



Bogatin Enterprises, LLC
Setting the Standard for Signal Integrity Training
Multimedia SI training: read it, watch it, hear it, do it,
answer it, ask it, *experience it!*
www.BeTheSignal.com

Transforming complexity into practical design principles

Efficiently Designing Power Distribution Networks (PDN)

By Dr. Eric Bogatin, Signal Integrity Evangelist, Bogatin Enterprises, www.beTheSignal.com and Chuck Ferry, Product Marketing Manager, Mentor Graphics, www.mentor.com/hyperlynx

Simulation should be your friend

The earlier in the design cycle the right decisions can be made, the shorter the development time and the lower the development costs. This is probably the most important product development principle.

It is especially important when interconnects are not transparent and signal or power integrity is a key element holding back performance. In this regime, it's critical to find, fix and design out signal integrity problems before you commit to hardware. In our globally competitive environment, few product development programs can afford the multiple board spins that result when the "build it and test it" design approach is used.

The alternative approach is to incorporate analysis early and often in the design cycle. Three different analysis tools should be in every designer's tool box: rules of thumb, approximations and numerical simulation tools. Each is equally important and each has a different role to play in the design process.

Rules of thumb have value in feeding your intuition and helping to make big picture decisions. Approximations can be used to quickly and effortlessly explore design space to zero in on cost-performance design tradeoffs.

When you want to establish a specific design rule for the custom features in your board, the tool to use is a numerical simulation tool. Its chief value is in the accuracy to predict performance given the specific conditions of the physical design, material properties and driver models. The more flexible it is and easier to use, the more added value it can provide.

One of the ways of assessing the accuracy of a tool is by comparing its simulated predictions of test boards to the actual measurements. Mentor Graphics recently released a power integrity feature in HyperLynx 8.0. This article discusses its accuracy with a few examples.

PDN simulation

The power distribution network, PDN, includes all the interconnectors from the voltage regulator module (VRM) to the pads on the chip. Its purpose is to provide a constant voltage on the pads of the chips, even when the current through the chip varies by hundreds of amps and at frequencies from DC into the GHz range.

The primary design guideline to achieve this low voltage ripple is to keep the impedance from the chip pads to the VRM below a target value. When the peak current draw is on the order of 1A with a 1 volt rail, the load impedance the chip represents is about 1 Ohm, and the PDN impedance should be less than about 0.1 Ohm to keep the voltage ripple on the chips' pads less than 5%. Larger current draws require even lower PDN impedance targets. High end graphics chips that draw 100 amps, require a PDN impedance less than 1 mOhm, from DC to the GHz regime.

Luckily for board designers, the equivalent lead inductance in the chip package usually limits the bandwidth of the impedance the board level features can affect to under 100 MHz. This means board level design needs to keep the impedance of the PDN low only up to about 100 MHz. This is done by power and ground planes and decoupling capacitors, integrated with as low equivalent series inductance as possible.

A PDN simulator needs to be able to include the important interactions of the capacitors and planes including the spreading inductance effects of Swiss cheese planes and odd shapes, common with split plane applications.

For example, the impedance of a capacitor by itself, including its equivalent series resistance, ESR and its equivalent series inductance, ESL, based on how it is mounted onto the board, behaves like a simple RLC circuit. The impedance profile of a plane is more complicated, with spreading inductance and resonance effects. Examples of two such impedance profiles are shown in Figure 1.

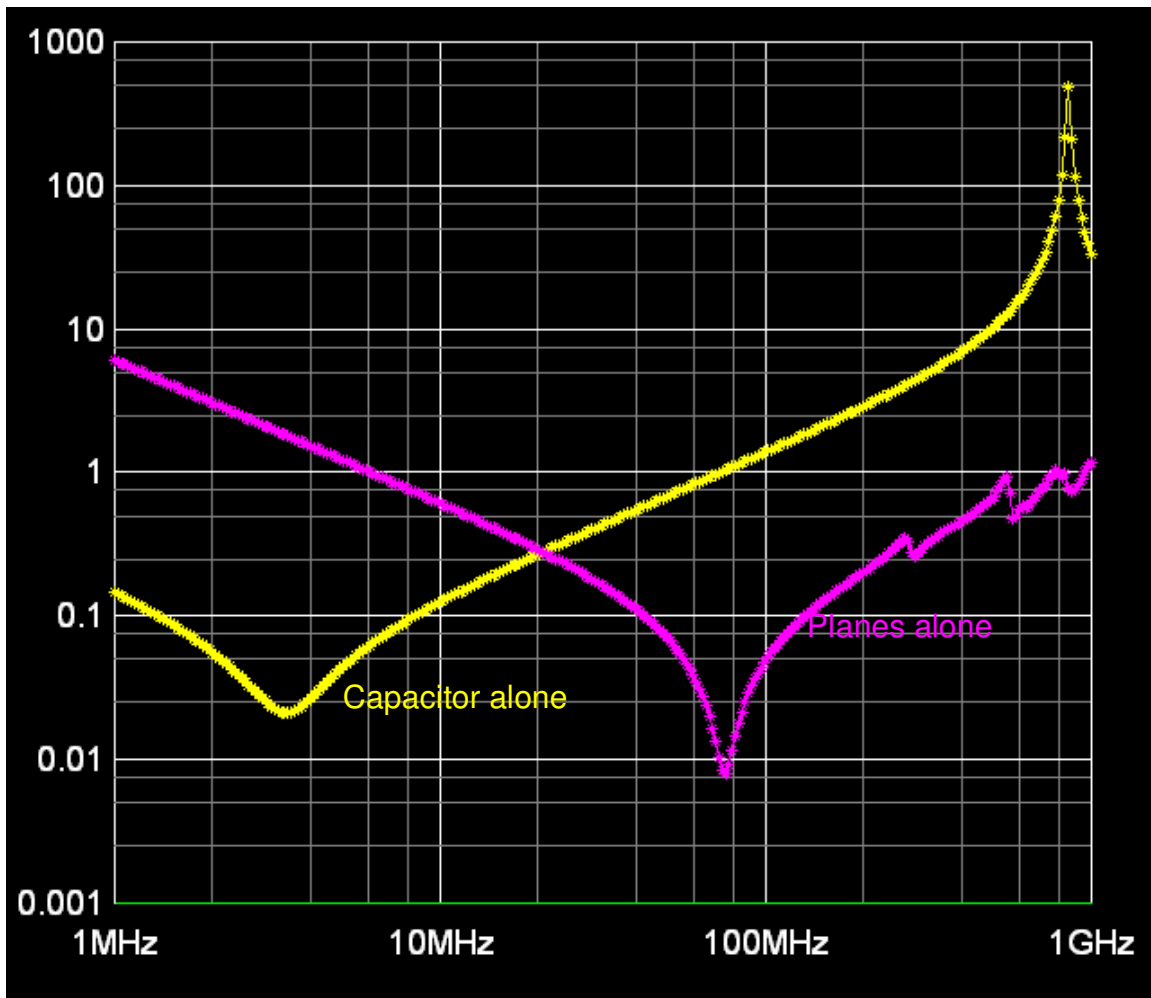


Figure 1. Impedance of a 1 uF capacitor and a 5 inch x 5 inch power and ground plane pair, simulated with HyperLynx.

When the capacitor is mounted on the power and ground planes there are complex interactions. The simulated impedance of the combination is shown in Figure 2.

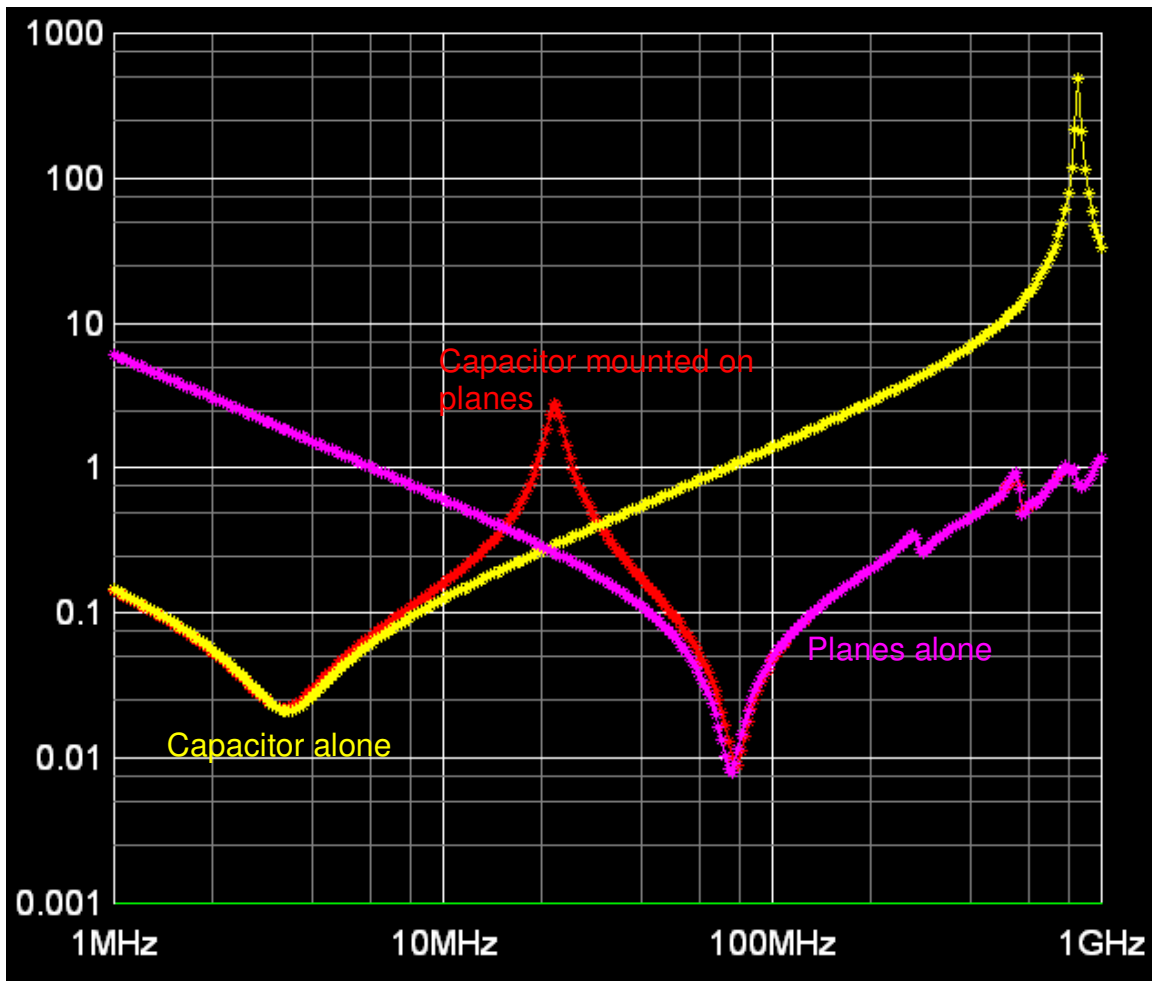


Figure 2. Impedance profile of a capacitor mounted on a power/ground plane pair showing the parallel resonance peak, simulated with HyperLynx.

The Test Vehicle

To establish a level of confidence in the accuracy of the impedance simulation tool, a simple six layer test board was build up with 16 different modules with various combinations of plane layers and 0603 decoupling capacitors mounted to the top surface. A close up of one of the modules is shown in Figure 3.




Figure 3. Close up of the top surface of one of the test modules showing an 0603 capacitor with via in pad and the connections for the 2-port impedance measurement.

Measuring the impedance as seen by a package pin connected to the power and ground cavity and decoupling capacitors, is easy. Measuring it without fixture artifacts is tricky. In this test vehicle, we used the very powerful 2-port method to minimize the effect of the contacts to the planes. In this method, two SMA connectors make contact to the same layer pair in the board and a 2-port VNA is used to measure the 2-Port S-parameters to these connections.

Normally, a 1-port measurement can be used to measure impedance. However, if the calibration is done to the end of the coax cable, the impedance measurement will have in it the impedance of the connector to the board. When the impedance you want to measure is in the milliohm range, having a connector, which might contribute an ohm of impedance at 1 GHz, is an irritating artifact. It cannot be calibrated or compensated out.

The 2-port technique is analogous to the 4-point Kelvin technique for measuring very low resistance as a way of eliminating contact resistance artifacts. In the DC technique, one pair of contacts forces a known current while the other pair of contacts measures the voltage generated as this current flows through the resistance of the device under test.

With 1 A of forced current, and a measured voltage as low as 1 mV, it is easily possible to measure resistances below 1 milliohm.

In a 2-port VNA measurement, one port forces a known current through the device under test, while the other port measures the voltage across the impedance. The impedance information is in the S21 term. This technique enables the measurement of milliohms of impedance into the GHz range.

An example of the measured impedance of a capacitor mounted to the power and ground planes with a 1-port and 2-port measurement is shown in Figure 4.

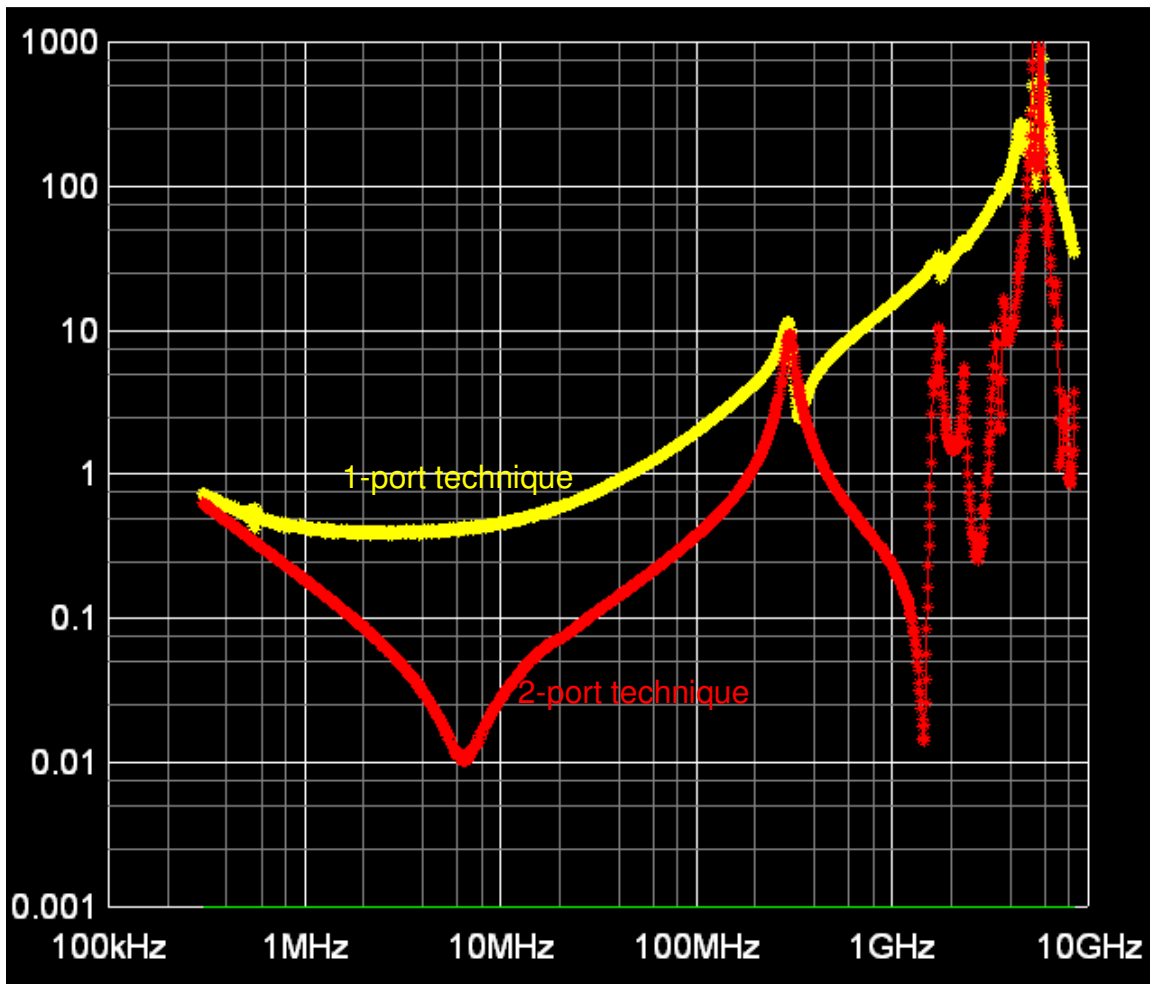


Figure 4. Measured impedance profiles of a test board and capacitor with 1-port and 2-port techniques. Done correctly, the 2-port method enables the impedance seen looking into the boards, independent of the connector artifacts. The measurement was done from 300 kHz to 8 GHz.

It is essential, when measuring the impedance of any element in the power distribution network, to use the 2-port VNA technique, otherwise, the measurements obtained will be dominated by artifacts of the connector and never match what is going on in reality.

Comparison of Simulation and Measurement

The old cliché, “Garbage in, Garbage out” was invented to describe simulations. The quality of a simulation can never be better than the quality of the input information about the material properties and geometry of the structures simulated or the model of the components added.

When simulating the bare board, the lateral dimensions, cavity thickness, depth of the cavity from the top surface and the dielectric constant and dissipation factor of the laminate materials are all important input parameters to the simulated impedance profile.

When the correct values of input parameters are used, the ability to predict the bare board impedance is excellent, as shown in Figure 5 for the case of one of the modules.

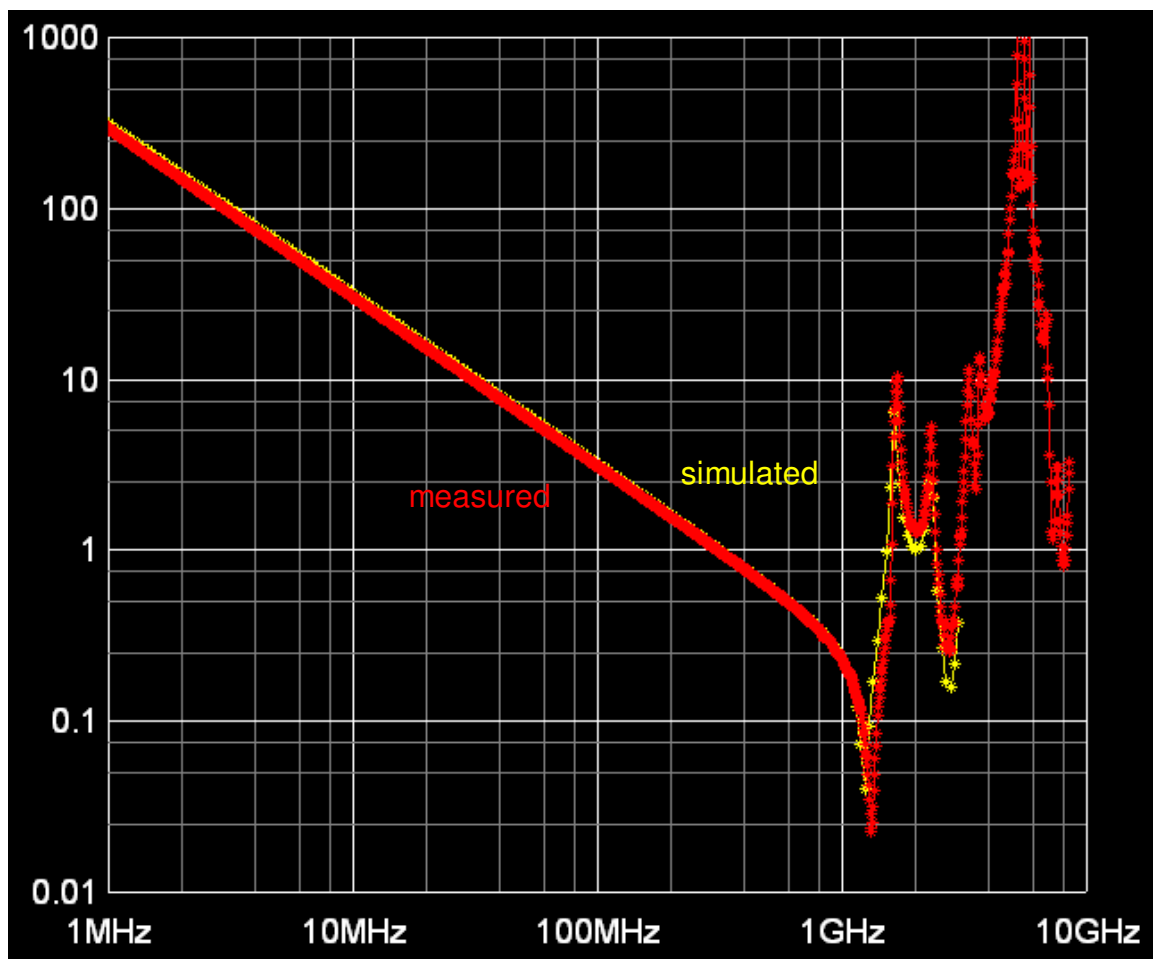


Figure 5. Comparison of the simulated and measured impedance profile of just the power/ground cavity of the test board. The impedance profiles are almost on top of each other. Simulated with HyperLynx.

When a capacitor is added to the board, the agreement between the simulated impedance profile and measured impedance profile is also very good. This is shown in Figure 6.

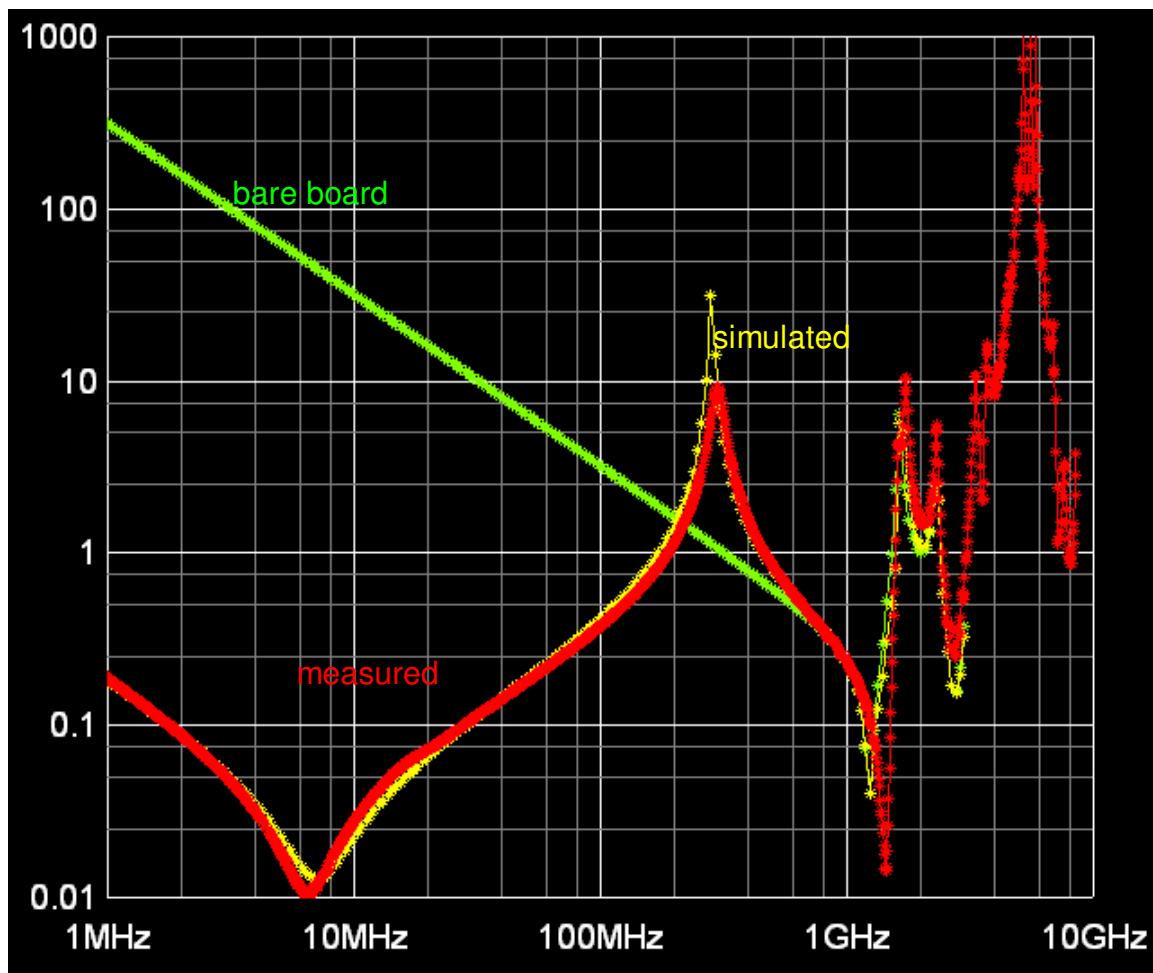


Figure 6. Measured and simulated impedance profile for a board with a mounted capacitor, and the measured impedance of the bare board for comparison. Simulated with HyperLynx.

Impact from the mounting structure

Using these test vehicles and HyperLynx 8.0, some of the recommended design features for mounting capacitors to boards can be evaluated. The goal in any PDN design is to arrive at the lowest impedance possible, especially below 100 MHz. While the capacitor selection plays a role at low frequency, it plays little role at high frequency. Rather, the high frequency impedance is all about the physical design of the interconnects of the capacitor to the power and ground cavity.

Three factors are important:

1. keep surface traces from the capacitor to the vias as short and wide as possible.
2. bring the power/gnd cavity as close to the surface as possible
3. use as thin a dielectric as possible between the power and ground planes.

By following these three important guidelines, the loop inductance associated with the capacitor to the package pins can be minimized and the impedance at high frequency can be minimized.

Many capacitor vendors will offer an inductance for their capacitors as a way of estimating the impedance profile their component may provide. These values are completely worthless as the loop inductance of a capacitor depends so strongly on the rest of the interconnects linking it to the power and ground planes.

In fact, the same capacitor can show more than an order of magnitude different mounted inductance, just by changing its integration with the board. In the last example, we compare the measured impedance of the same 100 nF capacitor, as seen by the 2 ports of the VNA, with two different configurations.

Case 1 is a good case: via in pad, power and ground cavity close to the surface and thin dielectric in the cavity.

Case 2, is an example of not doing everything right. There are 30 mil long surface traces from the capacitor pads to the via and the cavity is about 55 mils below the surface.

The measured impedance profiles as seen by the VNA, are shown in Figure 7.

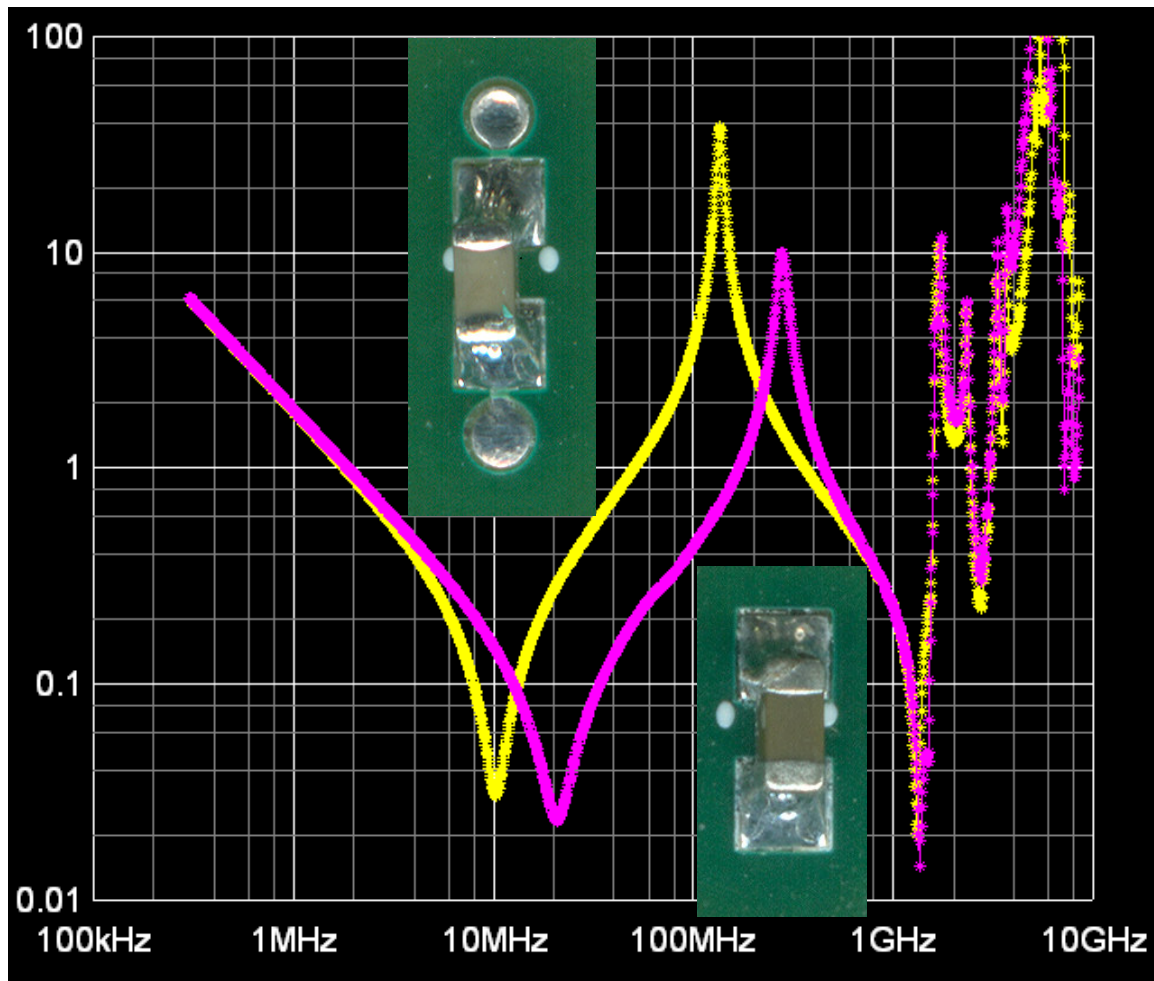


Figure 7. Measured impedance profiles of the same capacitor mounted in two different ways to a board, showing 5x difference in impedance from 10 MHz to 100 MHz.

The impedance of case 1 is less than 20% the impedance of case 2. This means, if you do the integration right, you could achieve the same target impedance with 1/5 the number of capacitors as would be needed if you used the mounting geometry of case 2. This could be a significant cost savings.

This is why there is no one magic formula for how many capacitors you should use- it depends so heavily on what mounting inductance you can implement.

Of course, you should always follow the design guidelines to minimize the impedance of the PDN. But sometimes, it's not practical due to the specific constraints of your product. The power plane may be an odd shape because it shares the same layer as some signals. There maybe a large clearance hole field due to a BGA. The stack up may not allow the power/gnd cavity near the surface.

This is why it is essential to explore design space for the optimum capacitor selection, number and distribution using an accurate simulation tool.

For additional information on PDN design, visit

www.beTheSignal.com

www.mentor.com/hyperlynx