

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 28,000 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
- PDN EMI Plot with EMC Limits. Frequency range up to 100GHz
- Extensive Capacitor Library – over 5,650 capacitors derived from SPICE models

Rock Steady Design

by Barry Olney

IN-CIRCUIT DESIGN PTY LTD / AUSTRALIA

How do we ensure that our high-speed digital design performs to expectations, is stable given all possible diverse environments, and is reliable over the product's projected life cycle? One word: Impedance!

For the perfect transfer of energy and to benefit from the highest possible bandwidth, the impedance of the driver must match the impedance of the transmission line and be constant along its entire length. Also, the power distribution network (PDN) must provide low AC impedance up to the maximum bandwidth and deliver an undisrupted, minimal return path for high-current switching devices. These two seemingly unrelated concepts are controlled by the PCB stackup configuration including mate-

rial selection and should be analyzed concurrently.

For a sufficiently large number of electronics products, failures are distributed in time as shown in Figure 1. This curve is called "the bathtub curve" and displays the typical reliability of diverse products regardless of their functionality. One would expect a product to fail after some years of service, but preferably long after the product becomes obsolete. Premature failures are of particular concern and are typically the result of poor design practice or sub-standard manufacture. This column will focus on the design aspects.

In a previous column, [Intro to Board-Level Simulation and the PCB Design Process](#), I men-

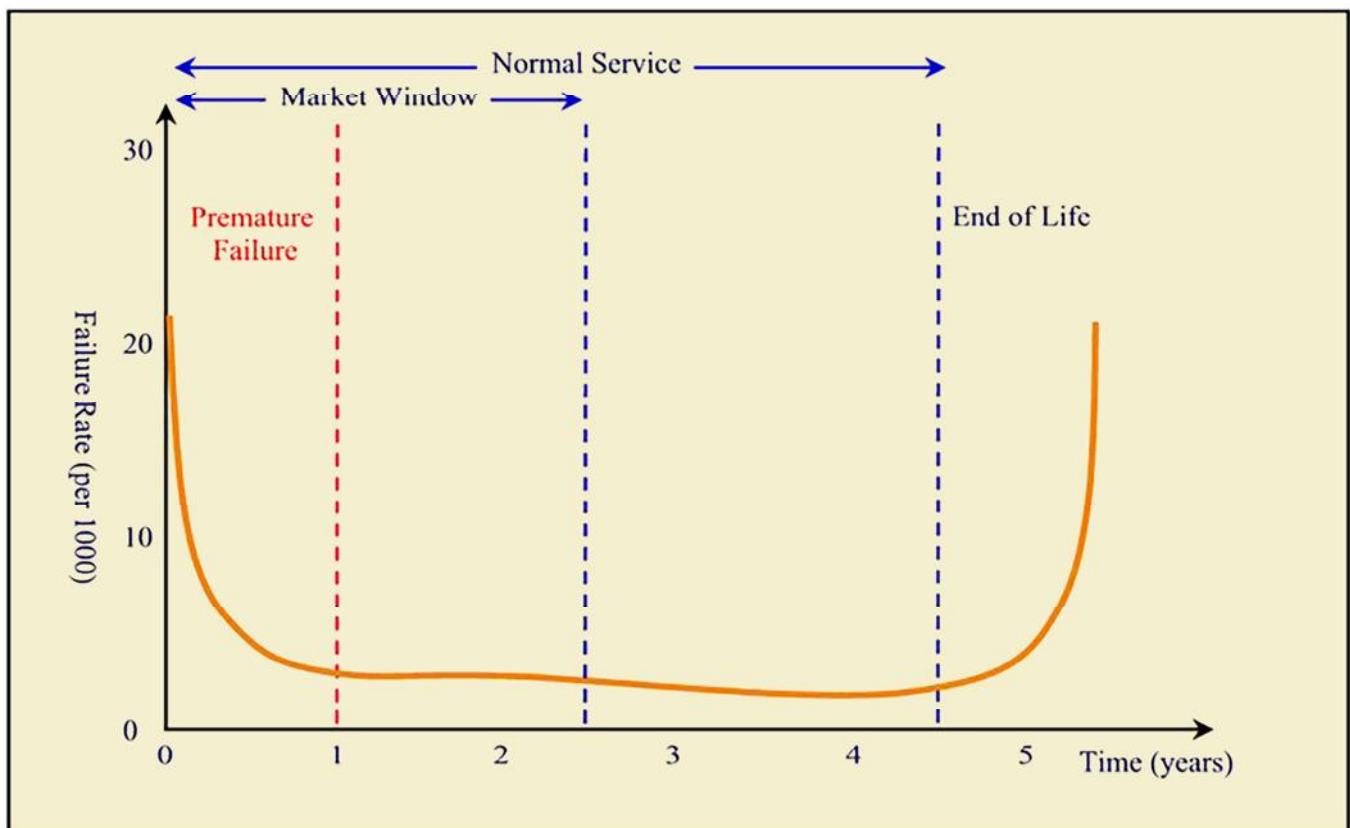


Figure 1: Product reliability life cycle—the bathtub curve.

tioned that the cost of development is dramatically reduced if simulation is employed early in the design cycle. If changes are made late in the design process, then it takes more time, people, material and therefore money to complete the project. The advantage of simulation is that it identifies issues early in the design process and rectifies them before they become a major problem. Design changes that occur:

- In the conceptual stage cost nothing;
- During the design stage requires just a little extra time;
- During the test stage means that you have to regress one stage;
- During production, or worse still, in the field, can cost millions to fix and possibly damage the company's reputation.

Reference designs are arguably the cause of many reliability issues. Many reference designs are developed by academics who are lacking the knowledge of DFM and reliability and have little appreciation for real-world industry expectations. And although their design may work on the test bench, it may not work in an adverse

environment or with a variance of vendor components where the margins become borderline. So, what are the key pillars of stability?

1. Stackup Impedance

First of all, one needs to plan the most efficient stackup configuration for the design. For a typical digital design, 50–60 ohms of impedance is recommended. However, more than one impedance is generally required to convey differential pairs of various technologies. Therefore, all combinations of impedances must be defined on the one substrate, which can sometimes be a difficult task unless you have the right tool.

Also, dielectric materials vary in both dielectric constant and dissipation loss with frequency. And, although we may be tempted to use the fundamental frequency of the clock as the operating frequency, the maximum bandwidth is determined by the signal rise time. An upper knee frequency of $0.5/T_r$, forms a crude but useful translation between time and frequency domains. So if, for instance, the rise time is 500ps, which is typical these days, then the upper bandwidth is actually 1GHz regardless of the clock frequency.

Layer No.	Via Span & Hole Diameter	Description	Layer Name	Material Type	Dielectric Constant	Dielectric Thickness	Copper Thickness	Trace Clearance	Trace Width	Current (Amps)	Characteristic Impedance (Zo)	Edge Coupled Differential (Zdiff)	Broadside Coupled Differential (Zdbs)
1	8 4 4 8	Soldermask	Top	Conductive	3.5	0.5	2.2	12	4	0.43	51.67	98.65	
2		Prepreg	GND	370HR; 1080; Rc=66% (1GHz)	3.97	2.9	1.4						
3		Core	Inner 3	370HR; 1-1652; Rc=43% (1GHz)	4.4	5	1.4	16	4	0.31	51.23	99.63	48.89
4		Prepreg	Signal	370HR; 2116; Rc=56% (1GHz)	4.14	4.8	1.4	16	4	0.31	51.23	99.63	48.89
5		Core	Signal	370HR; 1-1652; Rc=43% (1GHz)	4.4	5	1.4	16	4	0.31	51.23	99.63	48.89
6		Prepreg	PWR	370HR; 106; Rc=76% (1GHz)	3.74	2.3	1.4						
7		Core	GND	Conductive			1.4						
8		Prepreg	Inner 7	370HR; 1-1652; Rc=43% (1GHz)	4.4	5	1.4	16	4	0.31	51.23	99.63	48.89
9		Core	Inner 8	370HR; 2116; Rc=56% (1GHz)	4.14	4.8	1.4	16	4	0.31	51.23	99.63	48.89
10		Prepreg	GND	370HR; 1-1652; Rc=43% (1GHz)	4.4	5	1.4						
		Prepreg	Bottom	370HR; 1080; Rc=66% (1GHz)	3.97	2.9	2.2	12	4	0.43	51.67	98.65	
		Soldermask	Bottom	Conductive			2.2	12	4	0.43	51.67	98.65	
		Soldermask	Bottom	PSR-4000 HFX Satin / CA-40 HF LPI	3.5	0.5							

Figure 2: Typical 10-layer stackup.

Target Impedance	Tol +/- %	Layer No.	Trace Clearance	Trace Width	Current (Amps)	Characteristic Impedance (Zo)	Edge Coupled Differential (Zdiff)	Broadside Coupled Differential (Zdbs)
50/100 Digital	10	1	12	4	0.43	51.67	98.65	
50/100 Digital	10	3	16	4	0.31	51.23	99.63	48.89
50/100 Digital	10	4	16	4	0.31	51.23	99.63	48.89
50/100 Digital	10	7	16	4	0.31	51.23	99.63	48.89
50/100 Digital	10	8	16	4	0.31	51.23	99.63	48.89
50/100 Digital	10	10	12	4	0.43	51.67	98.65	
40/80 DDR3	10	1	12	6	0.58	42.53	81.33	
40/80 DDR3	10	3	12	6	0.42	43.15	82.31	37.58
40/80 DDR3	10	4	12	6	0.42	43.15	82.31	37.58
40/80 DDR3	10	7	12	6	0.42	43.15	82.31	37.58
40/80 DDR3	10	8	12	6	0.42	43.15	82.31	37.58
40/80 DDR3	10	10	12	6	0.58	42.53	81.33	
90 USB	10	1	6	4	0.43	52.07	91.41	
90 USB	10	3	8	4	0.31	50.99	90.58	48.89
90 USB	10	4	8	4	0.31	50.99	90.58	48.89
90 USB	10	7	8	4	0.31	50.99	90.58	48.89
90 USB	10	8	8	4	0.31	50.99	90.58	48.89
90 USB	10	10	6	4	0.43	52.07	91.41	

Figure 3: Target impedance chart for all required variations.

The stackup configuration in Figure 2 has been set up for multiple single-ended and differential impedances. The material selected in this case is Isola 370HR, which is a high-performance E-glass fabric for applications where maximum thermal performance and reliability are required. You should work closely with your preferred fab shop to ensure they stock the selected materials. As mentioned, the dielectric constant and dissipation factor vary with frequency so I have selected the data from the library at 1GHz to match the maximum bandwidth required. Blind and buried via spans have also been defined.

Figure 3, is a target impedance chart that includes characteristic, edge-coupled and broadside-coupled impedance that is exported as part of a fabrication drawing. This drawing also includes dielectric material definitions, trace/clearance properties and via spans.

2. Impedance Matching

Now that the stackup is finalized, let's look at impedance matching the source to the transmission lines. In a previous column, I pointed out all of the possible termination strategies. The easiest way is to use a resistive element. Termination can be made at the source or at the load. Series termination is excellent for point to point routes (such as data signals), one load per

net. It works well for traces that are electrically short and is also used to fanout multiple loads radially from a common source. Whereas, parallel termination is preferred for address buses.

The impedance of the trace is extremely important, as any mismatch along the transmission path will result in a reduction in signal quality and possibly the radiation of noise. Mismatched impedance causes signals to reflect back and forth along the lines, which causes ringing at the load. The ringing (as displayed in Figure 4) reduces the dynamic range of the receiver, eats into the noise budget and can cause false triggering.

As signal rise times increase, consideration should be given to the propagation time and reflections of a routed trace. If the propagation time and reflection, from source to load, are longer than the edge transition time, an electrically long trace will exist. If the transmission line is short, reflections still occur but will be overwhelmed by the rising or falling edge and may not pose a problem. But even if the trace is short, termination may still be required if the load is capacitive or highly inductive to prevent ringing. Generally, when the trace length exceeds one sixth of the electrical length of the rising edge rate, then termination is required.

Unfortunately, the driver impedance is always lower than 50 ohms and as such, does not

match the transmission line impedance. Given that the impedance of the transmission lines has been established, how do we calculate the value of the series terminator required for the configuration? The source impedance from the IBIS model of the driver needs to be determined. I defined this method in detail in a previous column. Basically, the series element is given by:

$$Z_{series} = Z_o - Z_{source}$$

e.g. for a 50-ohm transmission line:

$$Z_{series} = 50 - 26.16 = 23.84 \text{ ohms}$$

Impedance matching slows down the rise and fall times, reduces the ringing (over/under shoot) of clock drivers and enhances the signal integrity of a high-speed design.

3. Power Distribution Network Impedance

Now that the stackup impedance and terminations have been defined, to provide stable signal propagation, we need to ensure that the power planes and associated decoupling can handle the high switching current demanded by the processor and memory devices. Inadequate power delivery can exhibit intermittent signal integrity issues. These include high crosstalk and excessive emission of electromagnetic radiation, degrading performance and reliability of the product. The PDN must accommodate variances of current transients, with as little change in power supply voltages as possible. So, the goal of PDN planning is to design a stable power source for all the required power supplies. As with stackup planning, the PDN design is required before a single IC is placed on the board.




Figure 4: Ringing is reduced dramatically by adding a series terminator (simulated in HyperLynx).

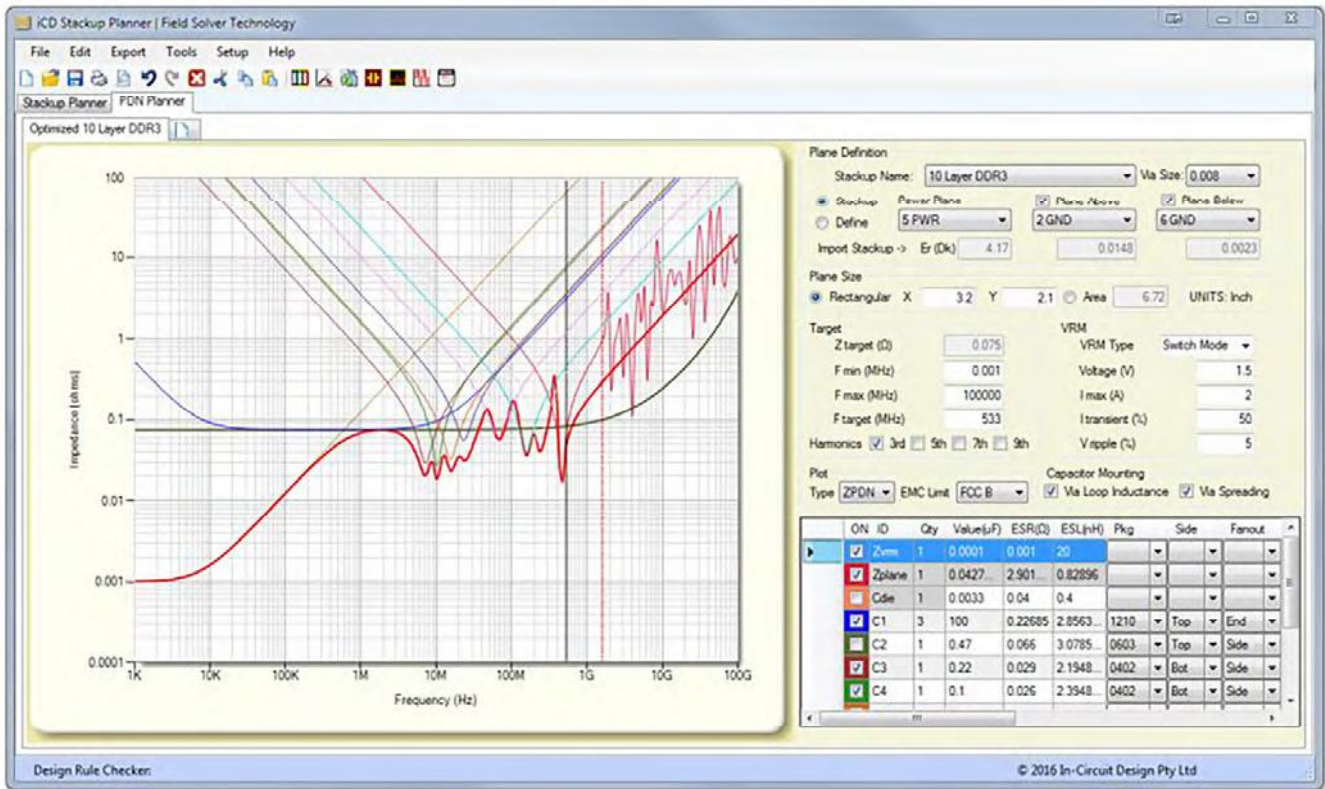


Figure 5: Optimized DDR3 PDN.

Also, the same PDN connections (planes) that are used to transport high transient currents are used to carry the return currents for critical signal transmission lines. If high frequency switching noise exists, on the planes, coupling may occur resulting in ground (supply) bounce, bit failure or timing errors. Many of the failures to pass electromagnetic compatibility (EMC) are due to excessive noise on the PDN coupling into external cables and radiating emissions.

Figure 5 illustrates a PDN with optimized capacitor values. This has 21 capacitors of different values and numbers to optimize the overall AC impedance. In this case, 21 capacitors from 100µF to 4.7nF are used. This approach gives a response close to the target impedance from DC to 700MHz. There are, in this case, a few anti-resonance peaks but they are way below the fundamental frequency and there is very low impedance right on 533MHz. When the capacitor’s self-resonant frequencies are spread, the parallel resonant impedance sets the limits to the PDN performance.

The PDN is linked to the stackup and therefore any adjustments to the stackup configuration whether it be materials, vias or trace, clearance or thickness parameters will also be reflected in the PDN. The PDN can be fine-tuned by adding more planar capacitance without affecting transmission line impedance. The optimization of the PDN is a trial and error process that needs to be done in conjunction with the stackup materials to fully exploit all avenues.

It is amazing how many designers do not get the basic key pillars of stability right. For very little extra effort, your design can have improved performance and reliability over a wide range of operating environments giving you greater confidence in your products performance for the projected lifetime.

Points to Remember

- The impedance of the driver must match the impedance of the transmission line.
- The PDN must provide low AC impedance up to the maximum bandwidth.

- The “bathtub curve” displays the typical reliability of diverse products regardless of their functionality.
 - Premature failures are typically the result of poor design practice or substandard manufacture.
 - The cost of development is dramatically reduced if simulation is employed early in the design cycle.
 - Reference designs are arguably the cause of many reliability issues.
 - The key pillars of stability: Firstly, plan the most efficient stackup configuration for the design. Dielectric materials vary in both dielectric constant and dissipation loss with frequency so the maximum bandwidth is used.
 - Secondly, impedance match the source to the transmission lines. To do this, the source impedance from the IBIS model of the driver needs to be calculated.
 - Finally, ensure that the power planes and associated decoupling can handle the high switching current demanded by the processor and memory devices.
- Optimization of the PDN is a trial-and-error process that needs to be done in conjunction with the stackup materials to fully exploit all avenues. **PCBDESIGN**

References:

1. Barry Olney’s Beyond Design columns: [Intro to Board-Level Simulation and the PCB Design Process](#), [Impedance Matching: Terminations](#), [Learning the Curve](#), [PDN Planning and Capacitor Selection, Part 1 & 2](#), [Master “Black Magic” with Howard Johnson’s Seminars](#).
2. Art and Engineering in Product Design, by Andrew Taylor.



Barry Olney is managing director of In-Circuit Design Pty Ltd (ICD) Australia. The company is a PCB design service bureau that specializes in board-level simulation. ICD has developed the ICD Stackup Planner and ICD PDN Planner software, which is available [here](#). To contact Barry, [click here](#).

DARPA’s Microsystems Technology Office Streamlines Contracting for Innovators

DARPA’s Microsystems Technology Office (MTO) has a proud history of making seminal investments in breakthrough technologies that ultimately became critical components in our electronics-filled world, from flash memory to radio frequency (RF) semiconductors to microelectromechanical systems (MEMS). But DARPA does not develop technologies on its own. The Agency’s approach is to set extremely challenging goals and then offer innovators at universities and companies the support they need to pursue those remote but exciting frontiers.

MTO’s commitment to that catalytic role—and the opportunity to engage with DARPA on the next generation of cutting-edge advances—is now expanding with the introduction of a simpler contracting approach for companies and other entities that have not previously worked with DARPA or had large contracts with the Defense Department.



This new approach aims to help the Agency and the nation take fuller advantage of the enormous depth and breadth of private-sector creativity that is currently brewing in the fast-evolving domains of networked sensors, spectrum access, machine learning, and hardware security. In particular, it aims to reduce barriers for innovative companies that don’t

engage in the standard federal contracting process.

For example, DARPA recently worked with two small companies on twin advances that could enable next-generation radio frequency (RF) arrays for both military systems and commercial wireless communications. The newly announced approach takes advantage of DARPA’s so-called Other Transactional (OT) authority, which grants the Agency certain alternatives to provisions in the standard Defense Department contracting rules, known as the Federal Acquisition Regulation, or FAR.