulated the pre-layout timing, we can now put values into the placement of the DDR2 memory chips with respect to the processor. For Figure 1, $X = 850$ mils and Y offset = 575 mils. These are the maximum values and should be reduced to allow for the length tuning of traces. It is best to reduce the X dimension to, say, 650 mils and keep the Y offset as large as possible to provide space for vias at the centre of the T section.

Termination

Placement of series terminators—if required—can also a debacle. We are told that series terminators should always be placed within 200 mils of the source. Since there is no termination at the far end of the trace, the series terminator has to absorb the entire reflection from the load.

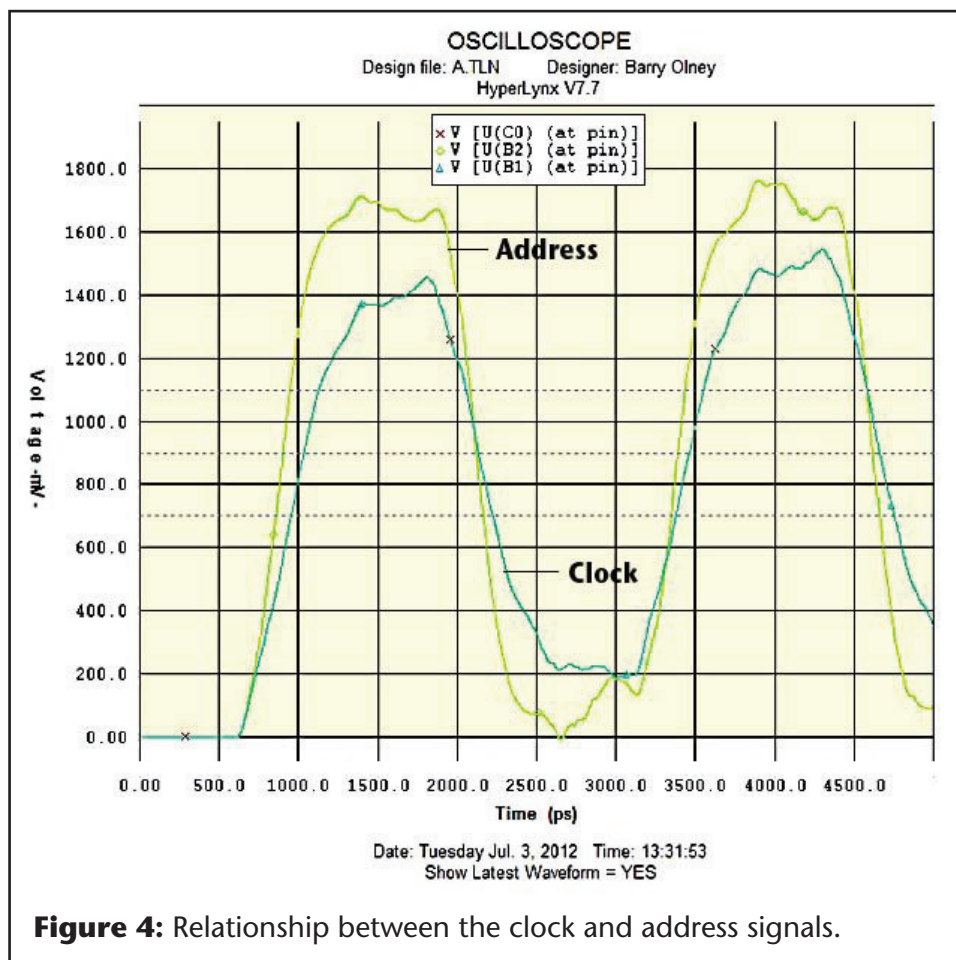


Figure 4: Relationship between the clock and address signals.

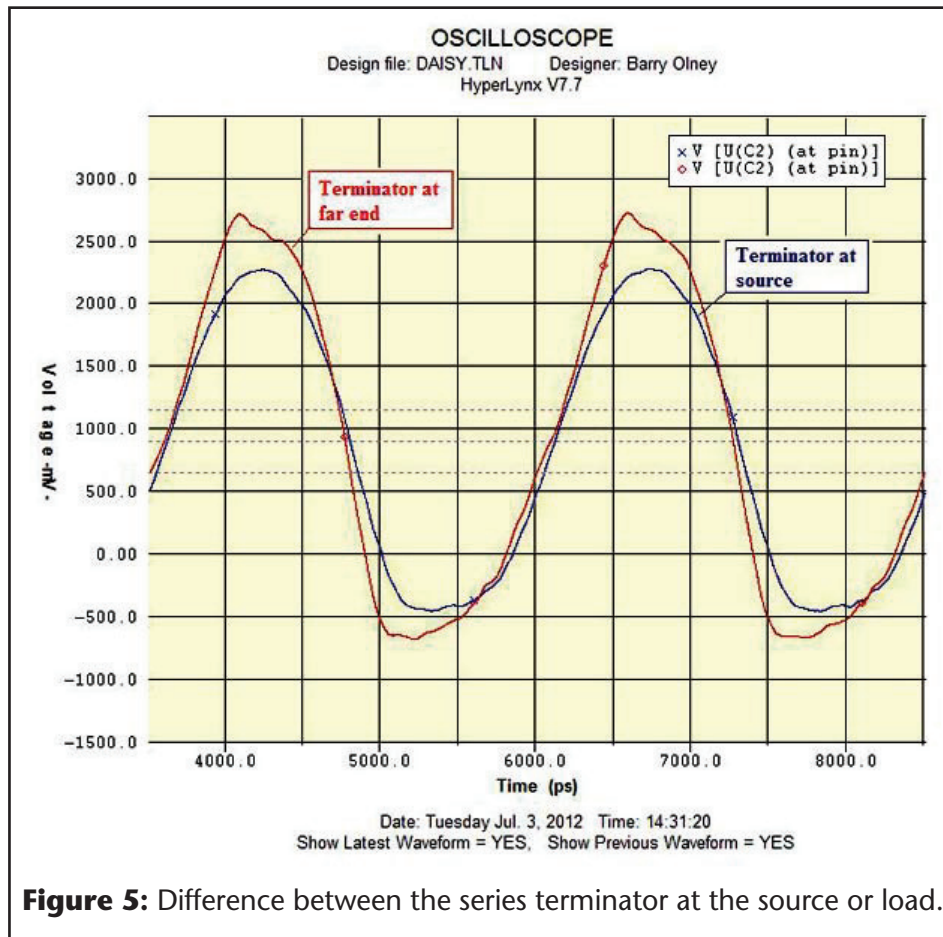


Figure 5: Difference between the series terminator at the source or load.

Figure 5 contrasts the placement of a series terminator close to the source—as typically done—and at the far end near the load—which is not recommended. It seems that either way the series terminator does its job—absorbing the reflection before it can propagate along the trace a second or third time. The only noticeable effect is that the far-end termination is more distorted whereas the near-end termination is a perfectly clean waveform.

There may also be a requirement to use parallel termination—address VTT pull-up. These are normally placed on the end of a daisy chain to pull the signal up to VTT which is the reference voltage for DDR2. This has to be at the end of the line.

Design Re-use

If the same processor and memory chips are used on consecutive designs, creating a library of matching re-use blocks for schematic and PCB makes the best use of existing design elements

for future designs. Simply add a sub-circuit block to the schematic, transfer to the PCB database, and load a predefined layout block including component placement, tracks, copper and text. Whether you use it for multiple channel designs, critical digital circuitry, RF circuit blocks, or just to replicate a commonly used layout pattern, design reuse will save time and ensure repeatability of design: a proven, tested, working solution to just drop into place.

Decoupling Capacitor Planning and Placement

The number, value, dielectric material and placement of decoupling capacitors are also critical for the proper DDR2 interface operation.

Manufacturer’s recommendations for decoupling capacitors are generally overrated in that they tend to overestimate the number of decaps required. This may be fine for short production batches, but a detailed power distribution network analysis is required to determine the exact number and value for mass production in order to reduce costs and optimize the PDN.

Figure 6 depicts the ICD PDN Planner analysis of the DDR2 1.8V supply (download from www.icd.com.au). The 30 x 100nF capacitors should be placed as close as possible to the DDR2 chip supply pins using thick (20 mils) traces to reduce inductance. A good capacitor dielectric choice is X7R with a footprint of 0402 or smaller. 4 x 22 uF Tantalums reduce the AC impedance around 10M, while the 100nF capacitors and tight interplane coupling (0.004” in this case) keep the impedance as low as possible at the fundamental clock frequency of 400 MHz.

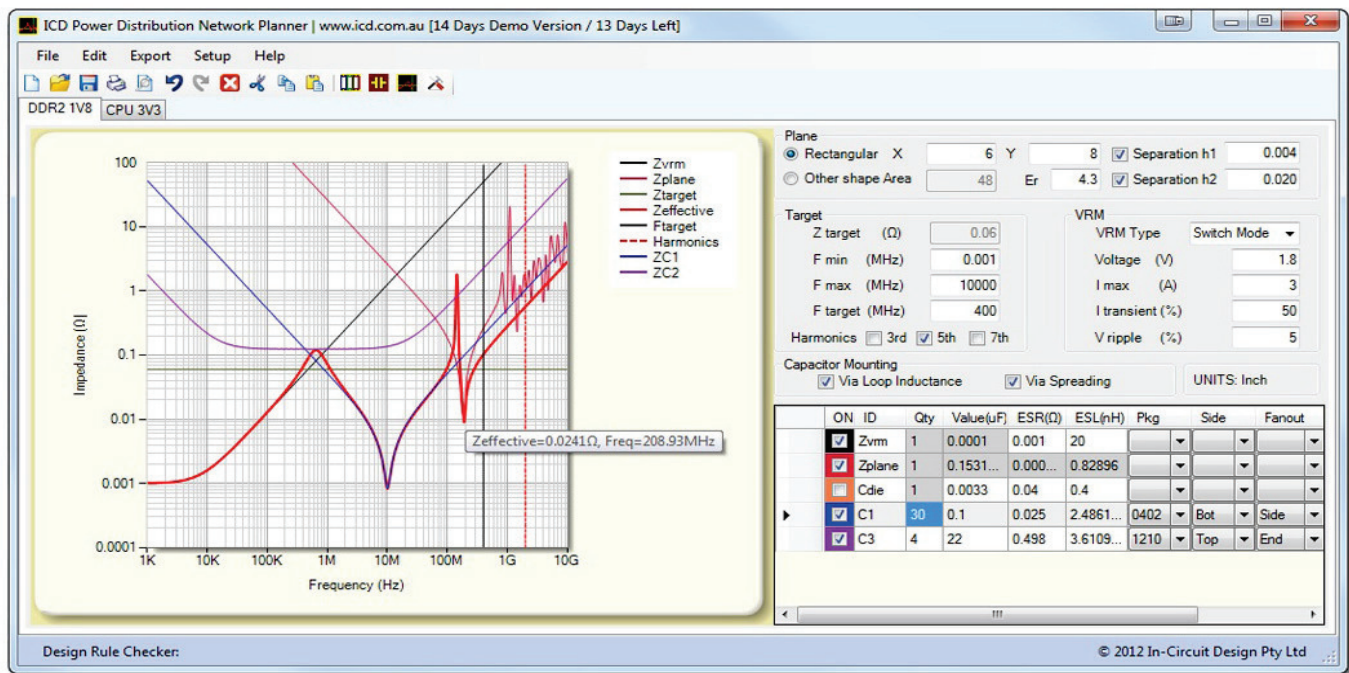


Figure 6: Illustrates the PDN requirements of DDR2 1.8V supply.

Points to Remember:

1. Mechanical constraints should be dealt with prior to component placement.
2. Consideration should be given to thermal conduction and convection, as well as radiation.
3. Controlling the maximum placement of devices limits maximum trace delay to about the longest Manhattan distance of the signals contained in the clock domain.
4. The main requirements of DDRx timing are that the clock must be longer than the address, control and command signals and the strobe must be longer than the data and data masks.
5. A cleaner waveform is delivered when a series terminator is placed close to the source.
6. Design re-use will save you time and ensure repeatability of design.
7. The quantity, capacitance, dielectric material, placement and mounting of decoupling capacitors are critical for the proper DDR2 interface operation. Detailed PDN analysis is required if you want to save on component costs.
8. For each DDR2 speed grade—with frequencies ranging from 200 to 400 MHz—inter-plane capacitance may be required to maintain a low-impedance supply of power. A software

tool, like the ICD PDN Planner, should be employed to analyze whether this design objective has been met. **PCB**

References

1. Advanced Design for SMT—Barry Olney
2. Beyond Design - Intro to Board-Level Simulation and the PCB Design Process—Barry Olney
3. PCB Design Techniques for DDR, DDR2 & DDR3, Part 1 & 2—Barry Olney
4. Beyond Design - Power Distribution Network Planning—Barry Olney



Barry Olney is managing director of In-Circuit Design Pty Ltd. (ICD), Australia, a PCB Design Service Bureau and board level simulation specialist. Among others, ICD was awarded "Top 2005 Asian Distributor Marketing," and "Top 2005 Worldwide Distributor Marketing" by Mentor Graphics, Board System Division. For more information, contact Olney at +61 4123 14441 or by e-mail: b.olney@icd.com.au.