

Frequency-Domain Power-Distribution Measurements – An Overview

Istvan Novak

Sun Microsystems, Inc.



Author Biography

Istvan Novak is signal-integrity senior staff engineer at SUN Microsystems, Inc. Besides signal-integrity design of high-speed serial and parallel buses, he is engaged in the design and characterization of power-distribution networks and packages for workgroup servers. He creates simulation models, and develops measurement techniques for power distribution. Istvan has twenty plus years of experience with high-speed digital, RF, and analog circuit and system design. He is Fellow of IEEE for his contributions to the signal-integrity and RF measurement and simulation methodologies.

Abstract

In today's advanced digital systems, the power-distribution networks often have to deliver hundreds of watts at low voltages, and the required low impedance of the power-distribution must be maintained over a wide frequency band. Traditional high-frequency measuring instrumentation has been tailored to handle impedances close to 50 ohms. The very low impedance values in the power-distribution networks create measurement and calibration challenges. This presentation explains the benefits of the frequency-domain method and impedance measurements with two-port VNA setups. Extraction of component parameters, calibration options for different instruments and frequency ranges, probe connections and constructions are explained. Various compensation methods are described, and examples are given to measure PDN components from single-elements to full working systems. The concept of attached impedance/inductance of bypass capacitors is introduced and we show data supporting the claim that for the high-frequency inductance of bypass capacitors the cover thickness of the part matters, the total height of the part does not.

Outline

- I. Introduction: Why frequency domain?
- II. How to measure PDN impedance
 - II.1. Why two ports
 - II. 2. Connection between self and transfer impedances
 - II. 3. Measuring the magnitude of $|Z| \ll 25$ ohm values
 - II. 4. Measuring arbitrary impedance values
- III. Extracting component parameters
 - III. 1. Extracting capacitance
 - III.1.1. Compensating for series inductance
 - III.2. Extracting Equivalent Series Resistance (ESR)
 - III. 3. Extracting inductance
 - III. 3. 1. Compensating for series capacitance
 - III. 3. 2. Compensating for parallel capacitance
 - III. 3. 3. Compensating for series and parallel capacitances
 - III. 4. Extracting parameters of PCB plane pairs
- IV. Calibrations and probes
 - IV. 1. No calibration
 - IV. 2. Through calibration
 - IV. 3. Full two-port calibration
- V. Making the proper connections
 - V. 1. Eliminating cable-braid-loop error at low frequencies
 - V. 1. 1. Ferrite clamps/sleeves
 - V. 1. 2 Isolation transformer or amplifier
 - V. 2. Measuring low-ESR bulk capacitors
 - V. 3. Measuring Voltage Regulator Modules
 - V. 4. Measuring low impedances at high frequencies
- VI. Putting it all together
 - VI. 1. Measuring a full PDN
 - VI. 2. Detailed characterization of PDN components
 - VI. 2. 1. Modeling the frequency dependent capacitance of low-ESR bulk capacitors
 - VI. 2. 2. Wide-band measurement of bypass capacitors
 - VI. 2. 3. Determining the inductance of multi-terminal bypass capacitors: the concept of attached inductance
 - VI. 2. 4. Impact of capacitor height on attached inductance of bypass capacitors
 - VI. 2. 5 Repeatability of data

Acknowledgements

References

I. Introduction: Why frequency domain?

The power-distribution network (PDN) of digital circuits has to feed the chips with DC power. The switching circuitry creates current transients, which across the PDN impedance generates voltage fluctuations. The voltage fluctuations (transient noise) must be kept below a predefined limit so that it does not interfere with the signaling. For each signal line, a time window can be identified, within which any supply-rail noise will reduce the noise margin. In synchronous systems this is the window defined by the setup and hold requirements. Outside the sampling window the noise will not harm the signal, except of extra large noise levels may drive the attached devices into nonlinear regions, or even cause breakdown in low-voltage devices. However, as noise propagation from multiple noise sources may be too complex to analyze in detail, a conservative approach assumes that the specified supply-rail noise limit should be kept at all times.

Because digital signaling is defined by its voltage or current levels in the time domain, it may seem obvious to characterize the behavior of the PDN also in the time domain. While the characterization itself is doable in the time domain, validation in a real system in the time domain is very hard. Here is why.

Having multiple active devices connected to the same PDN rail, the noise voltage at each location becomes the sum of the products of the appropriate self and transfer impedances and noise currents. Figure 1 shows a simplified sketch of a circuit, where the PDN consists of a PCB, one power-ground plane pair, three bypass capacitors and two active devices. For sake of simplicity lets assume that we are interested in the noise voltages at the active devices only, and the size of the devices and the highest frequency of interest allow us to use only one test point for each active device. We can further assume that we know from simulations or from measurements the impedance matrix of the PDN for the two test points. With very few exceptions, the PDN components are electrically reciprocal, therefore $Z_{12}=Z_{21}=Z_M$. Electrical symmetry, however, cannot be assumed in a generic case; therefore Z_{11} and Z_{22} are, in general, different. The noise voltages at test points 1 and 2, generated by the noise currents of $I_1(t)$ and $I_2(t)$ of the two active devices, can be expressed as:

$$\begin{aligned} V_1(t) &= Z_{11}I_1(t) + Z_M I_2(t) \\ V_2(t) &= Z_M I_1(t) + Z_{22}I_2(t) \end{aligned} \tag{1}$$

Note that while in this simple test case we assumed only two test points, still, the impedance matrix elements must be obtained for these two ports together with the three additional locations of bypass capacitors.

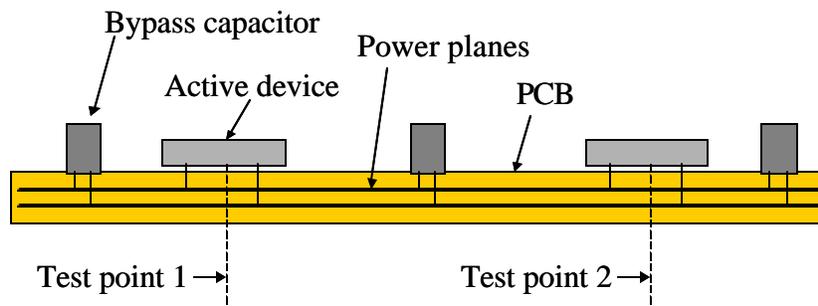


Figure 1 Simple sketch of PDN with two active devices, three capacitors and one pair of power planes.

For M active devices and test ports on the PDN, (1) can be extended, and the noise voltage at the k th test point becomes:

$$v_k(t) = \sum_{i=1}^M Z_{ik} I_i(t) \quad (2)$$

If we want to close the PDN design cycle with validation in the time domain, we have to know both the impedance matrix and the transient noise vector. The PDN is built of passive components (except the voltage regulator modules); therefore the PDN's impedance is less subject to statistical variations due to component tolerances. With well-behaving components, it is certainly not much time varying. In contrast, the transient noise in a complex system comes from many packages of different active devices, and the sources have their own timing and activity schedule. On a large board there may be hundreds of packages with thousands of active cells switching. Predicting or simulating the worst-case transient noise in the entire system is a daunting task. It is daunting, because transfer impedances and propagation delays among the noise sources and test points are usually not negligible, therefore the worst-case maximum transient noise will not necessarily occur when and if all the sources switch simultaneously at the same time.

If instead, we wanted to find the transient current by measurements, the difficulty is very practical: measuring current in a small-pitch printed-circuit board is not easy, as it would require a shunt element in series to each current path, or a current-measuring loop around each current-carrying conductor to be measured. And even if we could measure the current, it does not remove the problem stemming from the highly statistical nature of the noise current.

Recently several publications described ways to indirectly measure the impedance profiles and/or transient currents, see e.g., [1] and [2], which assume that the active devices can be exercised in a controlled manner.

Though eventually we may want to measure and validate the noise in the time domain, it is more straightforward to segment the task, and do the design and validation separately for the impedance matrix, followed by time-domain measurements/validations, if necessary.

A suggested new methodology starts with the step transient responses corresponding to the test points, and calculates the absolute worst-case transient noise magnitude [2]. In this case, obtaining the impedance matrix by measurements or simulations is sufficient, because the step responses can be obtained from the impedance matrix by calculations. This way we can eventually get the worst-case estimated noise without doing extensive time-domain measurements.

Finally, yet another reason suggesting frequency-domain measurements instead of time domain is the fact that any random noise from the environment can be suppressed more readily in frequency domain. Vector Network Analyzers operate with narrow bandwidth in a synchronous mode; therefore averaging will bring out the useful signal, and will suppress random noise. When we try to measure the worst-case transient noise instead, probably with an oscilloscope in infinite persistence mode, any random external noise getting into the measurement setup will corrupt our data.

II. How to measure PDN impedance

With increasing power levels and constantly dropping supply voltages, the PDN impedance of high-power systems must be low, sometimes in the milliohm range. The low impedance values themselves create unique challenges in the selection of instruments, setups and connections.

To measure the impedance of an unknown device over a wide frequency range, we can use LRC meters, impedance bridges, or Vector Network Analyzers (VNA). With just one port (two connection points), any of these choices will be limited to impedance values of a few hundred milliohms or higher and inductances of a few hundred pH or higher. The reason for this limitation is the practical difficulty of creating connections from the measuring instrument to the unknown impedance. One millimeter (about 40 mils) of wire can represent a $Z_{\text{connection}}$ impedance of a milliohm or more resistance at low frequencies, and hundreds of pH inductance at high frequencies. We can attempt to remove this extra impedance by calibration, but we will face a difficult task in defining the geometry and connection points with such a high accuracy that the calibration/deembedding could be effective at high frequencies and low impedance values. Figure 2 shows that with one-port impedance measurements we always measure $Z_{\text{DUT}} + Z_{\text{connection}}$.

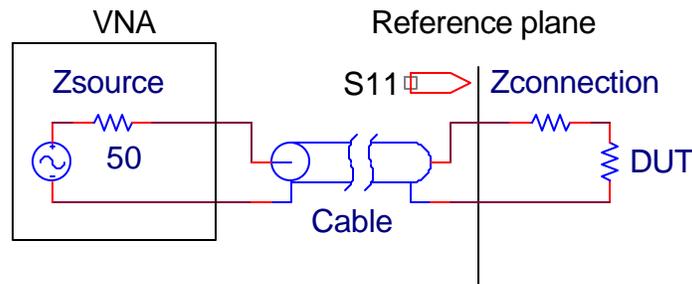


Figure 2. Setup with VNA, for one-port impedance measurements. One-port impedance measurements are limited to a few hundred milliohms mid-frequency error.

II.1. Why two ports

VNAs, on the other hand, offer a convenient way of measuring low impedances, similar to the four-wire DC resistance measurement setup.

As shown in Figure 3, two-port VNA connections use Port1 to launch a known signal current into the unknown impedance, and Port2 is used to measure the voltage drop. The same extra $Z_{\text{connection}}$ impedance that appears directly in series to the unknown impedance in one-port VNA measurements, is now transformed into the two loops of VNA ports, each having a nominally 50 ohm impedance. Though we now have two connections to care for instead of one, the connections to the DUT will introduce very little error up to several GHz frequencies. For the error analysis on self and transfer impedances, see [3].

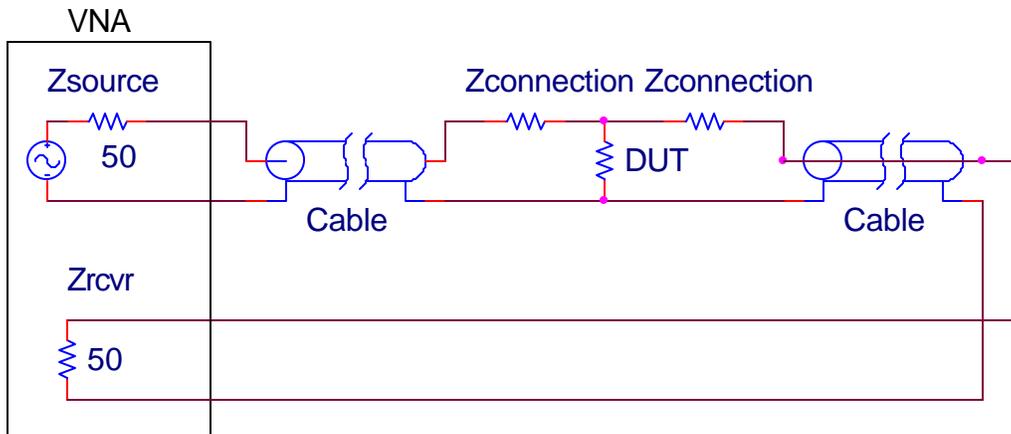


Figure 3. Setup with VNA, for two-port impedance measurements. Two-port VNA connection greatly reduces the error due to the series connection impedance, by moving them into the 50-ohm loops of VNA ports

II. 2. Connection between self and transfer impedances

By using the two-port connection of VNAs for impedance measurements, literally we always measure transfer impedance between the connections from Port1 to Port2. If, however, the two ports are connected to the same (or almost the same) location on the PDN, the transfer impedance becomes self impedance. This creates a single and unified way to measure both self and transfer impedances with the same instrumentation and setup.

The top two graphs Figure 4 illustrates the gradual change of transfer impedances into self impedance over a bare pair of square planes. The graphs assume a pair of 10"x10" square bare plane pair with 2-mil plane separation, and dielectric constant of 4. To show the effects clearly, dielectric and copper losses are ignored. The impedance curves are shown with one port being always in the middle, while the other port moves away along the diagonal, by the following distances along the x and y coordinates from the center: 0.25", 0.5", 1", 2", 3", 5". On the left, the impedance magnitudes are plotted over the 10MHz to 1GHz frequency range. Note that at low frequencies all curves follow the impedance of the static plane capacitance. At higher frequencies some of the curves hit a sharp minimum, some other curves go through a shallow bottom. Because of the loss-less assumption, all curves go through the same modal resonance peak values. Later we will show that when losses are not negligible, the modal resonance peaks do depend slightly on the location on the planes. On the right, the same set of curves is shown on zoomed horizontal and vertical scales, with labels identifying the individual curves. Label 0" indicates self impedance at the center; label 5" shows the transfer impedance between the center and corner. The labels show port separation along the x and y coordinates; the line-of-sight probe spacing is $\sqrt{2}$ times larger.

At lower frequencies the usual assumption could be that self and transfer impedances may not differ significantly, as long as the spatial dimensions are much smaller than the shortest wavelength of interest. At very low impedance values, however, the series losses of the distribution network together with the bypass capacitors may create different impedances at and between various points. This is illustrated on the measured low-frequency plots on the bottom of Figure 4. Similar measured graphs at high frequencies are shown later in section VI. 1.

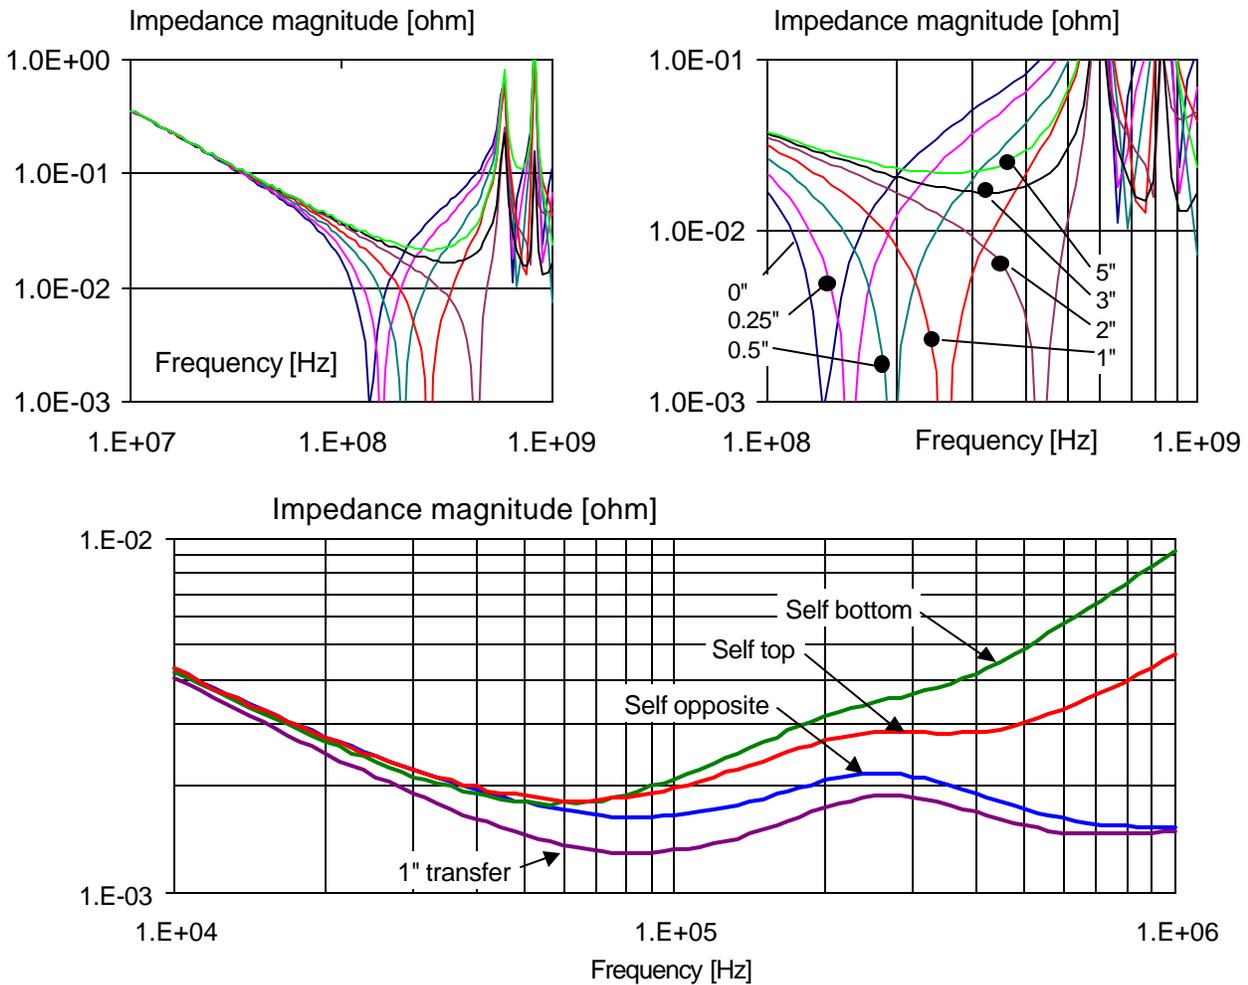


Figure 4. On top: simulated self and transfer impedances on a bare pair of power planes. Left graph: overall picture. Right graph: zoomed horizontal and vertical scales, with port connection spacing labeled. Note that the labels show the x and y distances: the port separation along the diagonals is $\sqrt{2}$ times larger. On bottom: measured self and transfer impedances on a populated board at low frequencies. There were several large bulk capacitors on the plane, creating the shallow minimum around 100kHz. The power/ground plane pair was close to the top in the stack up. The DC-DC converter was not powered up; hence we see the capacitive slope between 10 and 100kHz. There was one pair of test vias, where the VNA ports were connected. The curve labeled 'Self bottom' was measured with both VNA ports connected to the bottom pads on the test via pair, thus having the longer part of the via loop in series to the PDN impedance. At 1MHz, the equivalent inductance is approximately 1.3nH. The curve labeled 'Self top' was measured with both VNA ports connected to the top pads of the test via pair, thus going through the shorter part of the via loop in series to the PDN impedance. The equivalent inductance at 1MHz was 680pH. The curve labeled 'Self opposite' was measured with the two VNA ports connected to the opposite sides of the test via pair. The equivalent inductance at 1MHz is 100pH. The curve labeled '1" transfer' was measured with one VNA port connected to the test via pair, the other VNA port connected to a via pair 1 inch away. Note that at 100kHz, where the wavelength is 200 meters, a 1" (0.0254 meter) distance results in an impedance drop from 1.6 milliohms to 1.3 milliohms. This happens because the low ESR of the bulk capacitors forms an attenuator with the series resistance of the planes.

II. 3. Measuring the magnitude of $|Z| \ll 25$ ohm values

As long as the unknown impedance magnitude is much lower than 25 ohms, the complex voltage divider of Figure 3 between the two 50-ohm VNA connections and the unknown Z_{DUT} impedance can be simplified, and the magnitude of the unknown impedance can be approximated by

$$Z_{DUT} = 25 * S_{21}, \quad S_{21} = 10^{\frac{S_{21}^{[dB]}}{20}} \quad (3)$$

If we use the Bode-plot style output from the VNA, the dB values first have to be converted to ratios, as shown by the second expression in (3).

As long as we do not need high accuracy and don't need the phase of the unknown impedance, there is hardly any need for calibration with certain VNA types at low frequencies (see later). We just have to establish the nominal full-scale reading as 25 ohms.

II. 4. Measuring arbitrary impedance values

The impedance of an inductor or capacitor varies linearly or inversely with frequency, therefore chances are that sooner or later the impedance magnitude will not be much smaller than 25 ohms. In such cases we have to solve the voltage divider expression (4) for the unknown impedance. With the notions of Figure 3, and assuming V_{source} source voltage, and 50 ohms connecting impedance to both VNA ports, the V_{ZDUT} voltage across the unknown impedance can be expressed as:

$$V_{Z_{DUT}} = \frac{1}{2} V_{source} \frac{Z_{DUT}}{Z_{DUT} + 25} \quad (4)$$

To get the correct answer, now we cannot avoid calibration, because we also need the phase information to solve the complex equation. For medium accuracy, a simple through calibration is usually sufficient. To obtain the highest possible accuracy, a full two-port calibration should be done.

Due to the robustness of the two-port impedance measurement scheme, small discontinuities in the connections to the DUT can be neglected up to several GHz frequencies [3]. By rearranging (3), the real and imaginary parts of the unknown impedance can be expressed as:

$$\text{Re}(Z_{DUT}) = 25 * \frac{\text{Re}(S_{21}) * (1 - \text{Re}(S_{21}) - \text{Im}(S_{21})^2)}{(1 - \text{Re}(S_{21}))^2 + \text{Im}(S_{21})^2} \quad (5)$$

$$\text{Im}(Z_{DUT}) = 25 * \frac{\text{Im}(S_{21}) * (1 - \text{Re}(S_{21}) + \text{Re}(S_{21}) * \text{Im}(S_{21}))}{(1 - \text{Re}(S_{21}))^2 + \text{Im}(S_{21})^2} \quad (6)$$

Note that expressions (5) and (6) can be easily programmed in a spreadsheet that takes the measured VNA data.

The graphs of Figure 5 show the extracted impedance magnitude and phase measured on a small-size PCB with three paralleled plane pairs with and without complex solution of the voltage divider. Note that without solving the complex attenuator equation, the extracted impedance magnitude at low frequencies saturates at 25 ohms. Similarly, instead of the -90 degree phase angle that we expect from a capacitance, at low frequencies the uncorrected phase angle approaches zero degree.

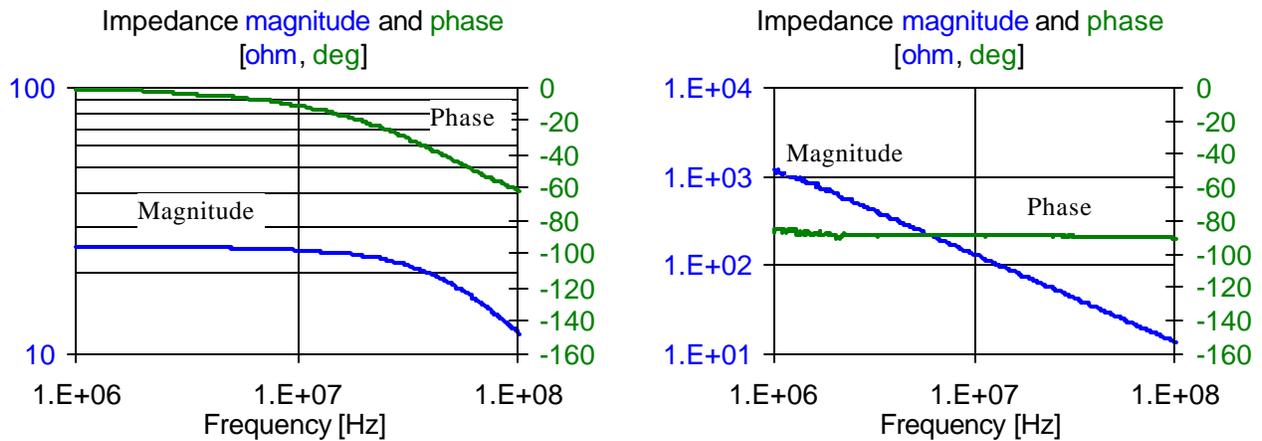
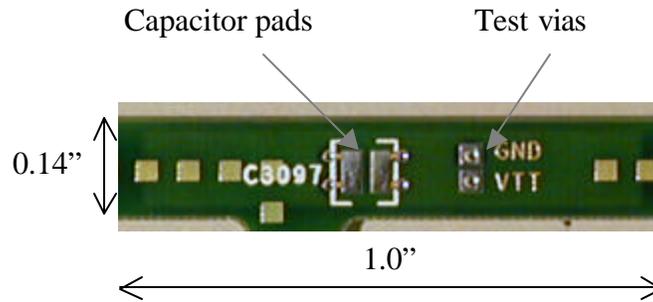


Figure 5. Measured impedance magnitude and phase of three paralleled bare plane pairs of 1"x0.14" size. On top: picture of the board with dimensions and connections. On lower left: impedance obtained without complex inversion of voltage divider formula, using (3). On lower right: same measured data using complex inversion with (5) and (6). On both graphs, left axis shows logarithmic impedance magnitude, right axis shows linear phase.

III. Extracting component parameters

Often times the measured PDN component can be approximated with a very simple equivalent circuit in a given frequency range. As shown in Figure 6, the most common equivalent circuits for PDN components are series C-R-L, or parallel C-R-L circuits.

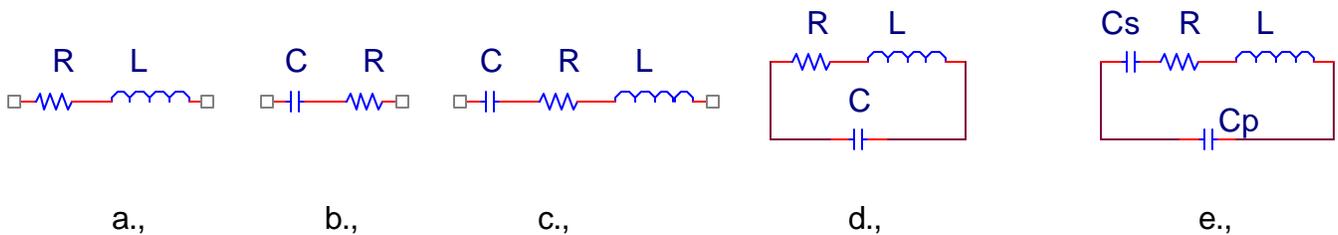


Figure 6. Simple equivalent circuits of PDN components.

Series C-R-L equivalent circuits can be applied to bare plane pairs (from very low frequencies up to close to the first modal parallel resonance), and to many bypass capacitors. Parallel C-R-L equivalent circuits can be applied to shorted plane pairs, and to capacitors mounted on plane pairs (around their parallel resonance). The parameter extractions can be built into spreadsheet calculations that capture the measured data. Under circumstances when a frequency independent R-L or C-R model is sufficient, the extraction is straightforward and accurate. When series or parallel C-R-L models have to be used, and/or when the components show non-negligible frequency dependency, an iterative solution is necessary to obtain the values for all of the equivalent-circuit elements.

III. 1. Extracting capacitance

Assume we have to measure a single piece of capacitor, or a one-dimensional transmission line (PCB trace) or two-dimensional transmission line (PCB parallel planes). Below their first series resonance frequencies, the measured impedance will follow the $Z_c=1/(2*\pi*f*C)$ trend. We can then reverse-calculate the capacitance from the imaginary part of impedance. As we need the phase information of the measured impedance, a minimum of through calibration is necessary before taking data (full two-port calibration is necessary for best accuracy). The capacitance estimate becomes:

$$C = -\frac{1}{2\pi f \text{Im}\{Z_{DUT}\}} \tag{7}$$

where f is frequency
 Z_{DUT} is the measured impedance of device.

Note that for low-loss capacitors, where the phase angle of the impedance is very close to -90 degrees, $\text{Im}\{Z_{DUT}\}$ can be replaced with $\text{Magnitude}\{Z_{DUT}\}$, and in that case there may be no need for calibration as long as $Z_{DUT} \ll 25\text{ohms}$.

Figure 7 shows the extracted equivalent capacitance of the same 1.0"x0.14" bare plane pairs, the impedance plot of which is shown in Figure 5. The graph on the left gives the equivalent capacitance calculated from the uncorrected impedance value using (3). Because the uncorrected impedance saturates at low frequencies at 25 ohms, it translates to a monotonically increasing equivalent capacitance with decreasing frequency. The graph on the right shows the capacitance extracted from the corrected impedance using (7) with the complex inversion formulas of (5) and (6). Note that at low frequencies, as the impedance becomes much higher than the 25-ohm normalization impedance, the extracted capacitance still becomes somewhat noisy, and shows the early signs of saturation. Apart the noise and slight saturation tendency below 2MHz, there is a definite negative slope on the data trace, which indicates the frequency dependence of the dielectric constant of the printed-circuit-board material.

Figure 8 is an example on a low-frequency bulk capacitor, showing its impedance magnitude and phase, as well as the extracted capacitance. Because in the given frequency range the impedance stays $\ll 25$ Ohms, there is no noticeable difference whether we plot the impedance from (3) and (7) or (6) and (7). However, in either case, the extracted capacitance has a sharp pole as frequency approaches the series resonance frequency.

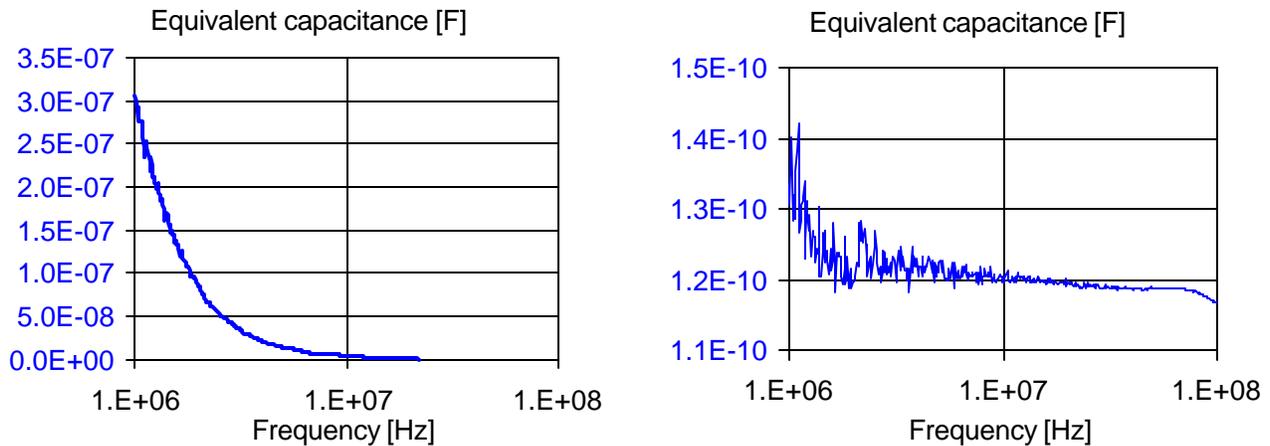


Figure 7. Extracted capacitance of a1.0”x0.14” bare plane pair, the impedance plot of which is shown in Figure 5. The graph on the left gives the equivalent capacitance calculated from the uncorrected impedance value using (3). The graph on the right shows the capacitance extracted from the corrected impedance using (6) and (7). Note the zoomed vertical scale to show the details.

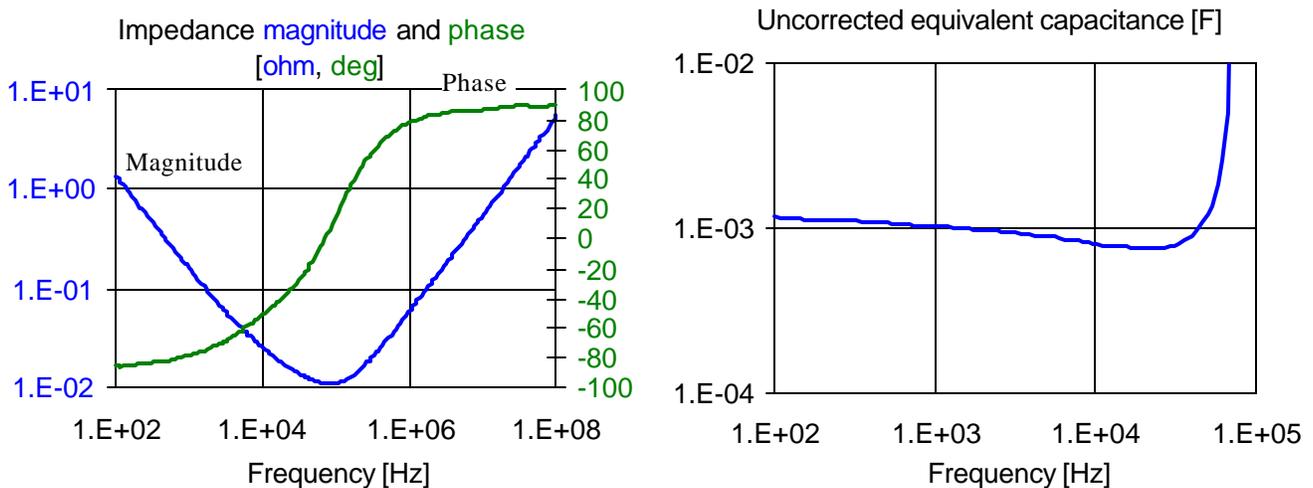


Figure 8. On the left, measured impedance magnitude and phase of a 1200uF polymer electrolytic capacitor. Impedance magnitude axis is on the left; phase axis is on the right. The right-hand graph shows the extracted capacitance versus frequency using (6) and (7).

III.1.1. Compensating for series inductance

As shown on the right-hand graph of Figure 8, the capacitance value given by (7) will show a gradually increasing error as we approach the series resonance frequency of the device. Note the sharp rise of capacitance just before 100kHz. As long as the inductance in the equivalent circuit can be estimated with reasonable accuracy, a correction can be applied to (7). With an inductance estimate of L , the corrected capacitance can be expressed as:

$$C = -\frac{1}{2\pi f(\text{Im}\{Z_{DUT}\} - 2\pi fL)} \tag{8}$$

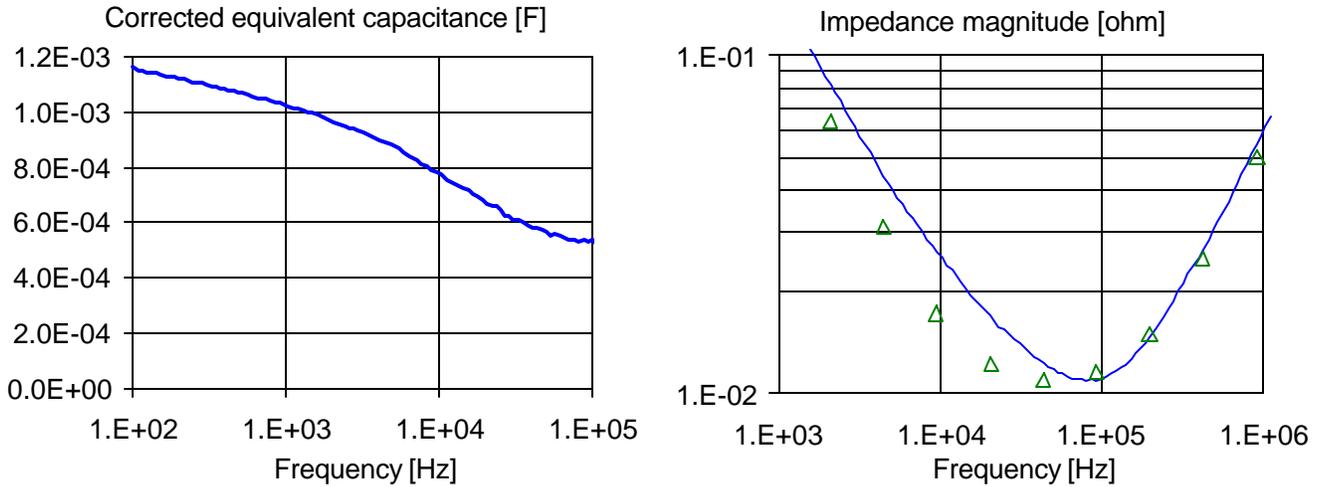


Figure 9. On the left, extracted capacitance of the 1200uF polymer electrolytic capacitor, using (8) with an estimated inductance of $L=9.4E-9$ H. On the right, impedance magnitude on zoomed scale: continuous line is measured impedance, triangles show the approximation with a series R-L-C circuit, with frequency independent values of $C=1.2E-3$ F, $R=0.01$ ohm, and $L=9.4n$ H.

Note in Figure 9 the frequency dependence of the extracted capacitance, and the large error between the measured impedance plot and simulated impedance assuming frequency-independent capacitance value. A frequency-dependent capacitor model will be shown later.

III.2. Extracting Equivalent Series Resistance (ESR)

Any one-port complex linear time-invariant circuit can be approximated with a single impedance or admittance at a given fixed frequency. In circuit theory, Equivalent Series Resistance (ESR) is the real part of the equivalent impedance. The equivalent impedance and its elements are, in general, frequency dependent. For capacitors, ESR usually means the value of impedance magnitude minimum as a function of frequency. Assuming dominant series losses in the component, it equals the real part of impedance at the series resonance frequency. If the real part of impedance varies with frequency, ESR will also depend on the connection inductance, as the loop inductance together with the capacitance determine the series resonance frequency.

ESR values many times are much smaller than 25 ohms; therefore in simple measurements there is no need to solve for the complex voltage divider. Dependent on the component and connection geometries, some parts will exhibit significant variation of equivalent series losses, because at higher frequencies the internal current path may change significantly [4]. Figure 10 shows the impedance magnitude and real part as a function of frequency for a 1200uF polymer electrolytic capacitor.

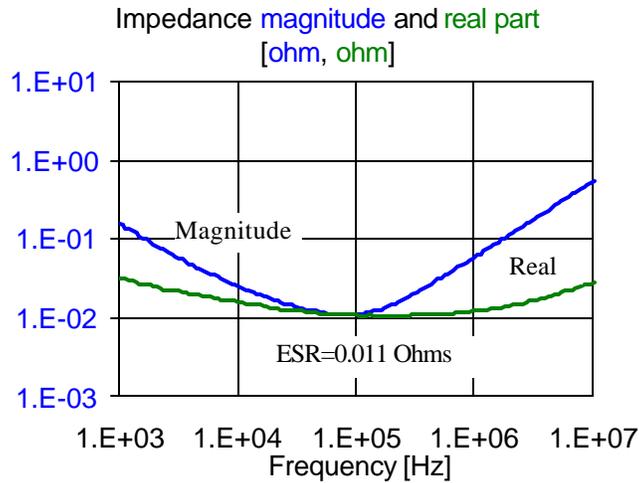


Figure 10. Impedance magnitude and real part of impedance of a 1200uF polymer electrolytic capacitor.

III. 3. Extracting inductance

Shorted traces and power/ground planes, as well as capacitors above their series resonance frequencies become inductive. The measured impedance will follow more or less the $Z_L=2*\pi*f*L$ trend. We can then reverse-calculate the inductance from the imaginary part of impedance. As we need the phase information of the measured impedance, a minimum of through calibration is necessary before taking data. For higher accuracy, full two-port calibration is preferred. The inductance estimate becomes:

$$L = \frac{\text{Im}\{ Z_{DUT} \}}{2\pi f} \tag{9}$$

where f is frequency

Z_{DUT} is the measured impedance of device.

Note that for low-loss inductors, where the phase angle of the impedance is close to 90 degrees, $\text{Im}\{Z_{DUT}\}$ can be replaced with $\text{Magnitude}\{Z_{DUT}\}$, and in that case there may be no need for calibration as long as $Z_{DUT} \ll 25\text{ohms}$.

III. 3. 1. Compensating for series capacitance

Figure 11 shows the extracted equivalent inductance of a 1200uF polymer electrolytic capacitor, shown in Figures 8 through 10. Similar to the extracted capacitance, the extracted inductance value will also show a false frequency dependency, an increasing negative error, as frequency approaches the series resonance frequency. As shown on the left-hand graph of Figure 11, the inductance value given by (9) will show a sharp decrease around 100kHz. As long as the capacitance at the series resonance frequency in the equivalent circuit can be estimated with reasonable accuracy, a correction can be applied to (9). With a series-capacitance estimate of C_s , the corrected inductance can be expressed as:

$$L = \frac{\text{Im}\{ Z_{DUT} \} + \frac{1}{2\pi f C_s}}{2\pi f} \tag{10}$$

The graph on the right of Figure 11 was calculated with (10) on the same measured data, using $C_S=5.5E-4$ F. Note that instead of the nominal 1200uF capacitance, the $C_S=5.5E-4$ F capacitance value at the series resonance frequency was used. This simple correction extends the valid frequency range of the extracted inductance by at least half a decade.

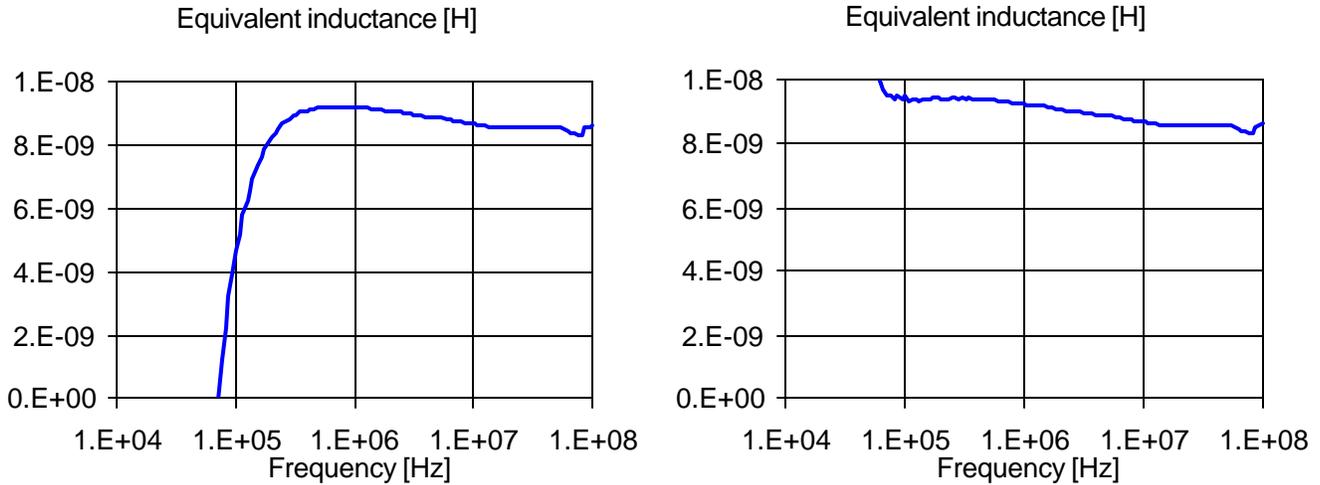


Figure 11. On the left, extracted equivalent inductance of a 1200uF polymer electrolytic capacitor, using expression (9). On the right, corrected inductance on the same set of data, using (10) with $C_S=5.5E-4$ F. Note that for the particular capacitor, the capacitance was found to be a strong function of frequency (see Figure 9): expression (10) requires the capacitance value around the series resonance frequency.

Note that the equivalent inductance after correcting for the series capacitance still shows a frequency dependency, a negative slope, but this is now real: it tells us that as frequency increases from 100kHz to 100MHz, the current loop changes such that inductance drops from 9.5nH to 8.5nH.

III. 3. 2. Compensating for parallel capacitance

In case of shorted planes, the inductance of short is in parallel to the static plane capacitance, which creates a parallel C-R-L circuit, similar to that of Figure 6.d. The extracted inductance using (9) again will show an increasing error as frequency approaches the parallel resonance frequency. Having an estimate on the C_P shunt capacitance, the extracted inductance can be corrected to take into account the susceptance of C_P :

$$L = \frac{\text{Im}\{Z_{DUT}\}}{2\pi f (1 + 2\pi f C_P \text{Im}(Z_{DUT}))} \tag{11}$$

Figure 12 shows the measured impedance of a 1.0”x0.144” plane pair with short over the capacitor pads.

Note that we have a similar equivalent circuit when a capacitor is connected to PDN planes, and the capacitor’s inductance creates a parallel L-C circuit with the static plane capacitance.

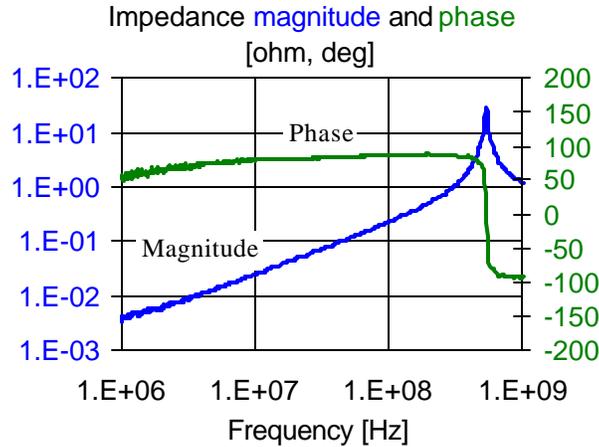


Figure 12. The impedance plot shows the impedance magnitude and phase of the small test board shown in Figure 5, measured at the test points, with the capacitor pads shorted.

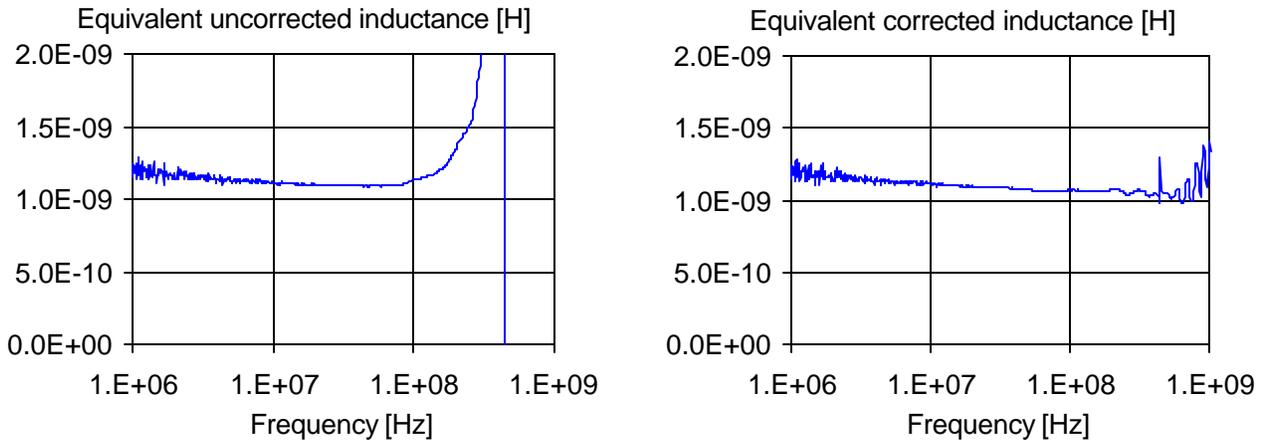


Figure 13. Extracted inductance of the 1.0"x0.14" plane pairs (shown in Figure 5) with short on the capacitor site. On the left: calculated from (9) with no correction for parallel plane capacitance. On the right: corrected for plane capacitance, calculated from (11).

III. 3. 3. Compensating for series and parallel capacitances

In case the high-frequency plane short comes from a capacitor mounted on the planes, the inductance of the capacitor is in parallel to the static plane capacitance, and we end up with a series-parallel C_S -R-L- C_P circuit, similar to that of Figure 6.e. The extracted inductance using (9) will show an increasing error close to both the series and parallel resonance frequencies. Having an estimate on the C_S and C_P capacitances, we can combine the corrections of (10) and (11):

$$L = \frac{\text{Im}\{Z_{DUT}\} + \frac{1}{2\pi f C_S}}{2\pi f (1 + 2\pi f C_P \text{Im}(Z_{DUT}))} \quad (12)$$

Note that this correction assumes that the series and parallel resonance frequencies are widely separated.

III. 4. Extracting parameters of PCB plane pairs

The procedure outlined below assumes rectangular planes, uniform and homogenous cross section and materials, and also assumes that the \mathbf{a} and \mathbf{b} dimensions of the planes (see Figure 14) can be obtained with sufficient accuracy either from the board file or from mechanical measurements on the finished board. The same procedure can be applied to separate plane pairs paired up and connected in parallel by vias, as long as the dielectric material in all pairs is the same. In this case the s separation will be the parallel equivalent of the individual laminate thickness values. For two parallel plane pairs with s_1 and s_2 plane separations, the equivalent plane separation is:

$$s = \frac{1}{\frac{1}{s_1} + \frac{1}{s_2}} \quad (13)$$

There are two parameters, however, the dielectric constant and the plane separation, which on a finished board cannot be measured directly without destructive probing.

To obtain the dielectric constant and the plane separation, we can use the following equations:

$$C_p = \mathbf{e}_0 \mathbf{e}_r \frac{ab}{s} \quad f_{res} = \frac{1}{2a\sqrt{\mathbf{e}_0 \mathbf{e}_r \mathbf{m}_0}} \quad (14)$$

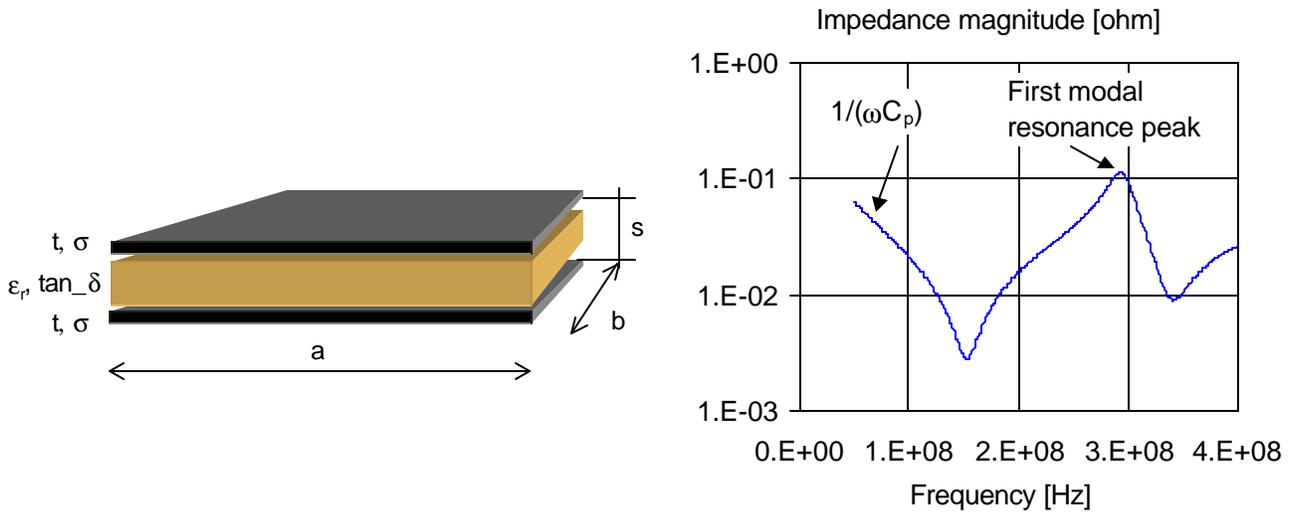


Figure 14. Parallel plane pair with dimensions and material constants on the left. On the right: impedance magnitude versus frequency plot, identifying the two parameters in (14).

By rearranging (14), we get:

$$\mathbf{e}_r = \left(\frac{1}{4a^2 f_{res}^2 \mathbf{e}_0 \mathbf{m}_0} \right) s = \mathbf{e}_0 \mathbf{e}_r \frac{ab}{C} \quad (15)$$

If conductors and dielectrics were all ideally lossless, the f_{res} first modal resonance frequency could be obtained by probing the plane pair almost anywhere. However, with conductive and dielectric losses, the impedance profile, and the modal resonance peaks do depend on the location over the planes. Furthermore, the peaks are not unique any more: the frequencies where the impedance magnitudes are the largest are not the same frequencies where the impedance of the phase is zero.

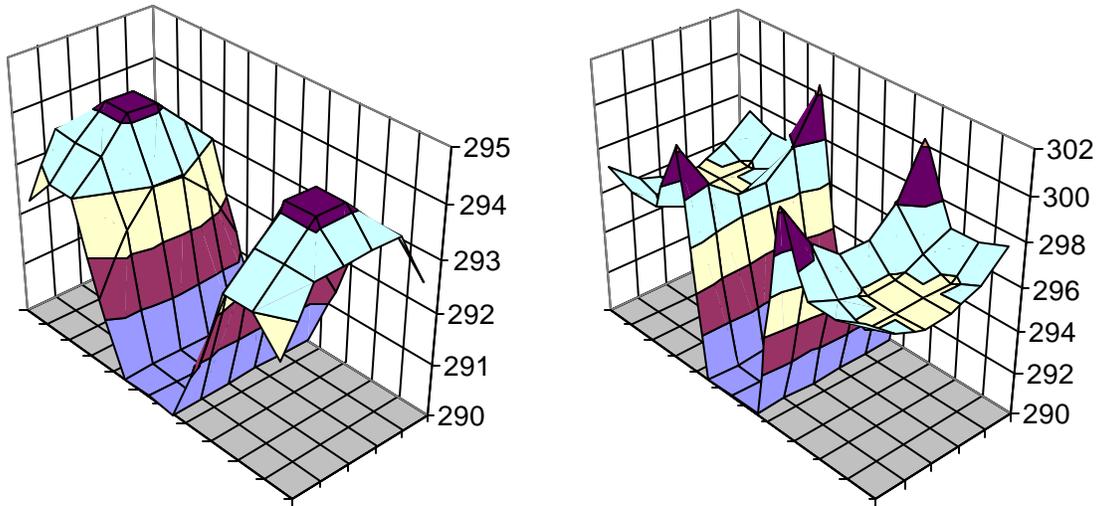


Figure 15. Extracted first modal resonance frequency from the simulated impedance of a lossy pair of planes with lossy dielectric. The graph on the left shows the frequencies where the impedance magnitude is the highest. The graph on the right shows the frequencies where the phase of impedance crosses zero. The gap in the middle reflects locations where the modal resonance is suppressed by the 2:1 aspect ratio. The floors of the charts represent the surface of the planes. Vertical scales: frequency in MHz.

Figure 15 shows the extracted first modal resonance frequency of a lossy pair of FR4 planes with one ounce copper on either side, plane dimensions $a=10''$, $b=5''$. Note that the frequency extraction is not unique, the values depend on the location on the planes. If, instead of impedance magnitude peak or phase zero crossing, we define the resonance frequency where the phase derivative has its extremum, the extracted frequency becomes unique at all locations where the modal peak is not suppressed.

Figure 16 shows the same pair of planes with the first derivative of phase plotted on the left at one given location on the planes, and the first modal resonance extracted from the phase derivative extremum (in case of first modal resonance: minimum) over the surface of the planes.

As a summary, for plane pairs where the losses are not negligible (increasingly the case with thinner laminates), the first modal frequency should be extracted from the extremum of the phase derivative. Note, however, that phase itself is already more noisy in measurements than magnitude. The derivative of the phase becomes even noisier, due to the high-pass nature of the derivative process. Therefore measuring the plane parameters based on this procedure requires averaging and narrow measurement bandwidth to sufficiently suppress noise.

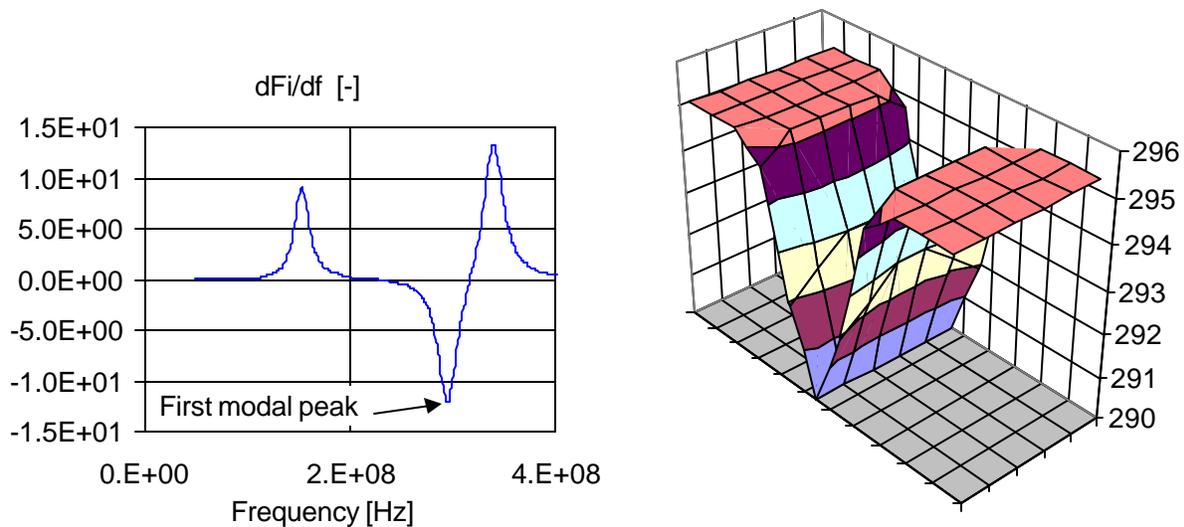


Figure 16. Extracted first modal resonance frequency from the simulated impedance of a lossy pair of planes with lossy dielectric. The graph on the left shows the first derivative of frequency at the corner of the planes. The graph on the right shows frequencies where the first derivative of the phase has its extremum. Vertical scale is frequency in MHz. The floor of the graph represents the surface of the planes.

IV. Calibrations and probes

Calibration is needed to remove the errors associated with the VNAs directional couplers and source flatness errors, as well as any error associated with cable and probe attenuations and mismatches.

Figure 17 shows the source flatness of three different VNA models. Model HP4395 covers the 10Hz-500MHz frequency range, and it does not have built-in directional couplers. The RF source and tracking receiver inputs are accessible at the front panel. Model HP4396 covers the 100kHz-1800MHz frequency range. Similar to model 4395, it also has direct access to the RF source and tracking receiver inputs. As an option, Transmission/Reflection kits or S-parameter test kits can be attached to this model, to measure S parameters directly. Model HP8720 covers the 50MHz to 20GHz frequency range with a built in directional coupler. This model measures S parameters directly, with no access to the RF source and tracking receiver input.

Note that at low frequencies the source flatness can be within one or two dB and cable attenuation can also be only a fraction of a dB. The direct access to the RF source and receiver input makes it easy to do quick transfer measurements with no calibration at all, or only with a simple through calibration. Higher-frequency VNA models with their built-in directional couplers exhibit more frequency dependency; therefore on those models calibration is always necessary.

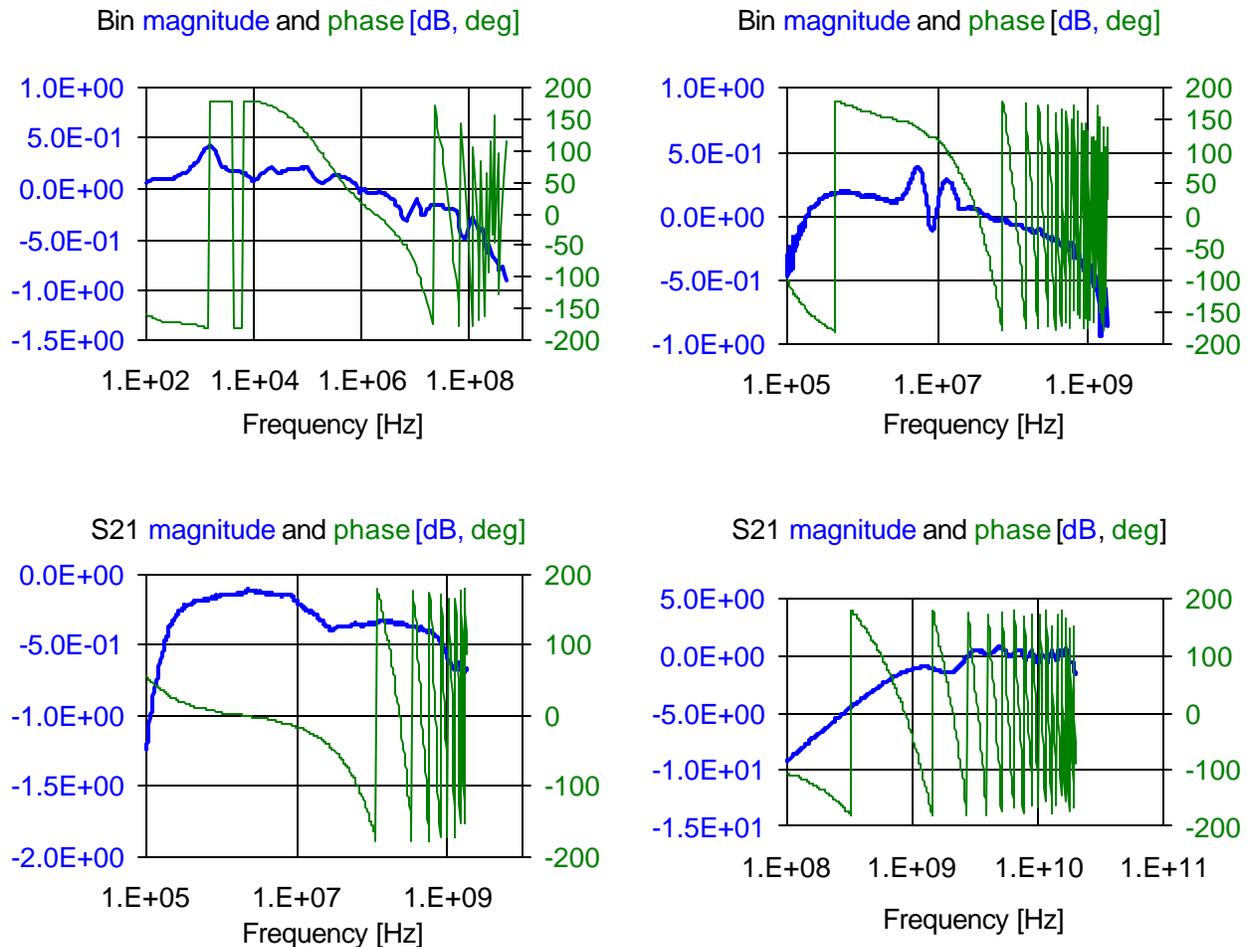


Figure 17. Measured source flatness of three different VNA models. On top left: HP4395 low-frequency VNA, RF out to B input. Top right: HP4396 mid-frequency VNA, RF Out to B input. Lower left: HP4396 mid-frequency VNA with HP85046 S parameter kit, uncalibrated S21. Lower right: HP8720 high-frequency VNA, uncalibrated S21. Note the different vertical and horizontal scales for the graphs.

IV. 1. No calibration

At low and mid frequencies, where the source flatness of the VNA is good, quick measurements can be done with no calibration at all. DUT impedance magnitude can be obtained with reasonable accuracy as long as its value is much less than 25 ohms.

At 1MHz, the wavelength in cables is 150 meters or more (depending on the dielectrics in the cable), so cables in a lab environment are usually much shorter than the wavelength. This means that mismatches in the cables and VNA-port impedances will result in a small constant error in the result, as opposed to the frequency dependent ripple error when cable length is longer than the wavelength.

We still have to watch though for cable-braid ground-loop errors (explained later in V.1). Figure 18 shows the measured impedance profile of the same 1200uF polymer capacitor that was shown earlier. The data was taken with the 4395 VNA with no calibration, no isolation transformer or isolation amplifier. The measured data may appear to be correct, but a reference measurement with the same

setup on a short (shown in Figure 18 on the right), reveals that the ground-loop error at low frequencies make the measured data questionable.

Note also that the phase information is totally irrelevant due to the lack of calibration; therefore capacitance and inductance could be extracted only from the impedance magnitude, further limiting the accuracy of data. Note, however, that measuring a smaller-value ($\leq 100\mu\text{F}$) and higher-ESR (say aluminum electrolytic) capacitor would be sufficiently accurate in this frequency range without any calibration.

For low-frequency measurements, there is little concern about probe size. Often times the probes can be replaced with short-pigtail soldered coaxial-cable connections. For reproducible measurements, a small fixture is recommended (see later).

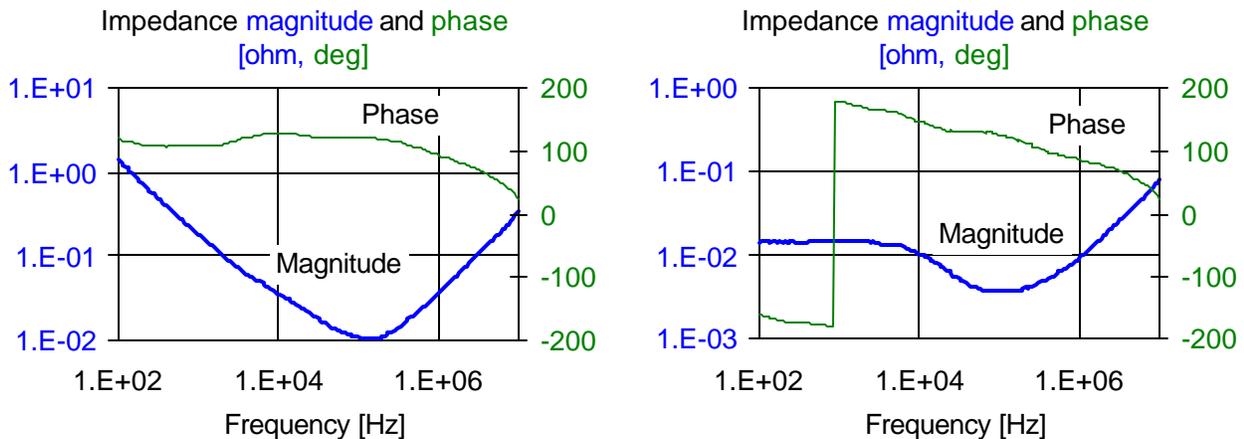


Figure 18. On the left: impedance profile of a 1200uF polymer capacitor, taken with no calibration. On the right: impedance reading with the same uncalibrated setup, with the capacitor location shorted (note the different vertical scales). Around 10kHz, the impedance reading of the capacitor is only three times bigger than the impedance reading of the shorted fixture, creating a visible ‘dent’ of the capacitor’s impedance profile around 10kHz. For such large capacitance and low-ESR parts, the cable-braid loop error has to be removed, as shown in V.1.

IV. 2. Through calibration

If cable and probe reflections can be neglected, errors of source flatness and cable/probe attenuation can be partly removed just by doing a simple through calibration. The through calibration establishes the full-scale impedance value and the phase reference. Inversion of the complex voltage divider becomes possible, and expressions (5) and (6) can be used to obtain the real and imaginary parts of the unknown impedance without the $\ll 25$ ohms restriction. Depending on the frequency range, and quality of cables/probes we use, the still uncalibrated impedance mismatches may show up at higher frequencies in form of an increasing ripple on the measured data. The left-side graph of Figure 19 is an illustration of the ripple error we can expect at high frequencies with having only through calibration. The right-hand side graph shows the same DUT measured at high frequencies with probes after full two-port calibration. Note the different horizontal scales: the graph on the left spans 100kHz-1GHz; the series resonance frequency of the capacitor is seen at 20MHz. The graph on the right spans 100MHz to 10GHz, in which frequency range we see a 600MHz parallel resonance of capacitor’s inductance and plane capacitance, followed by three plane resonance peaks.

The DUT in this case was a 0508-size ceramic capacitor mounted on a small square plane pair.

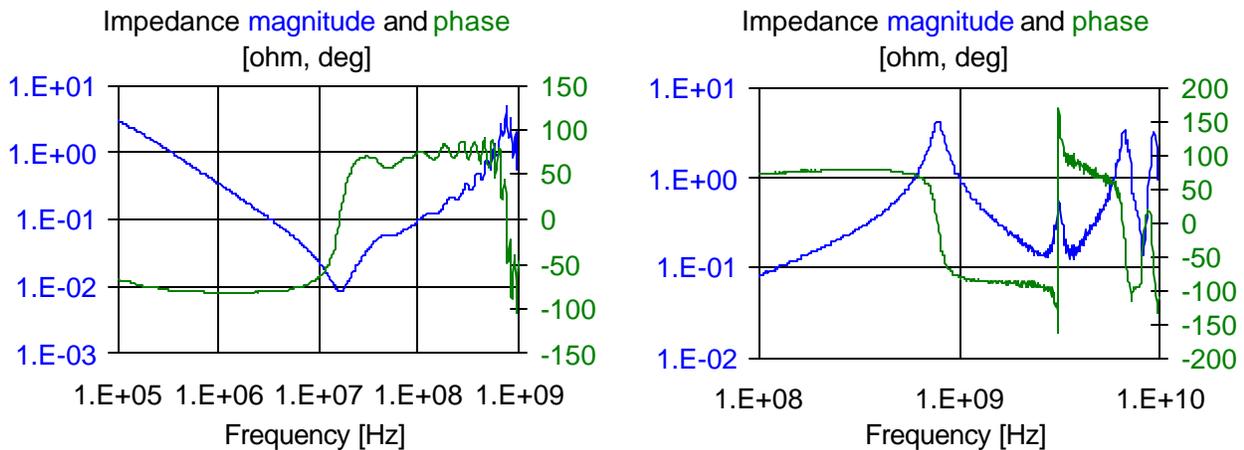


Figure 19. Measured impedance magnitude and phase of a ceramic capacitor soldered on a small test board. Low-frequency response with only through-calibration is shown on the left, high-frequency response of the same DUT with full two-port calibration is on the right.

IV. 3. Full two-port calibration

To probe PCB through-hole test points, short semirigid probes can be used to connect the VNA to the DUT. Figure 20 shows a simple home-made probe with 50-mil pin spacing, made of standard 0.082” diameter copper-sleeve semirigid coaxial cable. For maximum high-frequency isolation, there is a ferrite sleeve on the coaxial cable. The center wire of the coaxial cable serves as the hot pin. A short wire, soldered to the copper sleeve, is the ground pin. The pin-to-pin spacing is 1.25mm (50 mils).

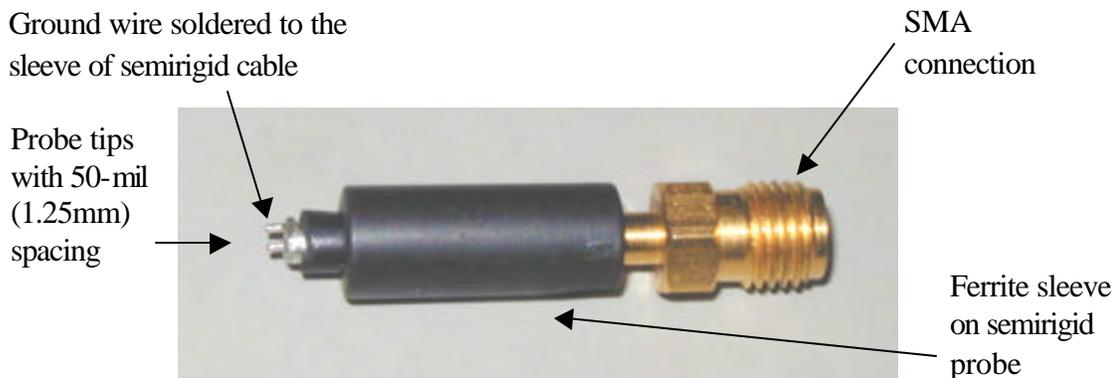


Figure 20. Short semirigid coaxial probe with SMA connection.

With semirigid coaxial probes, full two port calibration can be done up to the end of the two coaxial cables leading to the probes. Standard SMA open, short and load elements can be used for the reflection calibrations. For the through-calibration part, a semirigid SMA-SMA cable is used with a total length equaling the two semirigid probes. The only uncalibrated error is due to the impedance mismatch in the short semirigid probes, but mismatch and attenuation of the connecting cables are all removed.

With miniature two or three-pin probes, such as PicoProbes from GGB Industries, full two-port calibration can be done with the calibration kit available on ceramic substrate. Figure 21 shows probes with 450-micrometer pin pitch. The probes are positioned over the shorting pads of the calibration substrate. With the calibration substrate, calibration can be performed up to the tips of the probes.

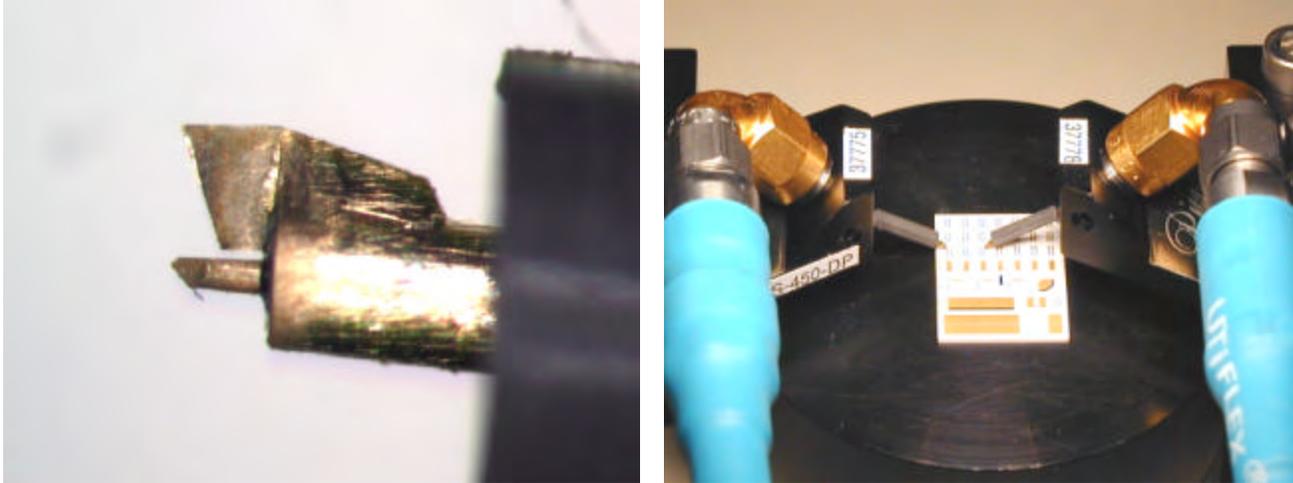


Figure 21. S-G 40-mil (0.45mm) PicoProbe is shown on the left. Photo on the right: two probes on ceramic calibration substrate.

V. Making the proper connections

The proper choice of calibration and probes is a necessary first step, but further considerations are needed to make sure that the connections between the VNA and probes do not impose limitations on the measurements.

V. 1. Eliminating cable-braid-loop error at low frequencies

V. 1. 1. Ferrite clamps/sleeves

The ground returns of Port1 and Port2 in the VNA are connected together inside the instruments. When we measure a low-impedance DUT, for instance an active DC-DC converter output, or a metal short, such as it is shown on the top of Figure 22, the current of Port1 will create a voltage drop across the parallel equivalent of the two cables' braid resistances. The low-frequency equivalent circuit for this scenario is shown on the bottom left-hand of Figure 22. The measured low-frequency impedance value, instead of Z_{DUT} , becomes

$$Z_{measured} = Z_{DUT} + R_{b1} \times R_{b2} \quad (16)$$

The cable-braid ground loop is not a limitation at high frequencies, because the cable inductance creates a -20dB/decade roll off of this residual reading above the corner frequency determined by the braid resistance and inductance. The inductance of the cable can be increased, and the corner frequency beyond which the error rolls off can be reduced, by placing ferrite clamps or ferrite beads around the cables. The right-hand graph of Figure 22 on the bottom shows the impedance magnitude measured

with different cable configurations, with short across the test points. All cables for this illustration were 24" long, type RG174 flexible coax and 0.082" semirigid coax. The plots show the residual error with and without ferrite clamps around the cables. The measurements were done with an HP4395 VNA, between RF Out and Input B. Figure 23 shows a coaxial cable with ferrite sleeves.

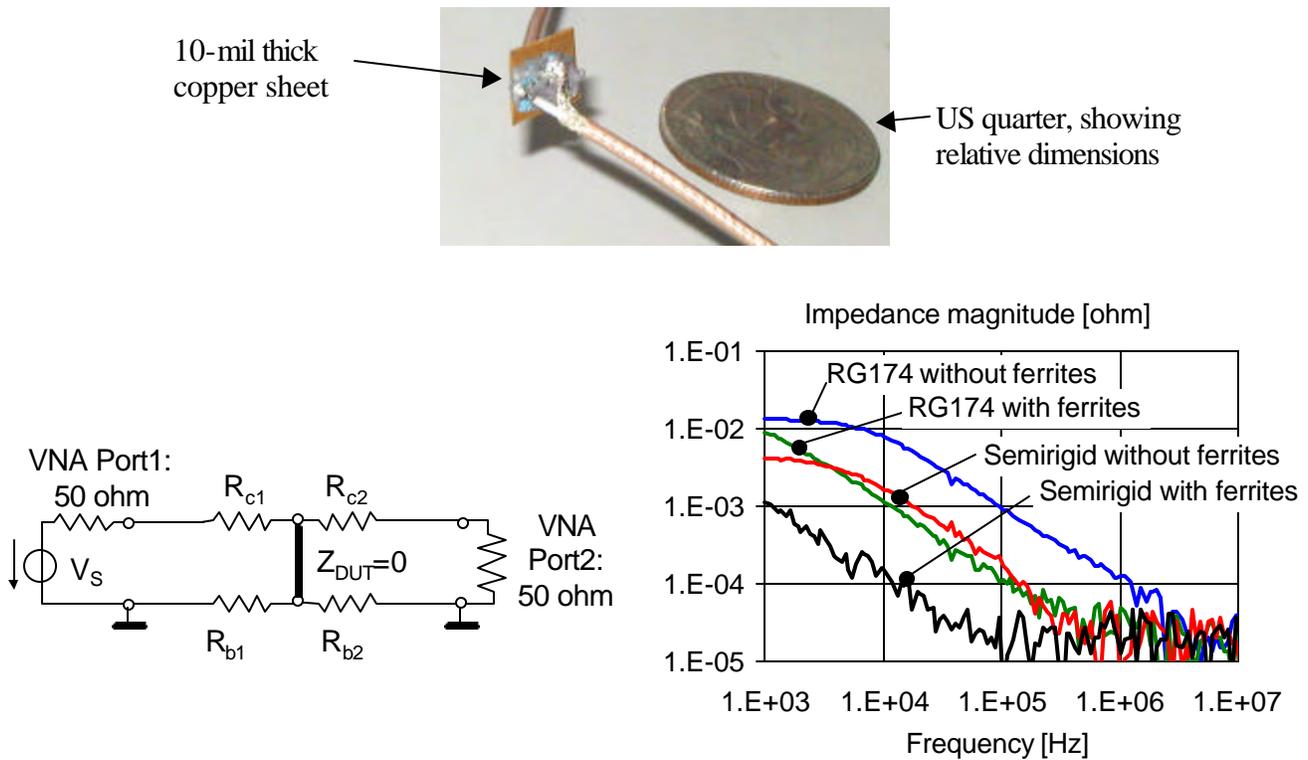


Figure 22. On top: two coaxial cables soldered to a 10-mil thick copper sheet. On bottom left: equivalent circuit of cable-braid loop. On bottom right: measured impedance magnitudes of residual error on shorts with different cables, with and without ferrite sleeves. For this illustration, no isolation transformer or isolation amplifier was used.

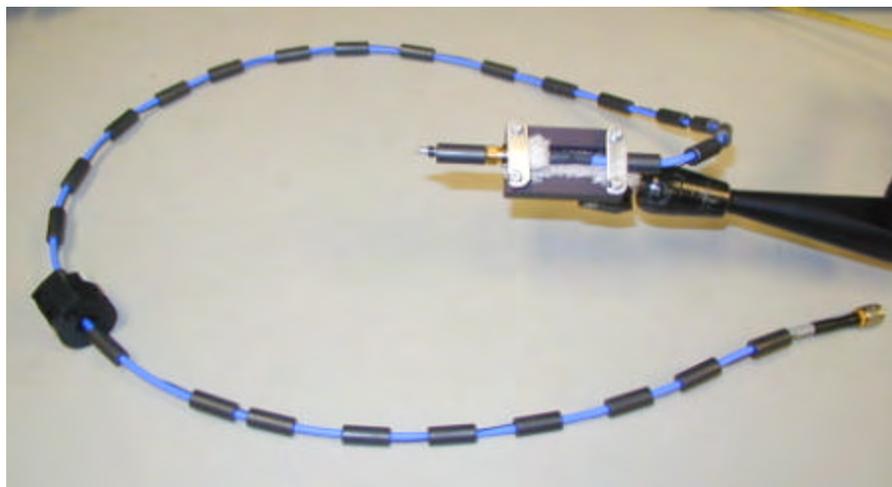


Figure 23. Flexible coaxial cable, with ferrite sleeves along its jacket. The short semirigid probe at the end of the cable is held by a positioner.

With the ferrite clamps the residual error still has a -20dB/decade roll off. This may be OK, when the measured DUT impedance follows the same pattern, namely for capacitors. It is then enough to reduce the residual error to at least 10dB below the expected impedance of the DUT. However, eliminating the cable-braid-loop error by ferrite clamps does not work in cases when we want to measure very low DC resistances combined with a low series inductance, such as shorted planes, shorted via loops and active voltage regulator modules (VRM).

V. 1. 2 Isolation transformer or amplifier

Another possible way of eliminating the ground loop of cable is to use an isolation transformer. Figure 24 on the left shows the photo of a home-made transformer, created on a Phillips ferrite toroid core, type TX51/32/19-3F3. The 52-mm ferrite core had two times 50 turns with AWG20 wires. The large-size core provides low distortion for the usual 0dBm VNA measurement level. The large size is also useful in measuring active devices, like VRMs, where the DC output current of the VRM may drive a smaller core into saturation. The number of turns is a compromise between the main inductance, which sets the lower corner frequency and the winding capacitance, which sets the upper corner frequency of the transformer response. The transformer described above had its -3dB points between 50 ohms at 500Hz and 3MHz . The main and stray inductances and the winding capacitance of the transformer create a frequency-dependent transfer response, which cannot be completely removed by VNA calibrations for a wide range of DUT impedances. Therefore the calibration should be done with a low-value standard (1 ohm or 0.1 ohms), which approximately matches the upper end of expected range of measured DUT impedances. This frequency dependent error can also be reduced by limiting the measurement bandwidth: for the given transformer, in the 1kHz to 1MHz frequency range, the residual error is negligibly low for any DUT impedance.

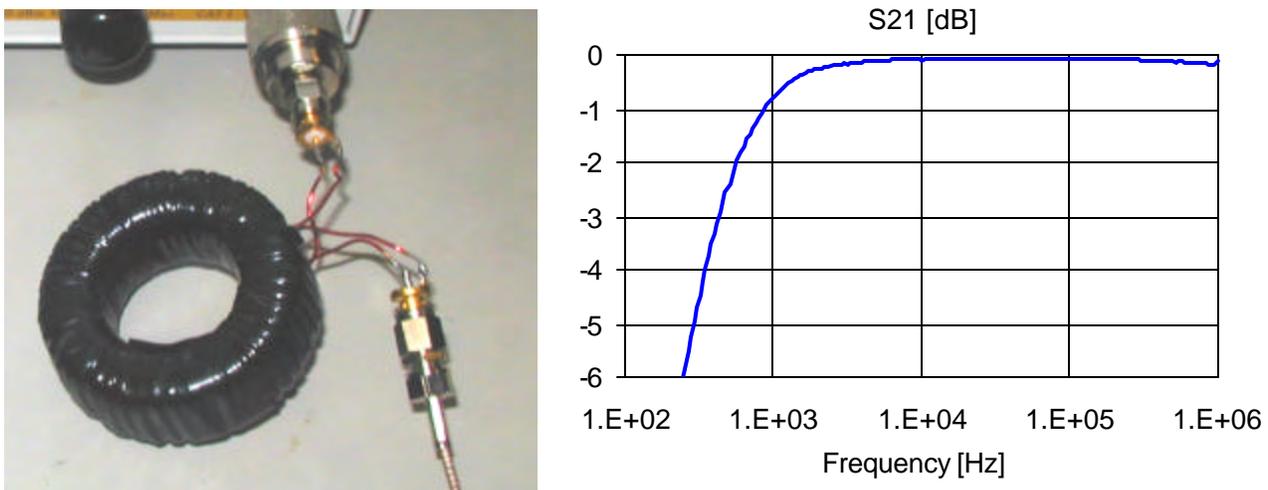


Figure 24. Isolation transformer connected to Input B of a VNA and its response, before calibration, between 50-ohm impedances.

Isolation amplifiers can also be used to break the ground loop, see Figure 25. The ground loop is broken at the floating differential input of the amplifier. It can be placed either in front of Port2 of the VNA, or at the output of Port1, before the DUT. Low-frequency response is usually not an issue in this case, as operational amplifiers are DC coupled. The high end of the frequency range is still limited by the frequency response and maximum slew rate of the amplifier. The circuit shown on the schematic had a flat frequency response with non-saturated slew rates up to a few MHz.

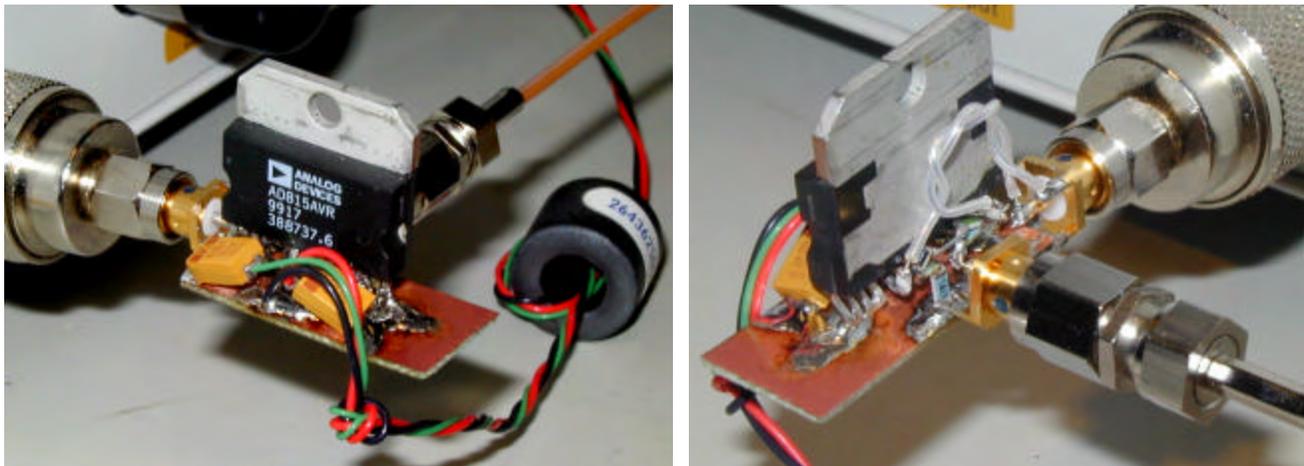
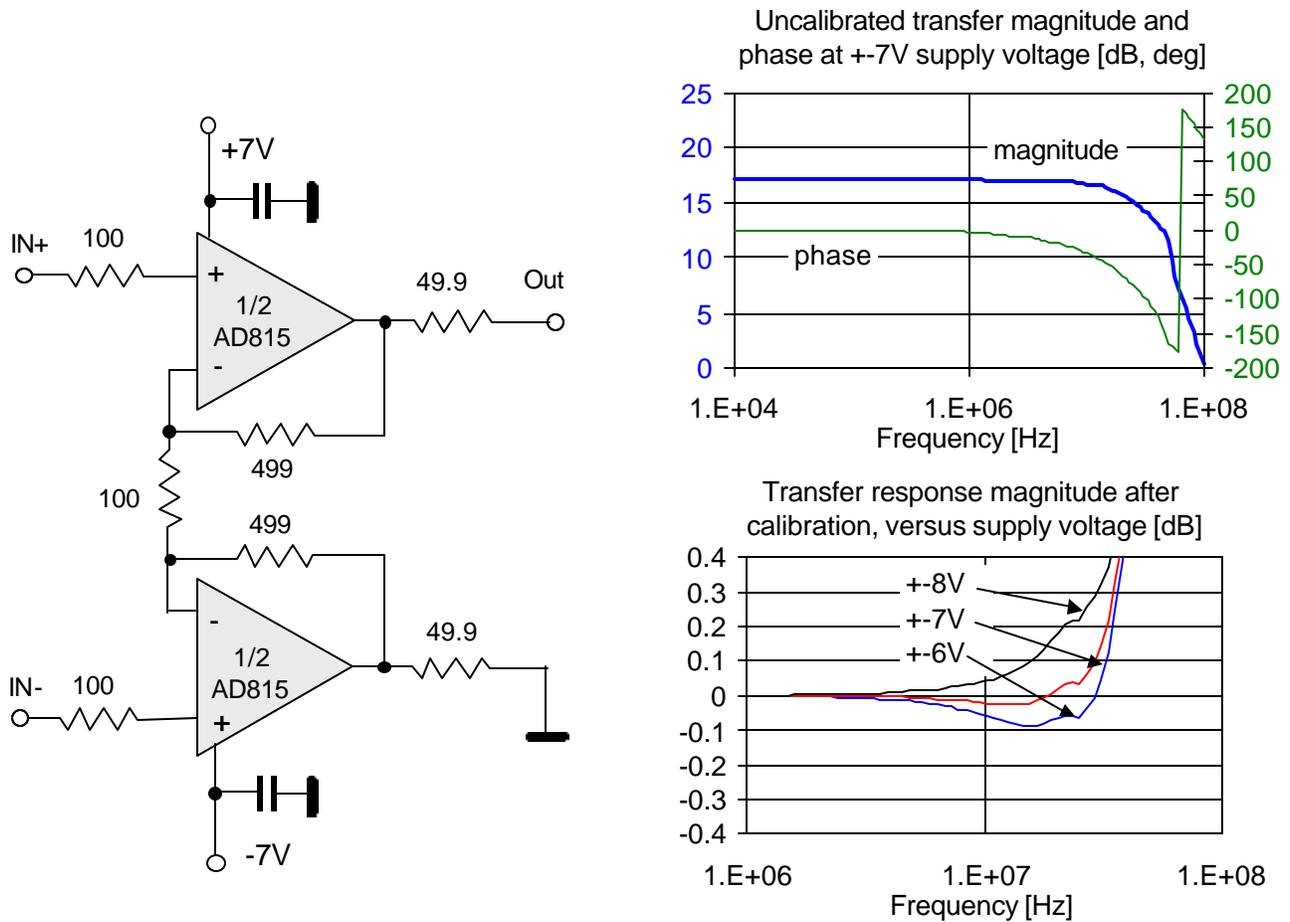


Figure 25. Differential-input single-ended output amplifier for the DC-10MHz frequency range. Top left: circuit schematic. Top right: measured transfer response at nominal supply voltage before calibration (with the VNA) and transfer response variation versus supply voltage after calibration. Bottom row: front and back photo of the amplifier, connected to the RF Out output of the 4395A VNA. Note that the three supply wires are fed through a ferrite ring. Without this isolation, there is a small (0.01 dB) variation in the transfer response with different input and output cable orientations.

The only drawback of the isolation amplifier is the need for external power supplies. The supply voltages must be selected such that the connected VNA input should survive the maximum saturated output voltage in case transients should occur. The lowest value of supply voltage is limited by the specified minimum rail voltage of the operational amplifier. With the circuit shown in Figure 25, proper operation was achieved at and above $\pm 4V$ supply voltage. For the HP4395 VNA model, the maximum input DC voltage is $7V$; this suggests a $\pm 7V$ maximum supply voltage for the differential amplifier. The output impedance of the amplifier is very close to $50\ \Omega$ in the entire $0-10\text{MHz}$ frequency range, and active DUTs (such as DC-DC converters and VRMs) can be connected to its output without any problem. Figure 25 shows the schematic and construction of the isolation amplifier, together with the frequency response and supply-voltage sensitivity curves. Note that the transfer gain is very stable below 1MHz , and varies less than $\pm 0.1\text{dB}$ at 10MHz as the supply voltage varies in the $\pm 6 \dots \pm 8V$ range. In this simple home-made setup, both the supply-voltage sensitivity and the absolute frequency response limits its application to frequencies up to 10MHz .

V. 2. Measuring low-ESR bulk capacitors

As was shown earlier, impedance magnitude of some low-ESR bulk capacitors may be so low that the cable-braid loop could limit the measurements. For lower capacitance values placing many ferrite clamps or ferrite beads around the cable sleeve may reduce the residual error sufficiently below the measured impedance values. For large capacitance values combined with low ESR, isolation transformer or isolation amplifier is needed.

The two-port VNA connection with isolation transformer or isolation amplifier helps to eliminate the measurement error in DUT connections, but connecting the probes to the device still raises questions. For through-hole parts, the probes can be soldered or pressed against the capacitor's leads. However, the actual connection points on the leads, their distance from the capacitor body, will eventually determine the measured value of ESR and inductance.

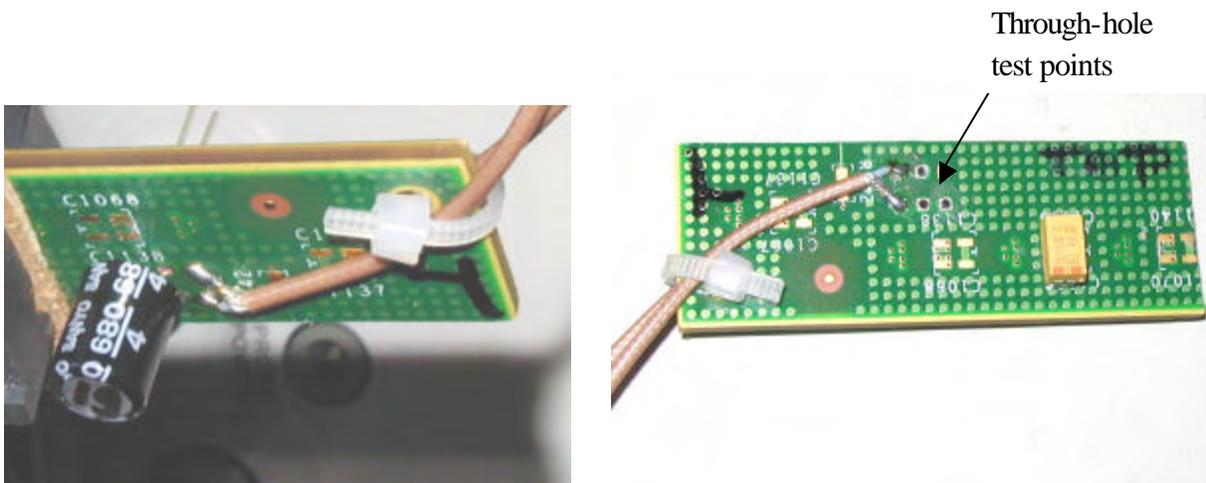


Figure 26. Small multi-layer test boards for measuring through-hole and surface-mount low-ESR bulk capacitors.

Usually we are interested in the performance of the part that we could expect in the actual environment it goes into. The capacitor is connected through vias and optional traces to planes or plane shapes, which connect the capacitor to the rest of the PDN. What we really need is the impedance presented by the capacitor to the rest of the PDN. The best way to achieve this is to create a small test board with the same stack up and PDN-plane allocation that we have in the final application. The capacitor can be soldered onto or pressed into the test board, and at a nearby via pair, the impedance can be measured.

The additional impedance of the plane connection from the probe points to the capacitor site can be established by measuring the same test board with the capacitor site shorted. Note that this also means that the inductance, and to a smaller degree, also ESR of the device will depend on the geometry of the application.

The stack up of the small printed-circuit board shown in Figure 26 has six copper layers, nominally with 0.093" total board thickness. Every copper layer is nominally one ounce. There are two thin cores with 2 mil dielectrics, 8 mils below the surfaces on either side. Only one of the plane pairs is connected to the capacitor pads and test vias. Therefore dependent on from which side we insert a through-hole part, the plane will be either 'near' or 'far'.

The setup used an isolation transformer. The transformer limited the upper bandwidth to 1MHz, and only through calibration was done with a one-ohm resistor soldered across the capacitor site. After calibration, the test boards resistance and inductance was measured by inserting a shorting wire into the through-hole test points.

The impedance plots are shown in Figure 27, separately for the short inserted from the 'near' and 'far' sides. The measured impedance profiles, separately for the same bulk capacitor inserted from the 'near'; and 'far' sides, are shown in Figure 28, with the extracted ESR and inductance values labeled. The resistance and inductance values measured on the shorted capacitor site can be subtracted from the measured capacitor data. ESR on the near side is $7.59 - 2.0 = 5.59$ milliohms. ESR on the far side: $8.47 - 2.42 = 6.02$ milliohms. Note that the difference can be attributed to the resistance of the lead-length difference between the near and far sides.

The extracted extra inductance (extra beyond the inductance corresponding to a shorting wire in the given fixture) can be calculated in a similar way. Extra inductance on the near side: $3.56 - 1.15 = 2.41$ nH. Extra inductance on the far side: $4.72 - 2.42 = 2.3$ nH. Note that this extra inductance is the same (within measurement errors) from the near and far sides, because this figure captures only the extra inductance beyond the shorting wire, which is purely determined by the capacitor body, and does not depend on which side the capacitor is inserted into.

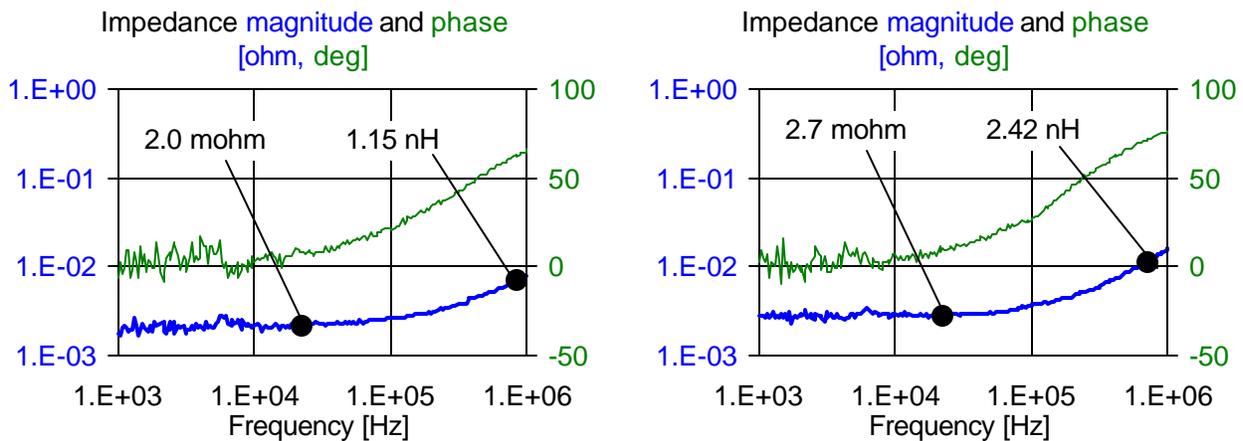


Figure 27. Impedance reading of reference short at the through-hole test points on the small test board shown on Figure 23. Graph on the left: short inserted from the 'near' side, where the plane pair is 8 mils below the surface. Graph on the right: short inserted from the 'far' side, where the plane pair is 93-8=85 mils below the surface.

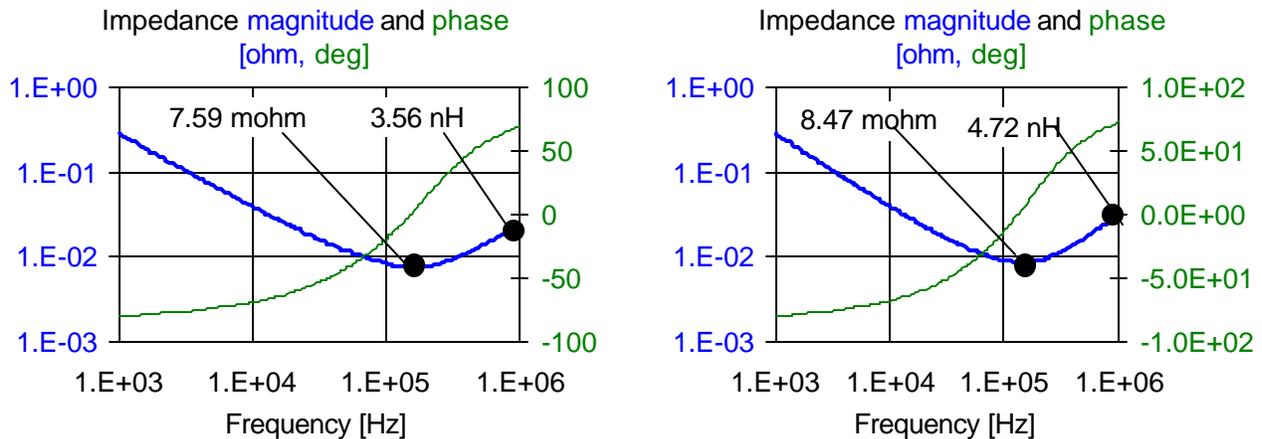


Figure 28. Graph on the left: through-hole low-ESR bulk capacitor in the small test board, inserted from the ‘near’ side. Graph on the right: through-hole low-ESR bulk capacitor in the small test board, inserted from the ‘far’ side.

V. 3. Measuring Voltage Regulator Modules

Voltage Regulator Modules (VRM) are DC-DC converters, which convert the incoming (e.g., 48V or 12V) DC to the final regulated voltages. VRMs should provide the specified DC current, and also at low frequencies, we expect them to provide low impedance. The bandwidth in which the low impedance can be maintained primarily depends on the switching speed. As a second-order effect, the loop transient response and bandwidth also depend on the style of the control loop and range of allowable load impedances. Currently the mainstream switching frequency is a few hundred kHz per phase, achieving way over a MHz effective switching frequency with multiphase converters. The output impedance in the fast point-of-load converters can be as low as a fraction of a milliohm up to at least 10kHz. VRM response, on the other hand, can be heavily nonlinear for large load-current steps. Guaranteeing the stability of control loop could also be a challenge for widely varying load currents and bypass-capacitor impedances.

For PDN-design purposes, the most convenient is if the VRM stays in the linear mode over the entire specified line and load range. In this case the simulated or measured output impedance is maintained under all operating conditions, there is no need to separately consider the large-signal conditions. One possible way of achieving this condition is to use voltage positioning with current-mode control loops [5].

Measuring the output impedance of VRM requires considerations similar to that of measuring low-ESR bulk capacitors. Because VRM control loops tend to have high gain at DC, the typical output impedance response is ‘inductive-like’, achieving very low impedances at low frequencies. Therefore the cable-braid ground loop must be opened by either transformer isolation or by differential-input amplifiers; ferrite loaded cables are usually not sufficient if we want to measure the VRM impedance below a few kHz.

If we use the differential-input single-ended output isolation amplifiers, no further precaution is necessary. If we have the amplifier in front of the VNA input, the input impedance of the amplifier is high enough that it will not interfere with the VRM output, and the differential-to-single-ended

conversion will suppress the DC output voltage of the VRM, so that the VNA input is not stressed with DC voltage. If the amplifier is placed to the Port1 output, in front of the DUT, the series 50-ohm resistor protects the amplifier's output from the DC voltage of the regulator. The proper DC load current of the VRM can be set either with a resistive load, or by using an active electronic load device.

Isolation transformers can also be used to measure an active VRM. In this case, however, we need to consider further factors. First, the transformer winding's resistance creates a DC load to the VRM. The isolation transformer shown on the previous figures has a few milliohms DC resistance, which would overload most VRM outputs if connected directly to them. To limit the DC load current, we can insert a few feet of thin (RG174 or RG178) coaxial cable, where the DC resistance of the inner wire will limit the DC load current, while the resistance of the wire is removed from the measured data by calibration. Second, unless we use two transformers, to isolate both VNA ports from the VRM, the non-isolated VNA port will get the full output voltage of the VRM. We always have to check the VNA specifications for maximum allowable DC input voltage: a few volts are usually not a problem.

Of course, if the VRM is not powered up, its output impedance can be measured similar to any low-frequency bulk capacitor.

Figure 29 shows the setup to measure the output impedance of a small-size socketed VRM. The DC input power was connected through wires on the left. The output voltage appeared at multiple pins on the socket. To maintain the low output impedance, a small double-sided thin printed-circuit board was soldered to the appropriate pins: all positive output pins were soldered to the top copper plane, all negative output pins were soldered to the bottom copper plane. The solid copper planes across a thin dielectric separation ensure that the VRM output pins are connected to the measurement points through a low-resistance and low-inductance path. The heavy isolated wire on the right shows the connection to the electronic load.

The VNA probes were soldered in the middle of the two-sided plane connection. The remote sensing wires were soldered to the same location.

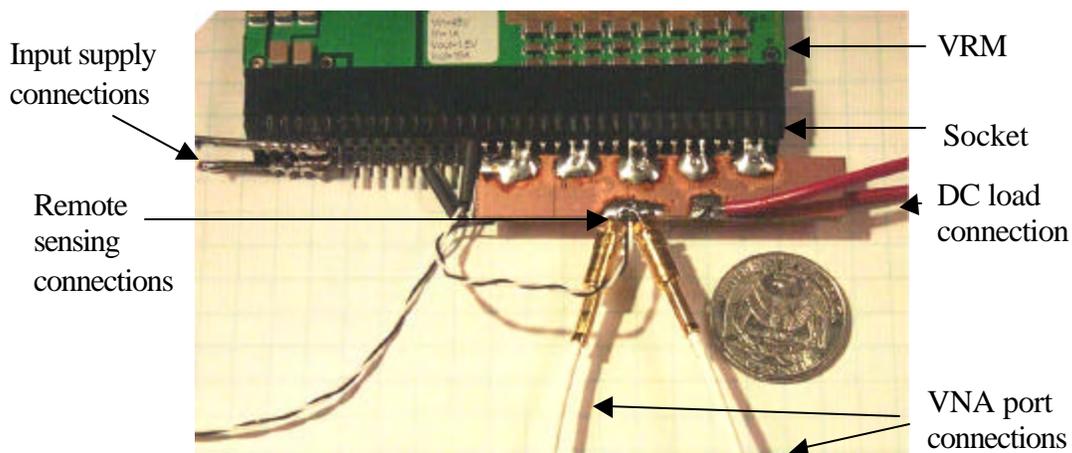


Figure 29. Setup and connections to measure VRM.

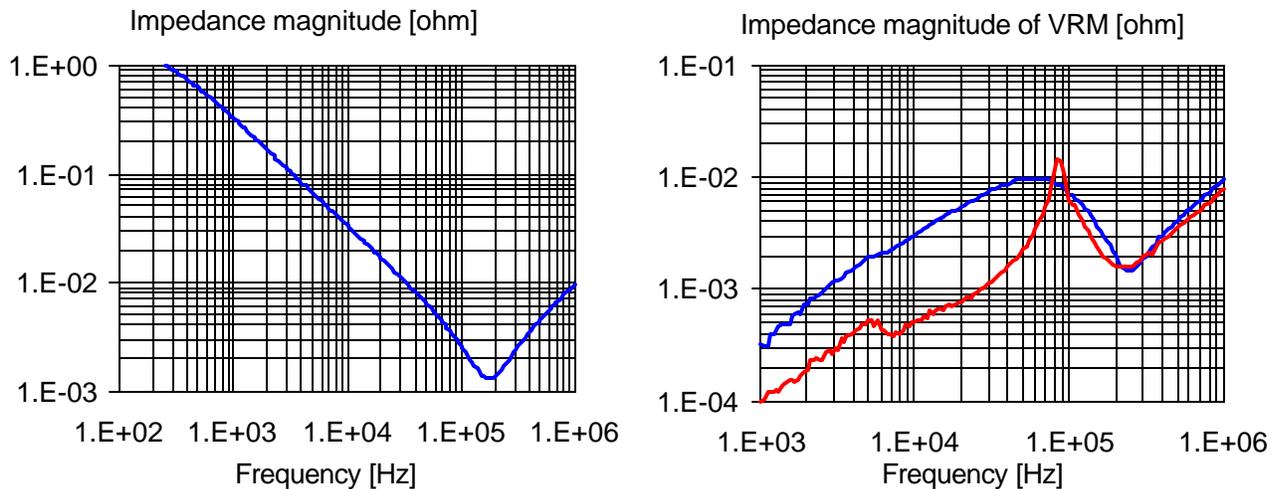


Figure 30. Measured low-frequency output impedance of a VRM with no input power (on the left) and with input power, for two different loop-compensations (on the right).

The VRM can be measured without input power (inactive loop) or powered (active loop). The left graph of Figure 30 shows the VRM’s output impedance with no input power. It follows the impedance profile of its filter capacitors at the output. Note that if the loop stability requires or assumes external capacitors, those should also be connected to the VRM output while we measure its output impedance.

On the right of Figure 30, two different versions of the same VRM are shown, with input power applied. Note the different vertical and horizontal scales for the two graphs. All data was taken with no additional bypass capacitors. The unpowered output impedance clearly follows the impedance profile of the on-board filter capacitors. The two versions of active output impedances illustrate the potential problem of instability. One of the curves maintains less than one milliohm up to 20kHz frequencies, but at 85kHz this version had a sharp impedance peak. While the magnitude of the impedance peak itself may not be a problem (it is around 15 milliohms), the sharp peak indicates very low phase margin, and it indicates the risk of self-oscillation under slightly different tolerance constellation. Separate analysis of the phase margin of the control loop indicated, in fact, low phase margin. This impedance peak can be viewed as the classic inter-capacitor ant resonance peaking, when a bypass capacitor’s inductive impedance creates a peak with another bypass capacitor’s capacitance. In this case, the inductive behavior comes from the decreasing loop gain in the VRM with increasing frequencies.

The second plot on the right-hand graph exhibits no sharp peaking. However, as the relative positions of the two plots show, unless the switching frequency is increased or the control-loop topology is changed, modifying an existing loop to increase phase margin almost always comes with an increase of output impedance below the peak frequency. In this particular case, the output impedance at 10kHz increased from 0.4 milliohms to 3 milliohms. Finally, above the peak, at about 200kHz and above, both versions’ output impedances follow the impedance curve of the output filter capacitors.

Note that the output impedance profile, while it does not provide any quantitative data about the loop stability, gives an easy and quick way to qualitatively check stability. Note also that measurement of loop stability usually requires opening up the control loop, while measuring the output impedance can be done without any modifications to the VRM. In fact, with the methodology described above, the VRM output impedance can be easily measured in a fully assembled working system.

V. 4. Measuring low impedances at high frequencies

Cable performance limits the two-port VNA measurements at low frequencies due to the cable-braid ground loop. At high frequencies, besides of the noise floor of the instrument itself, the finite surface transfer impedance and resonances in the cable braids set the error floor.

As PDN impedances are very seldom close to the 50-ohm instrument impedance, there is usually a big mismatch/reflection at the connection to the DUT. Through the finite surface transfer impedance of the cable braids, the reflections and leakage show up as erroneous peaks in the response.

The left graph of Figure 31 shows the equivalent impedance reading from two RG178 coaxial cables connected to a VNA port with their ends shorted with SMA caps. The two cables were a couple of inches apart. The same cables in the same position were also measured with different number of small ferrite clamps put around them. The graph also shows a trace with just a few ferrite clamps around the cable: the erroneous peaks are shifted toward lower frequencies and the peak values are reduced.

To achieve the lowest possible error floor, a flexible coaxial cable requires ferrite clamps or sleeves all along its lengths. The graph on the right in Figure 31 was measured across a plane short in a medium-size PCB, with the cables and setup shown in Figure 32. The few pH equivalent inductance of the error floor is achievable only if we cover not only the entire length of the cable braid with ferrite absorber sleeves, but also place ferrite sleeves around the semirigid probe and isolate the probe-holder metal parts from the cable.

There are coaxial cables available with good surface transfer impedances, which would reduce this error. Trouble is, cables with better braids and shields tend to be not only more expensive, but also more stiff. When we have to reach with two probes inside a bigger system, or to connect to a pre-mounted fixture, flexible cables are essential. Figure 33 shows a large probe station with high-performance coaxial cables. When measuring low-inductance PDN components, this setup still requires a few ferrite clamps along the cable to suppress cable resonances.

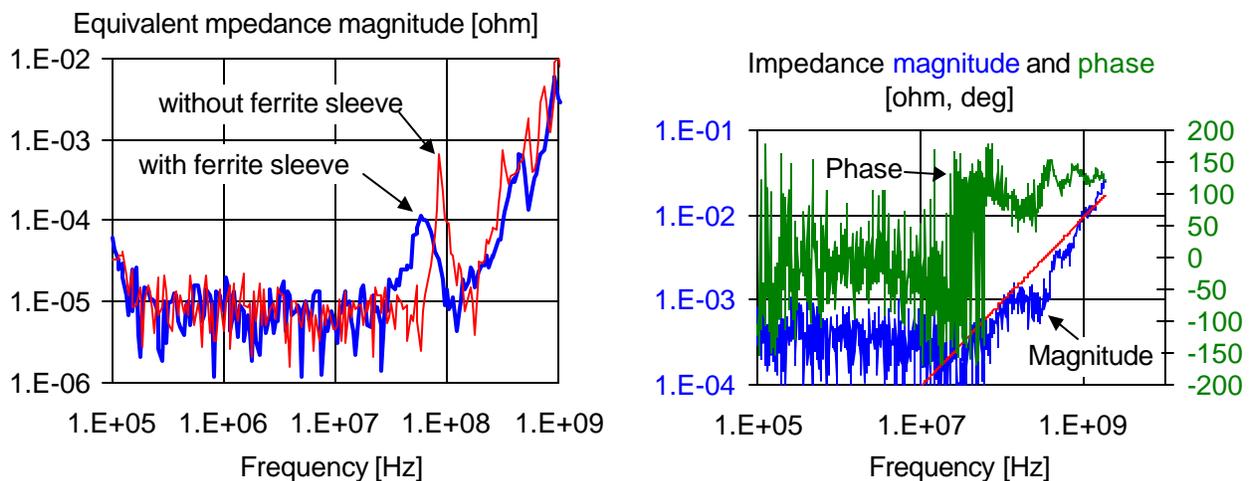


Figure 31. Illustration of cable performance at high frequencies. Left graph: equivalent impedance magnitude with two-port VNA measurements, the two cables are shorted at their ends, two inches apart, with and without ferrite sleeves. Graph on the right: residual impedance magnitude and phase measured on plane short with 50-mil semirigid coax probes

and ferrite covered coaxial cables (shown in Figure 32). The thin straight slope shows the impedance magnitude of a 1.5pH inductance.

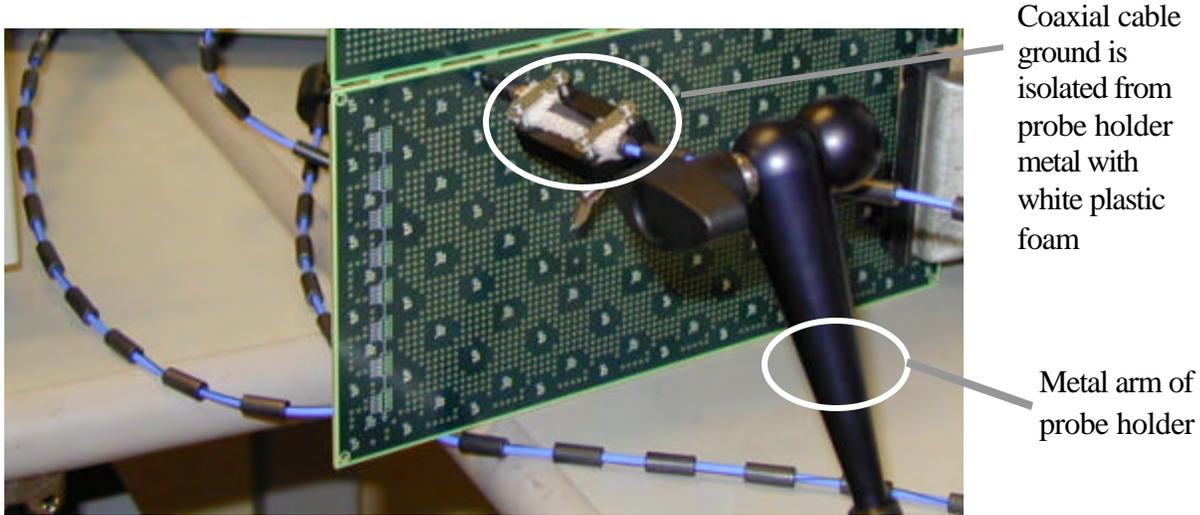


Figure 32. Ferrite-covered flexible coaxial cable, semirigid probe with ferrite sleeve. For lowest residual error, probe and probe holder metals are isolated in the probe-holder head with plastic foam.

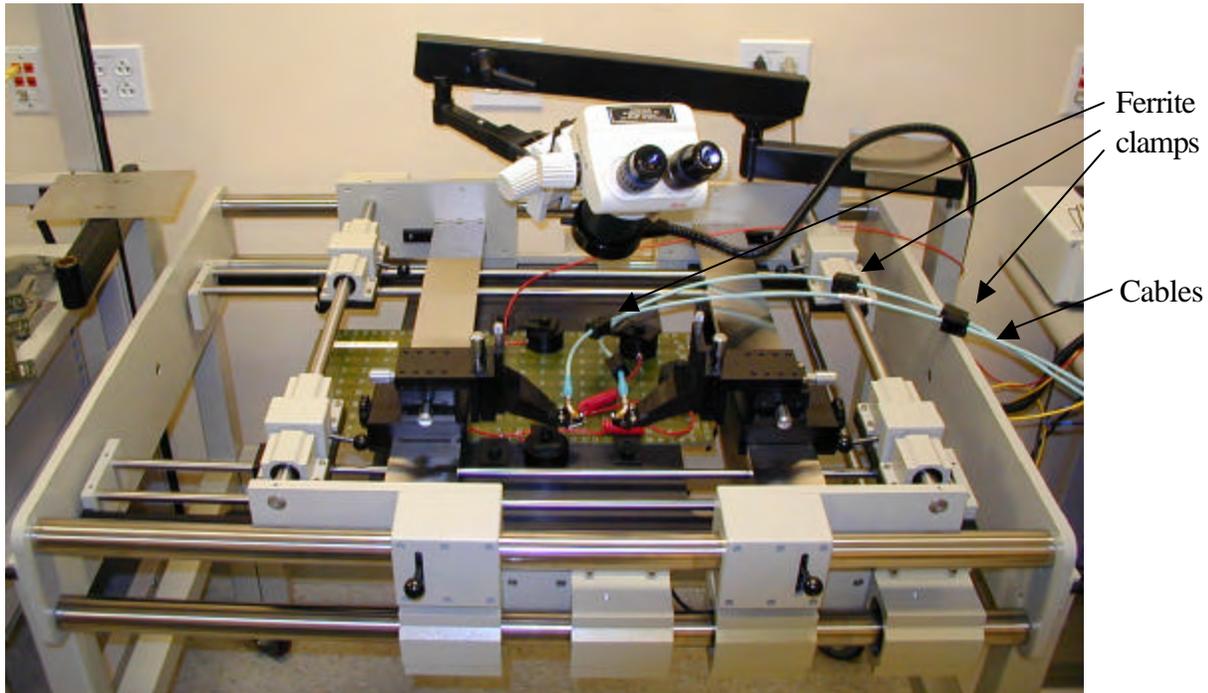


Figure 33. Large probe station with high-performance coaxial cables with a few ferrite clamps. Note the ferrite clamps at all locations where the cables otherwise would touch either each other or the metal frame of the probe station.

VI. Putting it all together

After having looked through the fundamentals of PDN measurements, we can apply the rules and tools to do actual measurements and characterization. In this section we show complex measurements and characterizations.

VI. 1. Measuring a full PDN

Figure 34 shows the setup for measuring the PDN of a powered board. The semirigid coaxial probes provide reliable measurements up to about 1-2GHz. The coaxial probes are mounted on small probe positioners.

For self-impedance measurements the probes are inserted into plated-through holes from the opposite sides. Measuring the self impedance from the opposite sides will ensure that the residual coupling through the probe-pin loops will not limit the measurement. Furthermore, by measuring from the opposite sides, we will measure the impedance as it appears across the PDN planes. If we measure the impedance at a via pair from the same side, what we get is the PDN impedance at the planes in series to the impedance of connecting vias. For transfer impedance measurements the probes can be inserted into the through holes from the same or from the opposite sides. Note the ferrite sleeves on the coaxial cables.

For low-frequency measurements with active VRMs, the coaxial cables are connected to a low-frequency VNA (in this case an HP4395) through an isolating transformer and an extra length of RG178 coaxial cable (to increase the DC resistance), or an isolation amplifier. For mid-frequency measurements, the same coaxial cables are connected to another VNA (in this case HP4396).

Calibrations can be performed with both VNAs, and measurement can then be performed over the entire frequency range only by switching cables at the VNA inputs: the probes can stay connected to the DUT.

The top two graphs of Figure 35 shows the measured self-impedance curves on one supply rail on a powered-up board. The two graphs together cover the 100Hz to 1.8GHz frequency range. The bottom two graphs illustrate the difference between PDN impedances measured from the same side versus opposite sides.

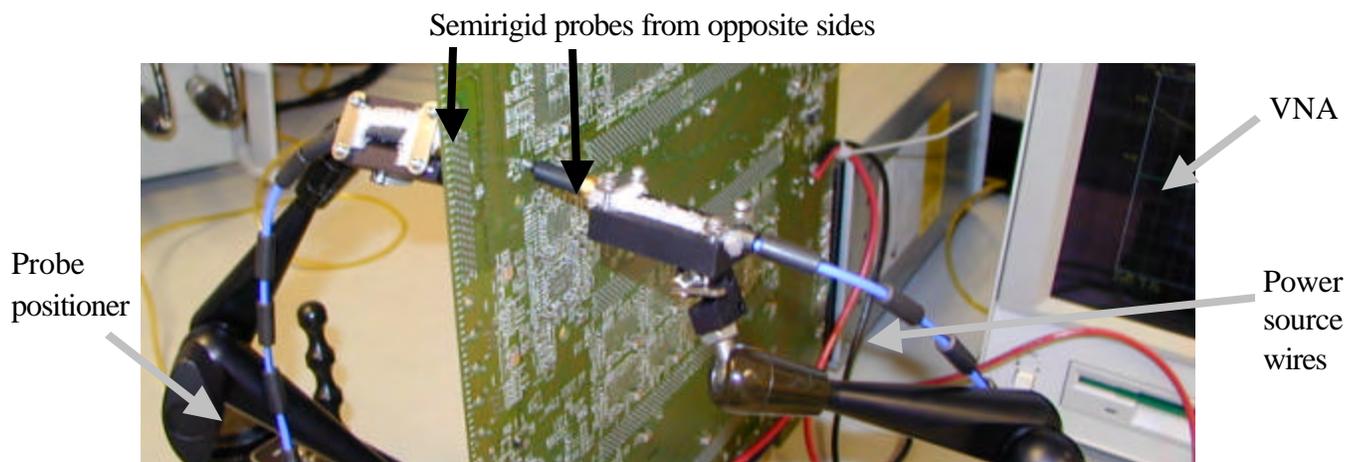


Figure 34. Probe connections for full PDN self-impedance measurements.

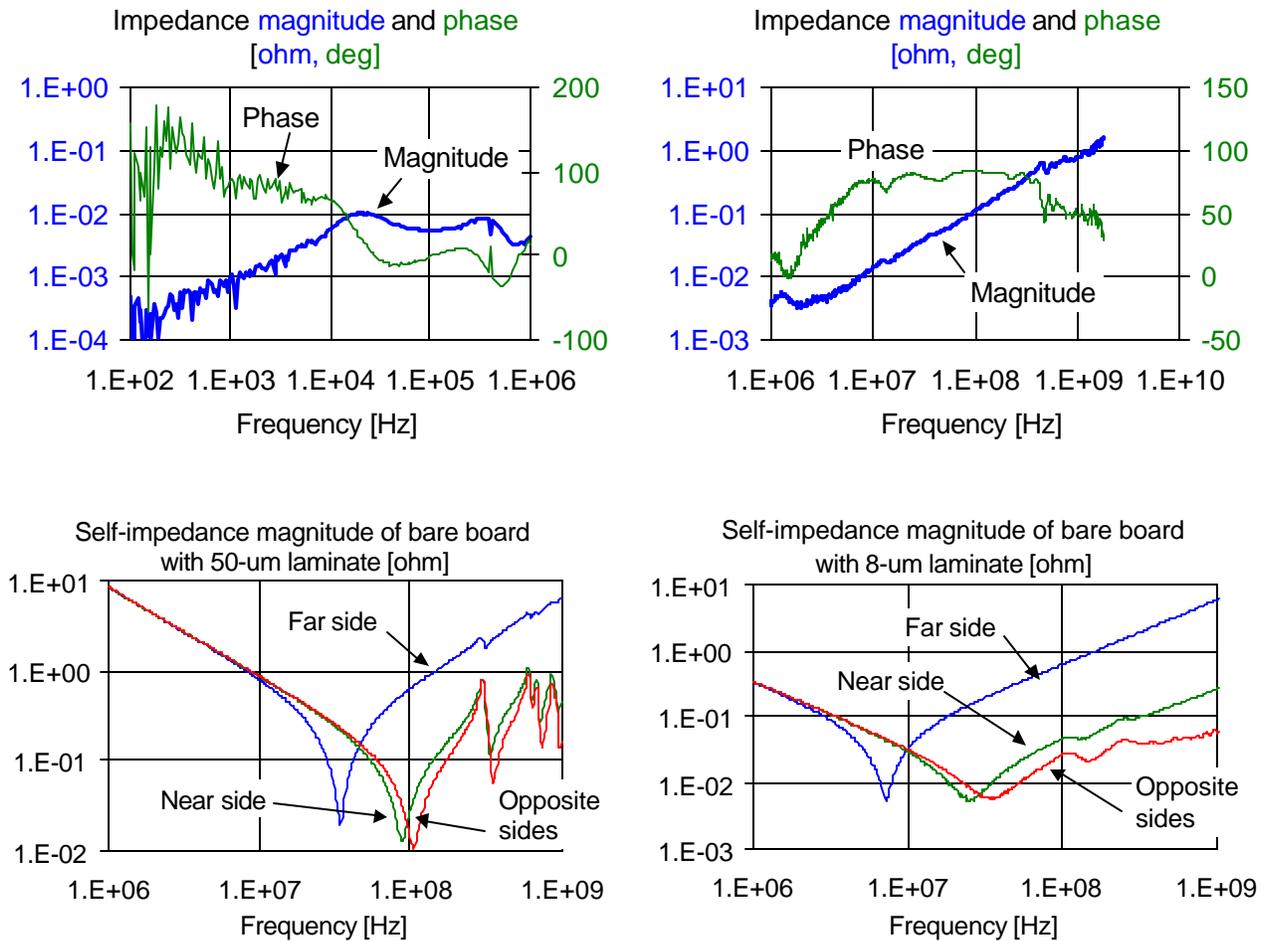


Figure 35. On the top: measured impedance profile on one of the supply rails of a powered board. Low-frequency impedance on the left, measured with an HP4395 VNA and isolation transformer. High-frequency self-impedance shown on the right, measured with an HP4396 VNA at one pair of test vias, with probes inserted from the opposite sides. Note the different horizontal and vertical scales on the two graphs. On the bottom: high-frequency illustration of the difference between PDN impedances measured from the same side versus opposite sides. Two 10"x5" boards were measured. Both boards had one plane pair in the PDN, 3 mils from the top surface and 90 mils from the bottom surface. On the left, the plane spacing was 50um; on the right the plane separation was 8um. When the impedance is measured through the 90-mil long vias from the bottom (labeled: far side), the impedance fluctuations are completely masked out by the via inductance, and we see no performance difference between the 50-um and 8-um laminates. When measured from the top (labeled: near side), the via interferes much less with the measured self impedance, but in case of the 8-um laminate, the inductance at high frequencies is still three times higher than that of measured from opposite sides.

VI. 2. Detailed characterization of PDN components

VI. 2. 1. Modeling the frequency dependent capacitance of low-ESR bulk capacitors

In case of bulk capacitors, the capacitance (and also the real part of impedance) may be strongly frequency dependent. This is due to the combined effect of large capacitance and low ESR. Simple spreadsheet approximations reveal that a three-element C-R-L equivalent circuit with frequency independent values has noticeable error of impedance prediction close to the series resonance frequency. This is illustrated on the right-hand graph of Figure 9, where the zoomed impedance plots of the measured data and frequency independent C-R-L impedance approximations are shown.

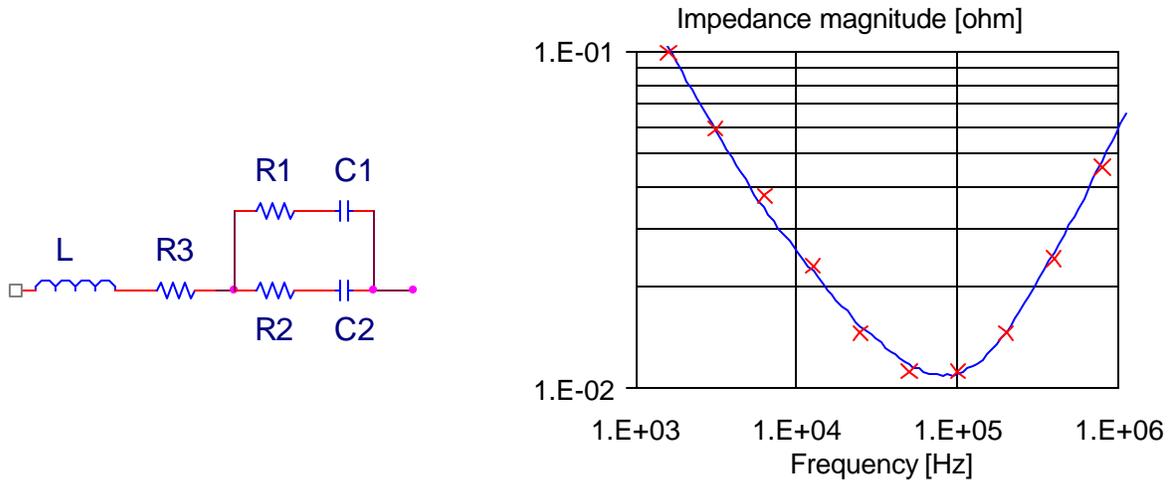


Figure 36. Five-element equivalent circuit on the left, and its correlation on the right, for the 1200uF polymer capacitor. Impedance magnitude on zoomed scale: continuous line shows measured data; crosses indicate simulation with the five-element equivalent circuit.

To achieve a better approximation, we could use frequency dependent C, L, and R values, however, this would limit the application of the equivalent circuit to AC simulations. Alternately, instead of one R and one C component, we can use an R-C ladder to describe the change of capacitance and resistance. For the given capacitor, two capacitors and three resistors in the equivalent circuit give reasonably accurate approximation of the measured data. Figure 36 shows the equivalent circuit on the left and the correlation on the right with $C1=5.5E-4$ F, $C2=5.5E-4$ F, $R1=5E-3$ ohm, $R2=9E-2$ ohm, $R3=6E-3$ ohm, and $L=9.4E-9$ H.

VI. 2. 2. Wide-band measurement of bypass capacitors

Figure 37 shows the measured impedance magnitudes of two multi-layer ceramic capacitors. Both capacitors had a nominally 1uF capacitance, and had the 0612 reverse-geometry form factors. One of the capacitors was the regular low-ESR type; the second sample was low-Q high-ESR capacitor, reported in [6]. The capacitors were soldered on a small test-fixture PCB, with approximately 1"x1" thin laminate in the board, close to the surface. To allow the measurement with one-sided probe stations, there were two pairs of vias, one on either side of the capacitor site. The distance between the via pairs was approximately 100 mils. The transfer impedances between the two via pairs were measured under different conditions. The fixture was first characterized as is, nothing mounted. The impedance profile of the bare test fixture follows the static capacitance up to beyond a GHz, followed by modal resonance peaks at 3.6GHz and 8GHz frequencies. Second, the test fixture was measured with the capacitor site shorted. Lastly, the capacitor sample was soldered on and measured.

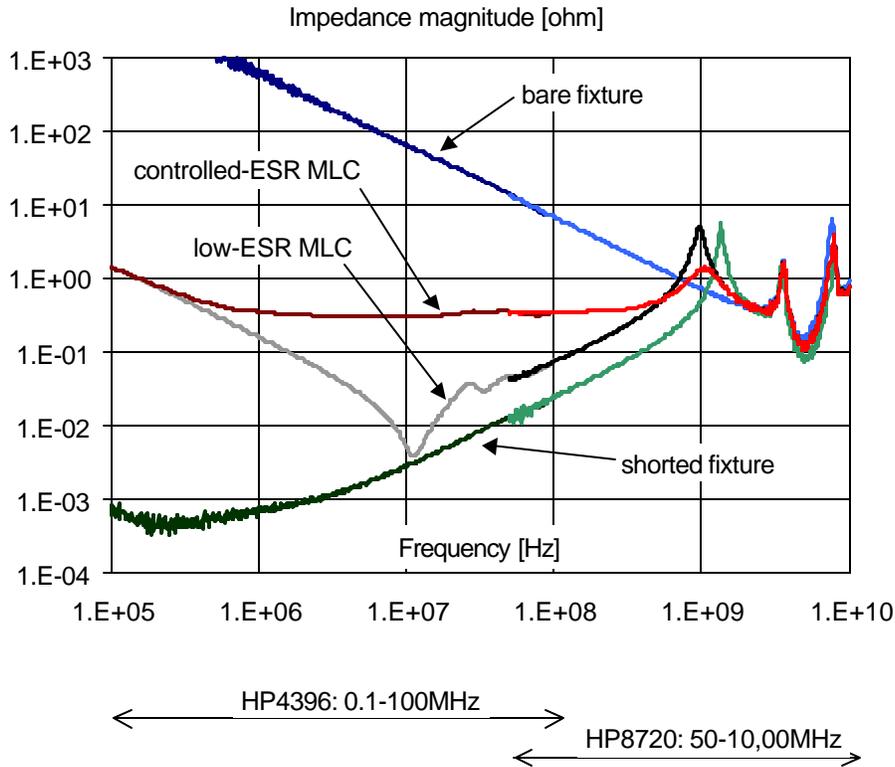


Figure 37. Wide-band characterization results of 0612-size 1uF multi-layer ceramic capacitors.

Note that data was taken with two independent instruments: in the 0.1-100MHz frequency range the VNA was HP4396, in the 50-10,000MHz range the data was taken with an HP8720 VNA. 450um-pitch PicoProbes were used with full two-port calibration on both VNAs. The data in the 50-100MHz range was taken with both VNAs, the smooth continuation of traces indicate good repeatability and accuracy of the independently recorded data sets.

VI. 2. 3. Determining the inductance of bypass capacitors: the concept of attached inductance

Inductance in real circuits is always realized by current loops. While the concept of partial self and mutual inductances allows us to break down the loop inductance into components, measuring these inductance components directly and separately would not be easy. In all of the previous examples the 'extracted equivalence inductance' meant some sort of loop inductance.

For PDN design and validation, however, the loop inductance may not be the most convenient parameter to characterize the high-frequency behavior of a bypass capacitor. This will be explained through Figure 38. The figure shows one bypass capacitor connected to a pair of planes through vias. The full construction will provide a Z_{loop} impedance with its corresponding inductance, which in general is frequency dependent. The loop inductance can be expressed in two different ways, and if done correctly, both partitioning should yield the same loop inductance value.

If we follow the concept of partial self and mutual inductances, we get:

$$L_{loop} = L_{capacitor-self} + L_{plane-self} - 2 * L_{capacitor-plane-mutual} + 2 * L_{via-self} - 2 * L_{via-mutual} \quad (17)$$

Here the only neglected contribution is the interaction of the orthogonal vias and plane and capacitor body. The two vias are assumed to have the same geometry; otherwise their self and mutual inductances could be accounted for separately. We have to recognize, however, that the current loop does not close on the upper plane only: it has to go through the dielectrics between the two planes and through the antipad taking the via connection to the lower plane. Therefore the $L_{\text{plane-self}}$ component in (17) does depend on the plane separation, antipad and via geometries. Moreover, the equivalent inductance of a plane pair is frequency and location dependent. This means that the loop inductance does depend on where we are on the plane. This also means that for simulation purposes, we would need to generate a different simulation model for the same capacitor depending on where it is connected.

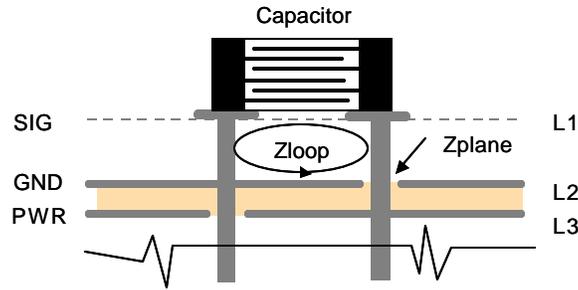


Figure 38. Sketch of a bypass capacitor mounted on a PCB and connected to an internal plane pair through vias.

We can generate a location and plane-independent simulation model for the bypass capacitor if we break down the loop inductance differently, as shown in (18).

$$L_{\text{loop}} = L_{\text{attached}} + L_{\text{plane}} \quad (18)$$

In (18), the L_{attached} component captures the partial self inductance of the capacitor, plus the self and mutual inductances of the vias, and the mutual inductance between the planes and the capacitor body. The attached inductance can be either larger or smaller than the partial self inductance of the capacitor body itself: with aggressive mounting, the mutual inductance between the capacitor and planes will create a low attached inductance.

Note that if the attached inductance is much larger than the plane inductance, the attached inductance and the loop inductance will closely match. As we use more aggressive mounting techniques, however, the attached inductance can become low compared to the plane inductance and the difference should not be neglected any more.

The attached inductance still depends on the mounting geometry, and we have to generate a separate simulation model for each of the capacitor/pad/via geometry we want to use. But the resulting L_{attached} will not depend on the location on the planes and will not carry the frequency dependency of the plane inductance.

We will illustrate the usefulness of the attached inductance concept through the modeling of a multi-terminal capacitor, mounted on a PCB, shown in Figure 39.

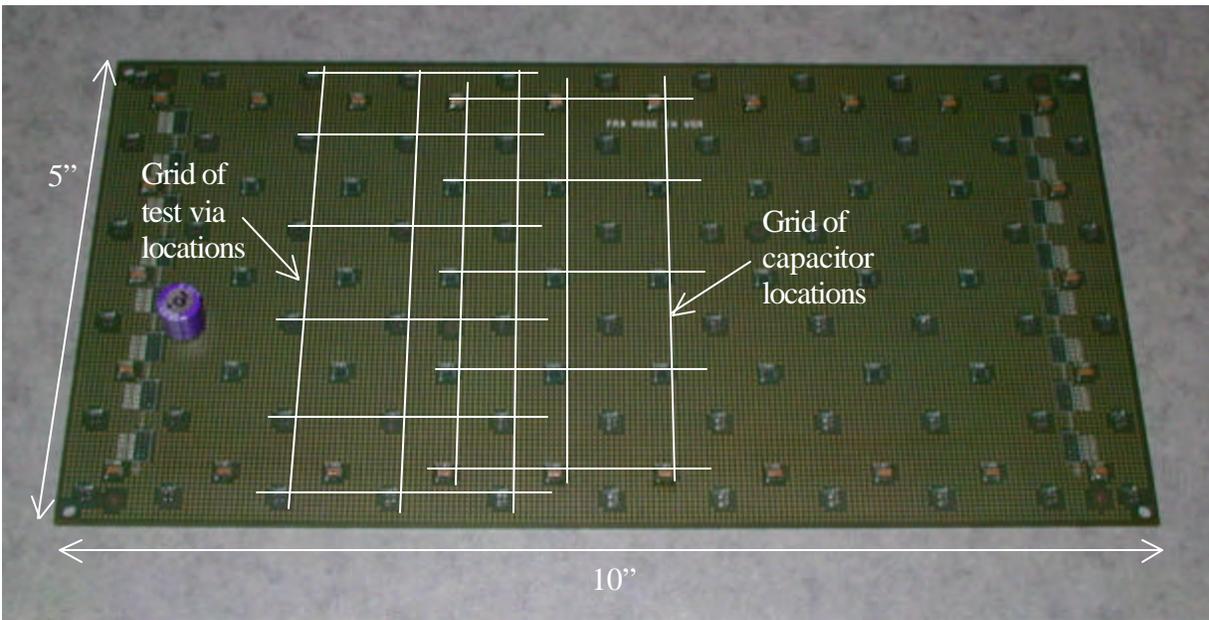


Figure 39. Photo of a 10"x5" test board, with the test via and capacitor grids identified. Besides of the single bulk capacitor, there are pads for 1206-size 8-terminal IDC parts.

The test board had a plane pair with 2-mil laminate 3 mils below the surface. Cross sections showing the top portion of the stackup with mounted capacitors of two different makes are shown in Figure 40.

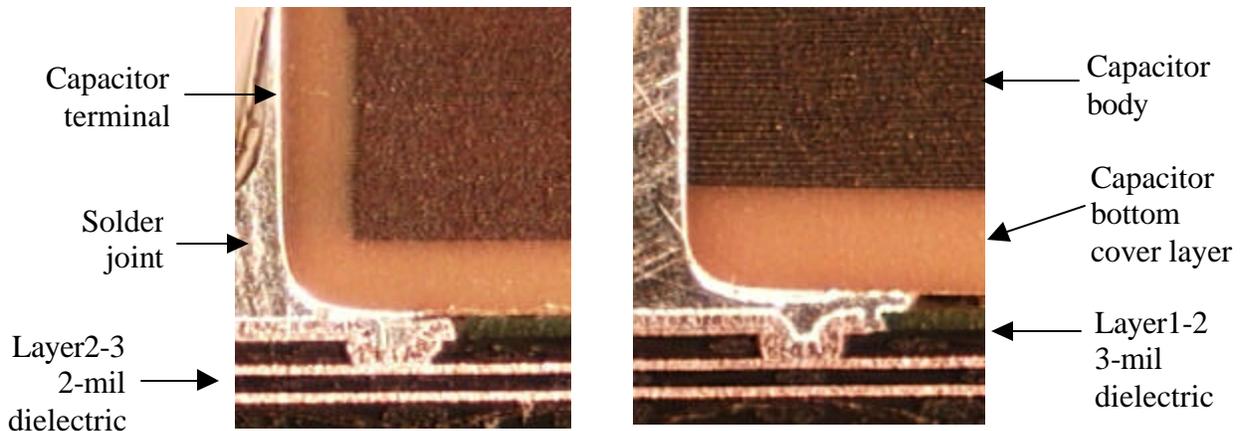


Figure 40. Cross sections of the top three metal layers of the test board, with two different capacitors. Only a corner of each capacitor and the top three layers of the boards are shown. Note the blind via connections under the capacitor terminals connecting the pads to Layer 2. Note also the difference in the bottom cover thickness of the two capacitors.

With the procedure outlined in III.4, the dielectric constant and the actual plane separation were determined. For the particular test board it was $s = 2.06$ mils/layer, $\epsilon_r = 4.22$. A detailed bed-spring simulation model was built only for the PCB planes, and it was correlated to measured data. The good agreement is illustrated in Figure 41. The simulation model used the conductivity of one ounce electrodeposited copper for the planes, and 0.02 for the loss tangent of the dielectric material.

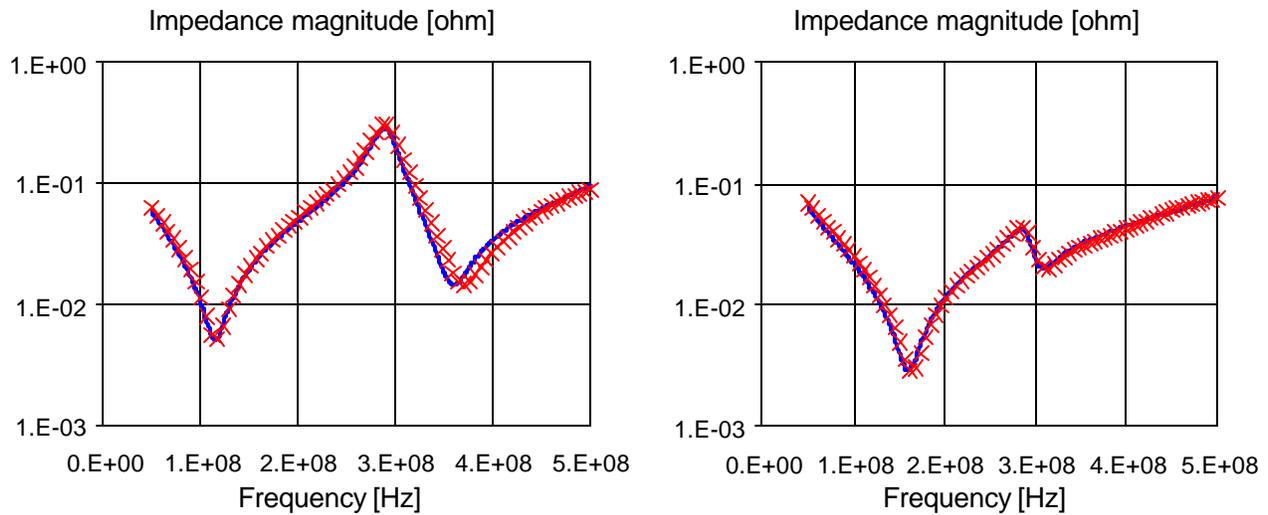


Figure 41. Simulated (heavy continuous blue line) and measured (red crosses) self-impedance magnitude at two different locations on the test board.

Using the accurate simulation model of the PCB planes, the model of the shorted capacitor pads and vias was obtained by fitting a simulated R-L model to the measured data. It was found that 0.5 milliohm DC resistance, $1.6E-7$ ohm/sqrt(f) skin resistance and 47 pH inductance gave good agreement for all combinations. The attached impedance of the shorted pad thus was:

$$Z_{attached} = 5E-4 + 1.6E-7 * \sqrt{f} + j2\pi f 47E-12 \text{ [ohm]} \tag{19}$$

Self and transfer impedances were measured and simulated with short placed over capacitor-pads at various locations. The graph on the left of Figure 42 shows three representative impedance plots with shorts at different locations and measured at various distances from the short. Note the slightly different slopes and significantly different resonance frequencies. The test board had two symmetry axes, so it was possible to prove that the variations in the shorted impedance curves were not due to local variations in via and pad geometries: the pad locations at the same symmetry points produced the same curves; the differences were due to the frequency and location dependent plane-pair impedance.

It was also found that the same expression of attached impedance given by (19) produced good correlation in simulations. The graph on the right of Figure 42 illustrates the agreement between simulated and measured impedances at one of the combinations.

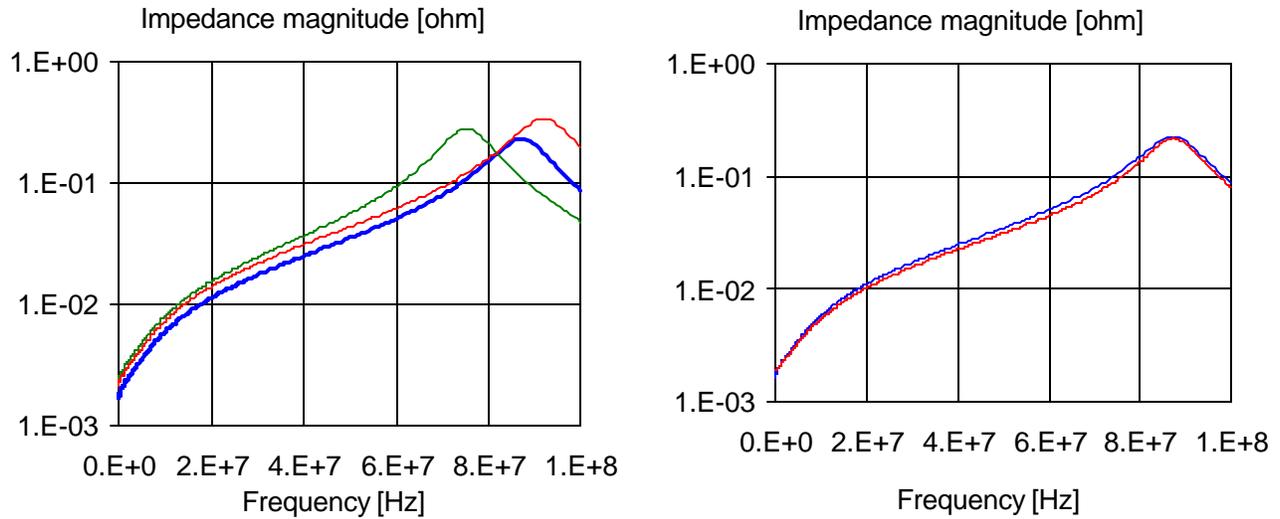


Figure 42. Left graph: measured self and transfer impedance plots at various locations of shorts and measurement points. Right graph: correlation between measurement and (19).

In the next step, the same piece of an eight-terminal 2.2uF IDC part from AVX was soldered onto the test board at various capacitor-pad locations, and self and transfer impedances were measured on the board with the single capacitor attached. Some of the representative impedance curves are shown on the left graph of Figure 43. The measured impedance profiles show the capacitive down slope at low frequencies, followed by a series resonance, and –after an inductive up slope– a parallel resonance. The equivalent circuit corresponds to that in Figure 6.e. Although the capacitor-pad and via geometries were the same at every location, and the same piece of capacitor was moved from location to location, there was a variation in the series and parallel resonance frequencies as well as in the impedance magnitudes at the resonances. The reason is the same as in case of the shorted plane example in Figure 42: the impedance and equivalent inductance of the plane pair is frequency and location dependent, whereas the capacitor pad/via geometry and the capacitor piece were the same for all of these measurements.

Similar to the shorted-pad case, we can establish a best-fit capacitor model that gives a location-independent characterization of the capacitor on the pads. Note that the model will describe the capacitor **and** the via/pad geometry together. Because the ESR of the capacitor was much higher than the via resistance, the skin-resistance portion of the via model does not show up separately. However, the ESR and attached inductance of the capacitor are both show noticeable frequency dependence.

Curve-fit models were created separately at the series resonance frequency (around 8MHz) and around the parallel resonance frequency (around 60MHz). The best-fit model for attached impedance was 6milliohms with 160pH at 8MHz (20), and 11milliohms with 120pH at 60MHz (21).

$$Z_{attached@8MHz} = 6E-3 + j2\pi f 160E-12 \text{ [ohm]} \tag{20}$$

$$Z_{attached@60MHz} = 11E-3 + j2\pi f 120E-12 \text{ [ohm]} \tag{21}$$

The graph on the right in Figure 43 illustrates the good correlation between the measured and simulated data at one of the locations.

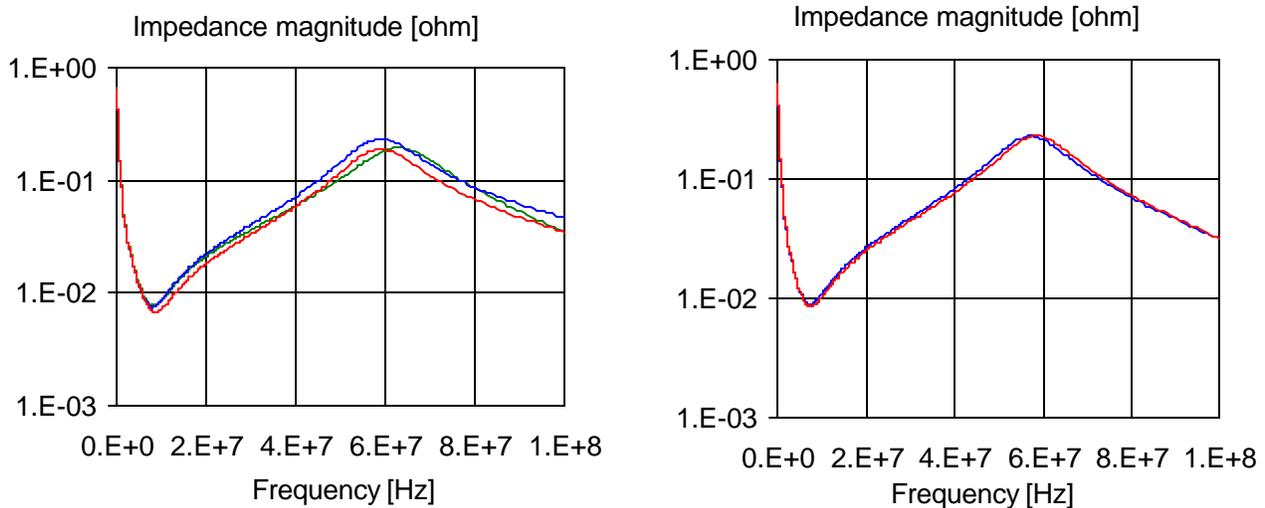


Figure 43. Left graph: measured self and transfer impedance plots at various locations of IDC part and measurement points. Right graph: correlation between measurement and (20) (21).

As it was shown above, the attached inductance can be determined separately for shorted capacitor pads and for the mounted capacitor. The advantage of the attached impedance/inductance concept is that it gives a model for a given capacitor and pad/via geometry, which is location and plane independent.

Note also that while the attached impedance of a given pad/via combination may show consistent results all over the board (because the geometry can be tightly controlled), the attached impedance/inductance of a capacitor will also depend on the height of the lowest capacitor plates above the PCB planes. This height has two major variables: the solder thickness between the pad and capacitor bottom, and the bottom cover thickness of the capacitor. The photos of Figure 40 show that the cover thickness of capacitor can be actually bigger than the dielectric layer height separating the first plane from the surface, and the variation of the cover thickness can be big. Therefore the good correlation shown in Figure 43 is just an illustration of the consistency of the attached impedance/inductance concept: if different pieces of capacitors are used at the various locations, the statistical variations of their geometry must be taken into account.

VI. 2. 4. Impact of capacitor height on the attached inductance of bypass capacitors

The size of the current loop, which determines the attached inductance of the part shown in Figure 44, is frequency dependent. The numerical illustrations of this frequency dependency were given in section VI.2.3. We can also notice that almost all of the extracted equivalent inductance graphs shown in earlier sections show some negative slope.

For bypass capacitors, this brings up the question: how the attached inductance depends on the geometry of the capacitor. Most specifically, for a given capacitor body size, how the inductance depends on the height of the cover thickness and the total height of the capacitor body.

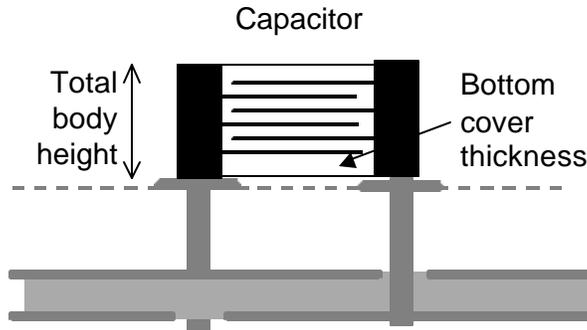


Figure 44. Sketch defining the total capacitor height and the bottom cover thickness of a capacitor.

To get the answer, a series of measurements was done on a small-size pair of planes with capacitor pads, similar to that shown on the top of Figure 5. The small plane pair was characterized with open (nothing soldered on it) and shorted (the capacitor pads shorted with a sheet metal). One capacitor was soldered on the pads, and on the nearby via pair, the self-impedance of the DUT was measured. From the self impedance, the equivalent inductance was extracted. One by one additional capacitors were soldered piggybacked, on top of the previous capacitors. This way we made sure that the geometry between the fixture's planes and closest capacitor plates were left unchanged, only the total capacitor height was increased.

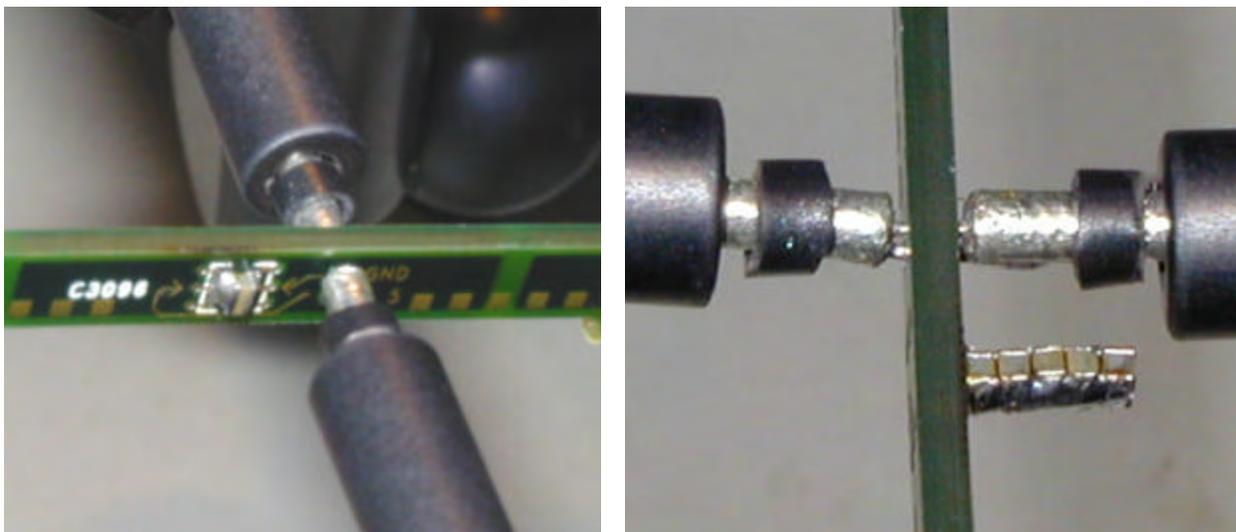


Figure 45. Small fixture with one and five stacked capacitors on the pads.

The fixture and the mounted capacitors are shown in Figure 45, with one and five capacitors on the pads. The capacitors were AVX reverse-geometry 0508 parts with 4.7uF nominal capacitance. The impedance profile was taken with an HP4396VNA after full two-port calibration up to the semirigid probes. Figure 46 shows the results. On the left, the traces show the impedance magnitude of the full DUT (small plane plus capacitors) with a reference trace, where the capacitor pads were shorted. The shorted trace runs somewhat below the other traces, indicating lower inductance.

With one, two and three stacked capacitors, the series resonance frequency was within the 1MHz-1800MHz measurement range. The table below lists the measured series resonance frequencies and the equivalent loop inductance based on the series resonance frequency and total capacitance.

	Capacitance [F]	Series resonance frequency [Hz]	Equivalent total inductance [H]
1x4.7uF	4.7E-6	2.77E6	7.01E-10
2x4.7uF	9.4E-6	1.69E6	9.38E-10
3x4.7uF	1.41E-5	1.17E6	1.31E-9

Note that as expected based on [4], as the total height of the capacitor stack increases, the inductance at the series resonance frequency also increases.

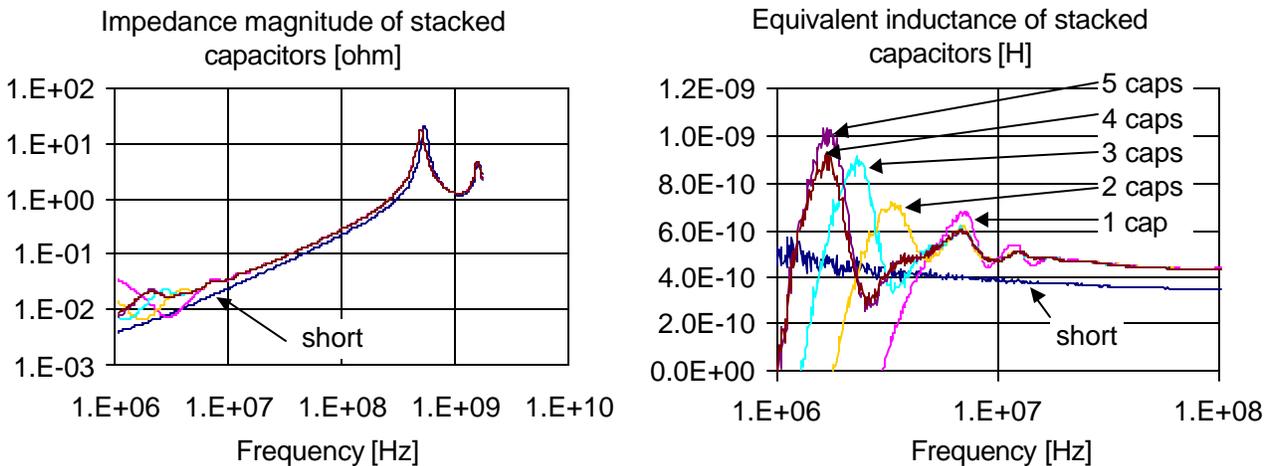


Figure 46. Measured impedance magnitude and extracted inductance of the small fixture with shorted pads and one, two, three, four and five stacked capacitors on the same pads.

The right-hand graph in Figure 46 shows the extracted inductance versus frequency. The inductance with shorted pads varies between 500pH (at 1MHz) and 350pH (at 100MHz). The inductance with different number of stacked capacitors varies strongly close to the series resonance frequency, but it quickly stabilizes at values running about 100pH above the trace of the shorted pad inductance. This 100pH difference comes from the extra current-loop size created by the cover thickness of the lowest capacitor in the stack. Above 20MHz there is no noticeable difference among the inductance values with the various number of stacked capacitors.

As the measure of effectiveness of bypass capacitors is their high-frequency inductance, we can conclude that it primarily depends on the cover-layer thickness and not on their total height.

VI.2.5 Repeatability of data

In making measurements on low-impedance structures, the repeatability of the connections has to be looked at. There is a valid concern that the finite contact resistance between the probe tips and DUT connection points may alter the data. To illustrate the robustness of the two-port measurement setup, a low-ESR bulk capacitor was measured repeatedly in the same small PCB fixture without soldering. The sample was inserted into the through holes of the fixture, and was held in place by slightly twisting the capacitor body, so that the leads make connections to the through-hole walls. Figure 47 shows the setup and the result. Note that the extracted impedance minimum (ESR) shows a variability of about one milliohm, but the inductance reading is very consistent. This suggest that to measure low-ESR bulk capacitors and/or active VRMs, soldered connections may be needed. But at higher frequencies, where inductance dominates the path, the repeatability is sufficient without soldered connections.

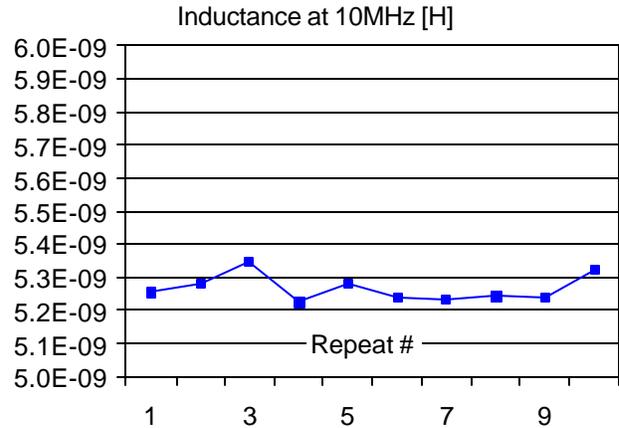
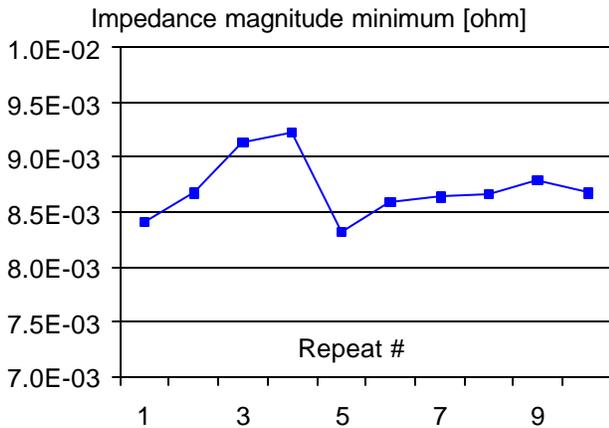
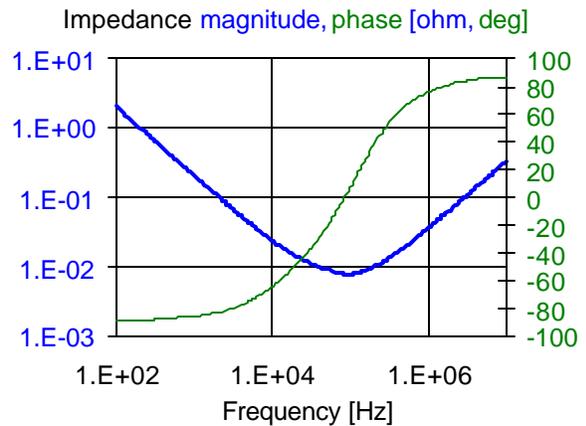
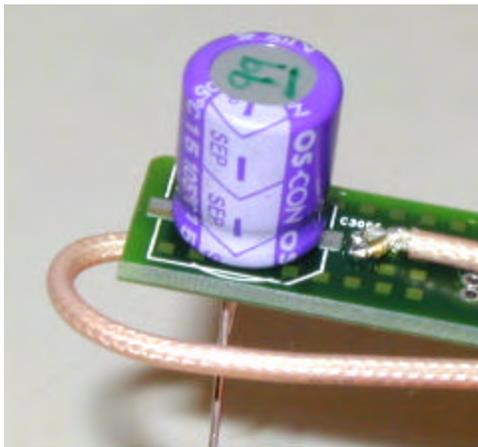


Figure 47. Illustration of repeatability of data at low frequencies in fixtures without soldering. On top left: photo of the fixture with one of the samples inserted. On top right: impedance magnitude and phase versus frequency. On bottom: minimum of impedance magnitude on the left, and extracted inductance at 10MHz on the right. The same sample piece was measured ten times in the same fixture, without soldering.

Acknowledgements

The author wishes to thank Sreemala Pannala, Leesa Noujeim, Jason R. Miller for their contributions and to Eric Blomberg, Mark Kanda and Kevin Hinckley for the useful discussions.

References

- [1] Greg Taylor, Craig Deutsche, Tawfik Arabi, Brian Owens, "An Approach to Measuring Power Supply Impedance of microprocessors," Proceedings of the 10th Topical Meeting on Electrical Performance of Electronic Packaging, October 29-31, 2001, Cambridge, MA, pp. 211-214.
- [2] Isaac Kantorovich, Chris Houghton, Steve Root, Jim St. Laurent, "Measurement of Milliohms of Impedance at Hundred MHz on Chip Power Supply Loop," Proceedings of the 11th Topical Meeting on Electrical Performance of Electronic Packaging, October 27-29, 2002, Monterey, CA
- [3] Istvan Novak, "Measuring Milliohms and PicoHenrys in Power Distribution Networks," DesignCon2000, Santa Clara, CA, February 1-4, 2000.
- [4] Larry D. Smith, David Hockanson, "Distributed SPICE Circuit Model for Ceramic Capacitors," Proceedings of the 51st Electronic Components and Technology Conference, May 29 - June 1, 2001, Orlando, FL
- [5] Richard Redl, Brian P. Erisman, Zoltan Zansky, "Optimizing the Load Transient Response of the Buck Converter," Proceedings of APEC98
- [6] Istvan Novak, Leesa M. Noujeim, Valerie St. Cyr, Nick Biunno, Atul Patel, George Korony, Andy Ritter, "Distributed Matched Bypassing for Board-Level Power Distribution Networks," IEEE Transactions on Advanced Packaging, Vol. 25, No.2, May 2002, pp.230-243.