

Agilent Ultra-Low Impedance Measurements Using 2-Port Measurements

Application Note




Table of Contents

Ultra-Low Impedance Measurements Using a Vector Network Analyzer	3
Limitations of 1-Port VNA Impedance Techniques	5
4-Point Kelvin Technique for Ultra Low DC Resistance Measurements.....	10
2-Port Measurements Reduces Fixturing Parasitics.....	11
Measurement Examples Using the Two-Port Measurement Technique	16
The Power Distribution Network (PDN).....	20
Decoupling Capacitors	34
Planes and Capacitors	40
How to Set Up a Measurement System	42
Conclusion.....	51

Ultra-Low Impedance Measurements Using a Vector Network Analyzer

In this application note, a 2-port network analyzer is used to measure impedances well below 1 Ohm. This regime is difficult to measure in practice with a conventional 1-port VNA due to real world limitations of signal to noise ratio and fixturing reproducibility.

With this new technique of using 2-ports and a conventional network analyzer, impedances as low as 1 milliOhm and inductances in the pH range can be routinely measured. This type of measurement is critically important for all the components that make up the power distribution network system.

- Vias
- VRM
- Capacitors
- Planes

Why is low impedance important?

Most interconnects used to transport signals have impedances in the 50 to 100 Ohm range. This is in the perfect range for measurement by conventional network analyzers with port impedances of 50 Ohms. But, for structures with ultra-low impedances, the mismatch with the 50 Ohm source impedance means that nearly all of the signal will reflect, and distinguishing 0.1 Ohm from 0.01 Ohms becomes extremely difficult.

It is predominately in the power distribution network (PDN), the interconnects from the voltage regulating module (VRM) that generates the precisely regulated voltage to the pads on the chip for the Vcc or Vdd rails, where ultra low impedance values are required. It is not uncommon in microprocessor based systems to have a target impedance for the entire PDN of less than 10 milliOhms from DC to a few GHz.

Each of the components that make up the PDN, the package leads, the ceramic capacitors, the on-chip capacitance, the power and ground planes of the circuit board, and even the VRM itself, must have impedances in the milliOhm range. It is not practical to measure them with the conventional return loss of a 1-port VNA, these are the sorts of applications for which the 2-port technique is critically important.

- Vias
- Package attach elements: wire bonds, solder balls
- Ceramic capacitors
- On chip capacitance
- Planes
- Power distribution networks
- Voltage regulator modules (VRM)

Limitations of 1-Port VNA Impedance Techniques

The simplest equivalent circuit model for 1-port of a VNA is a sine-wave generator with a source impedance of 50 Ohms. This signal transmits through an internal 50 Ohm coaxial transmission line to the front connector of the VNA and is launched into the device under test (DUT).

The amount of the incident sine wave which comes back to the detector and is measured as the reflected signal, depends on the miss match in impedance between the DUT and the 50 Ohm source. The reflection coefficient, S11, is the DUT impedance minus the 50 Ohms, divided by their sum. You can use this relationship to estimate the value of S11 for low impedance devices.




Figure 1. Reflection coefficient

When the impedance of the device is 1 Ohm, S11 is -0.96 or -0.35 dB. This is a reasonable value to be able to measure. When the impedance is 0.1 Ohm, S11 is -0.996 or -0.035 dB. This may seem like a very small amount of reflected signal, but in fact, 99.6% of the incident signal reflects. This is so close to 100% that it is difficult to measure the difference between 100% of the signal and the 99.6% of the signal. This creates a large relative uncertainty.

When impedances are less than a small fraction of an Ohm, the magnitude of the reflected signal is so close to 1 that it is difficult to distinguish it from 0 Ohms. This is a fundamental problem, as seen in Figure 2.

When measuring a low impedance:

$$S_{11} = \frac{1 - 50}{1 + 50} = \frac{-49}{51} = -0.96 = -0.35 \text{ dB}$$

$$S_{11} = \frac{0.1 - 50}{0.1 + 50} = \frac{-49.9}{50.1} = -0.996 = -0.035 \text{ dB}$$

Figure 2. Low impedance reflection coefficient

1-port impedance measurements are limited to $Z > \sim 0.1$ Ohms

If we look at the behavior of S_{11} , we find that when the impedance is less than about 0.1 Ohms, the value of S_{11} is close to 99% of 1. The difference between 1 and S_{11} is a useful metric for determining the difficulty of a measurement. The smaller the difference, the closer S_{11} is to 1, and the harder it is to measure.

Using the threshold of $S_{11} > 0.99$ as the practical limit for accurate measurements, Figure 3 shows that impedances less than about 0.1 Ohms are difficult to measure using a 1-port technique. This is the first, fundamental problem with low impedance measurements.

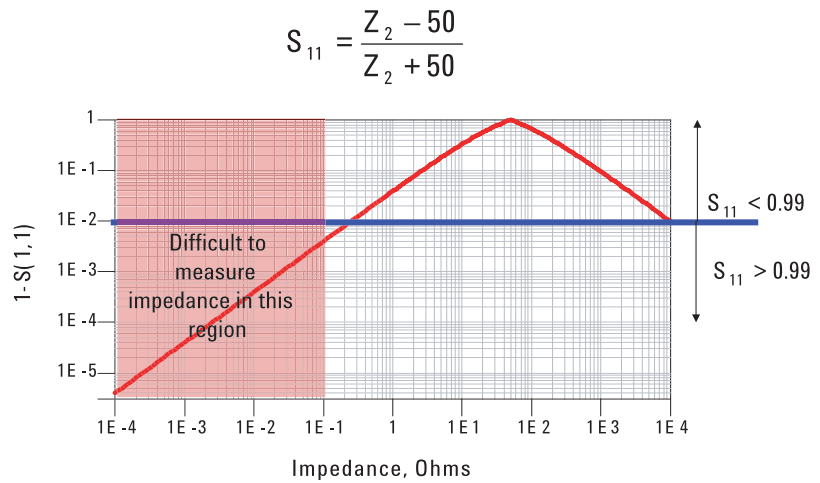


Figure 3. Low impedance measurements can be challenging for 1-port impedance measurements

Contact impedance and reproducibility

The second problem with low impedance measurements, is the reproducibility of the impedance of the fixturing to the DUT after calibration. Most calibration processes establish a reference plane at the end of the connector. Even in the case of a coaxial connection, the amount of torque used to make the connection to the DUT influences the residual, uncalibrated phase or series inductance to the DUT. Unless the electrical length of the test cable connected to the DUT is precisely known, there will be uncompensated inductance in the path that will be allocated to the DUT. It can either add or subtract series inductance to the DUT.




Figure 4. Probe model

In the case of microprobes, the reference plane can be established right at the tip of the probes by touching them to an open, short and load reference substrate. However, even after calibration, the flexing or over-travel of the microprobe will influence the uncompensated, residual inductance of the probe. With roughly 25 pH per mil of travel, and 2 to 8 mils of uncontrollable travel, this can be an uncompensated, variable inductance from the probe, in series with the DUT of as much as 200 pH of inductance.

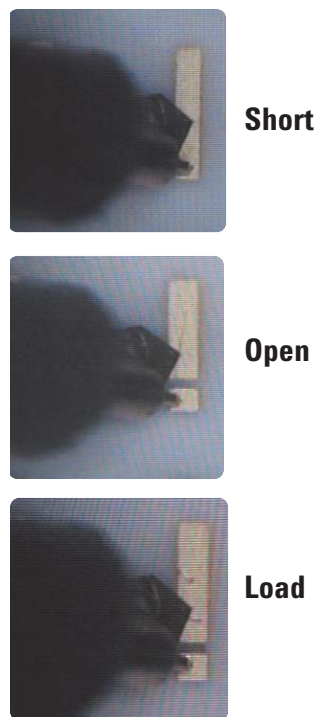


Figure 5. Short, open, load probe configurations

Further, depending on the metallization of the probe tip and the surface of the pads there can be a residual, uncompensated resistive impedance term on the order of 10 to 100 milliOhms due to contact resistance. Both of these factors contribute to a probe series impedance which cannot be calibrated out, on the order of about 0.1 Ohms and as much as 100 pH. This is the limit to the reliable measurements of a low impedance using a 1-port technique.

Residual R, L

- $R \sim 0.01\text{-}0.1$ Ohms
- $L \sim 0.05 - 0.1$ nH (2-8 mils length changes)
- Variation from calibration to measurement

Figure 6. Residual resistance and inductance

When comparing the log of the magnitude of a DUT's impedance to the log of the frequency, an ideal capacitor has an impedance that falls like $1/f$. An ideal inductor has an impedance that increases proportional to f . If the residual impedance limit of a probe is defined as 0.05 Ohms resistance in series with a 100 pH inductor, the impedance profile of this impedances defines a boundary between routine and difficult 1-port impedance measurements.

Figure 7 shows that trying to measure inductances below 100 pH or capacitances greater than 0.1 microF at frequencies above 20 MHz, is very difficult with a 1-port VNA. For these very low impedances or very low inductance structures, a new technique is needed that gets around these limitations.

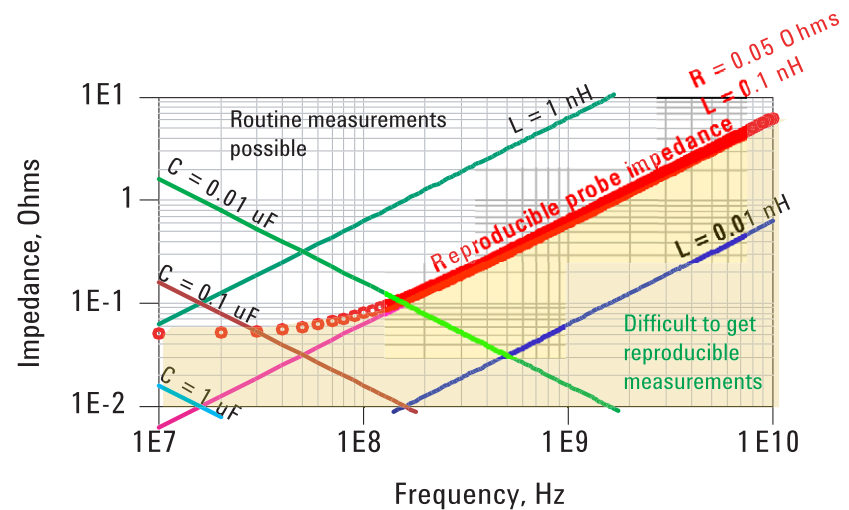


Figure 7. Range where contact impedance plays a role

4-Point Kelvin Technique for Ultra Low DC Resistance Measurements

Typical measurements of very low DC resistances share a common problem of contact resistance associated with the actual leads that make contact to the DUT. Due to contamination and oxide build up at the interfaces, or just due to the constriction resistance of the very tiny contact area, there is often a contact resistance on the order of 10 mOhms.

If two electrodes contact a very low resistance bar of copper, for example, the measured series resistance is not just the resistance of the rod, it is also the series resistance of the contacts on the two ends. These two resistances often dominate the total resistance and hide the true resistance of the rod.




Figure 8. Kelvin technique for a copper rod

Kelvin technique for DC resistance

In the late 1880's, William Thompson, also known as Lord Kelvin, developed a technique using 4 leads to get around the artifact of contact resistance when measuring ultra low resistances. The basis of his technique is to use two leads to force the current through the DUT and two separate leads to measure the voltage generated across the DUT from the IR drop of the current through the intrinsic, internal resistance of the DUT.




Figure 9. Separate leads for current force and voltage sense eliminate contact resistance effects

Figure 9 illustrates that when a current, I , is forced through the DUT, there is still contact resistance in series with the leads from the current generator. Due to the intrinsic series resistance of the DUT, the current, I , generates a voltage drop across the DUT, independent of the contact resistance, equal to $I \times R$ of the rod. This voltage is measured with separate contacts. Of course, the volt meter is high impedance, so the contact resistance of the voltage leads does not play a role. It can be as high as 1 MegOhm and still not influence the voltage measurement.

The resistance of just the copper rod is given by the measured voltage, V , divided by the current, I , forced. Neither of these two terms are at all affected by the contact resistance. The 4-point or Kelvin technique, which uses one pair of leads to drive the current and the second pair of leads to independently sense the voltage, is a powerful technique for measuring intrinsic resistances well below 1 milliOhm.

2-Port Measurements Reduces Fixturing Parasitics

The 2-Port VNA technique is the rf equivalent of the Kelvin DC technique. In this type of measurement, both ports of the Agilent N5230A vector network analyzer (VNA) are connected to the same pads of the DUT. Port 1 is used to drive the current through the DUT and port 2 is used to measure the voltage generated across the DUT.

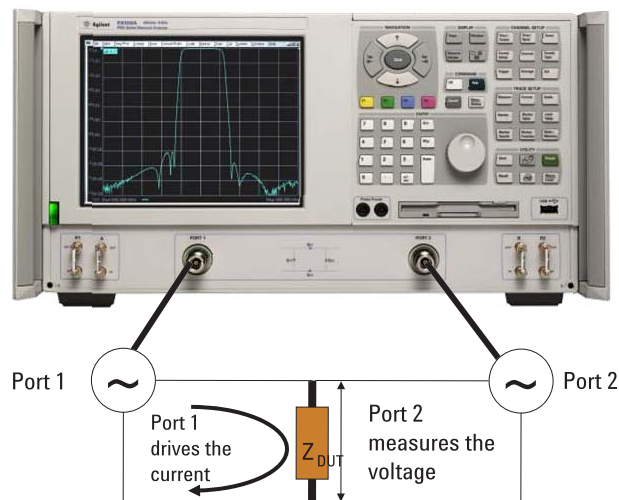


Figure 10. Shows how the 2 ports of the Agilent N5230A VNA are used to measure the voltage of the DUT

In this 2-port method, S_{11} and S_{22} see the impedance of the DUT, with a 50 Ohm impedance shunting across it, from the other port. The information contained in the return loss measurements is clouded by the low impedance of the DUT and the additional 50 Ohm shunt from the other port.

However, S_{21} has much more valuable information about the impedance of the DUT. In principle, the current from port 1 is limited by the 50 Ohm resistance of the VNA source. This current, through the low impedance of the DUT, creates a very small voltage across the DUT. This voltage is launched by the DUT into port 2 and is received. With a dynamic range of at least 80 dB, and in some cases as large as 110 dB, very small voltages can easily be measured by port 2. The ratio of the voltage at port 2 to the source voltage from port 1, S_{21} , has information about the very small impedance of the DUT.

First order analysis

Two simple analyses can relate the value of S21 to the impedance of the DUT. In the first order analysis, if we assume the impedance of the DUT is very small compared to 50 Ohms, then when the signal from port 1 hits the DUT, all of it reflects. The reflection coefficient of the voltage is -1 , the reflection coefficient of the current is just 1. This means that the total current through the device is the incident current plus the reflected current.

This total current produces a voltage across the device which is transmitted into port 2. From the voltage across the DUT that is picked up by port 2, you can estimate the impedance of the DUT. To the first order, the impedance of the DUT is 25 Ohms times S21. This is a very simple relationship and can literally be read off the front screen of the VNA.

The assumption for this relationship is that the impedance of the DUT is very close to 0. In practice, if the impedance of the DUT is greater than about 5 Ohms, this assumption is not very good and a more exact relationship should be used.

If $Z_{DUT} \ll 50$ Ohms

- Current out of port 1 = $V_{port1} / 50$ Ohms = 1 v / 50 Ohms ~ 20 mA
- Voltage across DUT ~ 0
- $\rho \sim -1$
- Current through DUT = 2×20 mA
- to first order, $V_{DUT} = V_{port2} = Z_{DUT} \times 2 \times V_{port1} / 50$ Ohms = $Z_{DUT} \times V_{port1} / 25$ Ohms

$$Z_{DUT} = 25 \Omega \times \frac{V_{port2}}{V_{port1}} = 25 \Omega \times S_{21}$$

Figure 11. Circuit characteristics when the output impedance is much less than 50 ohms

Second order analysis

The equivalent circuit of the VNA ports connected to the DUT is shown in Figure 12. Port 1 is a voltage source and 50 Ohm resistor in series. The DUT and the 50 Ohm of port 2 are part of the circuit. The voltage picked up by port 2, the voltage across the DUT, is very simple to evaluate based on the voltage divider circuit.

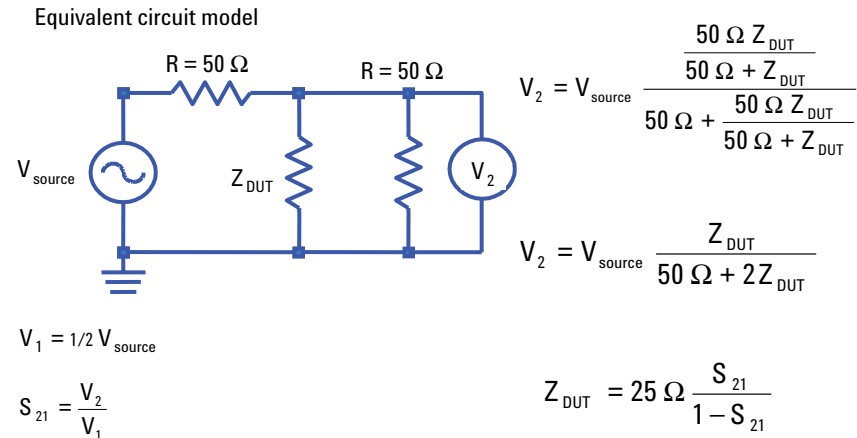


Figure 12. Equivalent circuit model

One important relationship is that the actual incident voltage to the DUT, transmitted from port 1, is not the source voltage; it is half the source voltage. Since S_{21} is the ratio of the measured voltage to the incident voltage, a factor of 2 is used. This results in the final relationship defining the impedance of the DUT as 25 Ohms times S_{21} divided by 1 minus S_{21} . This is an exact relationship and is valid for any value of the impedance of the DUT.

Recommendations for accurate 2-port measurements

This technique of using two ports to measure a very low impedance is very simple to implement. In practice, it is important to keep the phase length of the connections from the ports to the DUT as short as possible. As a rough rule of thumb, it should be within $1/20$ of a wavelength of the highest frequency measured. In FR4 type material with a dielectric constant of 4, this translates to the maximum frequency, in GHz, which will have accurate results should be less than 0.3 divided by the length, in inches. If the length of the connection is 1 inch, this limits the reasonable frequency regime for accurate results to about 300 MHz.

This is why, for high frequency measurements, it is important to use microprobes, rather than pigtailed. The microprobes, calibrated right to the tips, can enable accurate frequencies well above the GHz range.

When the fixture parasitics are low, the measurement of S_{21} can be directly interpreted as the impedance of the DUT, using the second order relationship. Even when the fixture impedance is not negligible, the effects of the fixture can be eliminated by using a good model to describe its behavior. In this way, the impact from the fixturing can be "de-embedded" from the measurements.

- Keep fixture parasitics as small as possible
 - Max accurate frequency, when fixture length $< 1/20 \lambda$
 - $F_{max} \sim 0.3/Len$, F_{max} in GHz, Len in inches
 - Microprobes and probe station recommended for > 1 GHz
- Interpret results directly with 2nd order method for $Z > 1$ Ohms
- It's always possible to use circuit topology de-embedding of the fixturing and S parameter measurements directly

Measurement Examples Using the 2-Port Measurement Technique

The sections that follow show how the two-port measurement technique can be used to measure the impedance of various devices, substrates and materials. The Agilent N5230A vector network analyzer is used with the Ultimatrix P4800 power delivery network analyzer (PDNA) to perform these measurements

Impedance of a via

In this first example, a via has been drilled in a microstrip shorting the signal line to the return path. The via is at the middle of the microstrip, with SMA connections on either side. The inductance of the via can be measured with both the 1-port and 2-port techniques.

As a general rule, a measurement or simulation should never be performed without first anticipating what to expect. The circuit board is 064 mils thick, which is the length of the via and the via diameter is 25 mils. So you can estimate the partial self inductance of the via using the simple rule of thumb that for a rod 1 mil in diameter, the inductance is about 25 nH/inch, and varies inversely with the log of the diameter.

For 25 mil diameter, the natural log of 25 is about 3, so the inductance should be about 8 nH/inch. For a length of 0.064 inches, the inductance of the via should be about 0.5 nH. At 100 MHz, the impedance of 0.5 nH is about 0.3 Ohms, which is difficult to measure with a 1-port VNA. There is a further complication, that a 1-port measurement will also include the impedance of the 0.35 inches of transmission line leading up to the via. In a 1-port measurement, how much of the measured inductance is the transmission line and how much is the via?

With about 9 nH/inch in a 50 Ohm line, 0.35 inches of transmission line has about 3 nH of inductance. This inductance will easily swamp the inductance of the via. However, the 2-port technique will separate the via inductance from the transmission line inductance.




Figure 13. Two-port short via

Why does 1-port measurement show $L \sim 10 \times$ the 2-port value?

With 1-port only connected to the DUT, S_{11} is measured and can interpret this as an impedance. Plotted on a log-log scale, the impedance has the general form of an ideal inductor. When measuring the series inductance of the SMA connector, the short length of transmission line and the via make it difficult to separate the via inductance from the transmission line inductance using a 1-port measurement.




Figure 14. In a 1-port measurement, it's challenging to separate via inductance from transmission line inductance

In the 2-port measurement, port 2 is measuring the voltage generated across the via, and is independent of the transmission line leading up to it. The S_{21} measurement can be directly interpreted as an impedance and plotted. Since the shape of the impedance profile matches that of an ideal inductor, we can interpret the impedance as an inductance, and plot the extracted inductance.

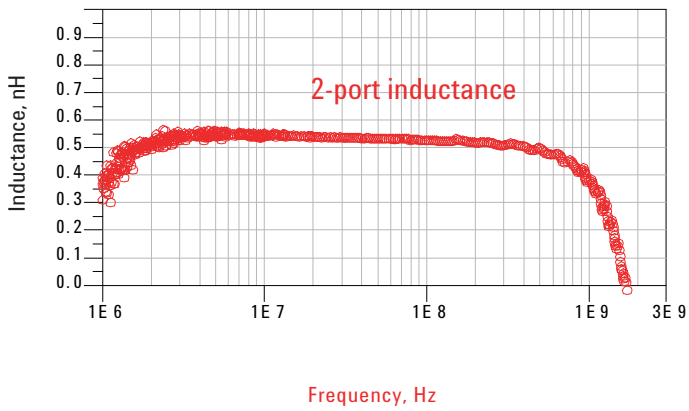


Figure 15. In a 2-port measurement, the voltage generated across the via is independent of the transmission line leading up to it

The inductance for each frequency value shows some frequency variation, but in the range from about 2 MHz to about 600 MHz, the inductance is very constant and is about 0.54 nH, which can be read right off the plot.

Given the length of the fixture to the via as 0.35 inches, you would expect accurate values of the DUT impedance for frequencies below about 1 GHz. This is exactly what is seen in the data. The drop of in inductance above about 600 MHz is not a real behavior of the inductance of the via, but is an artifact due to the length of the fixture leads. This can be verified by building a model of the DUT and the fixture and de-embedding just the via from the fixture.

Resolve root cause from a model of via short

The first step in de-embedding the fixture is to build a simple equivalent circuit model for the via and the fixture. In this case, the fixture is modeled as an ideal, lossless transmission line, with a characteristic impedance and a time delay. To simplify the model, assume that the same transmission line is on both sides of the via and that the via itself is an ideal inductor with a series resistor.

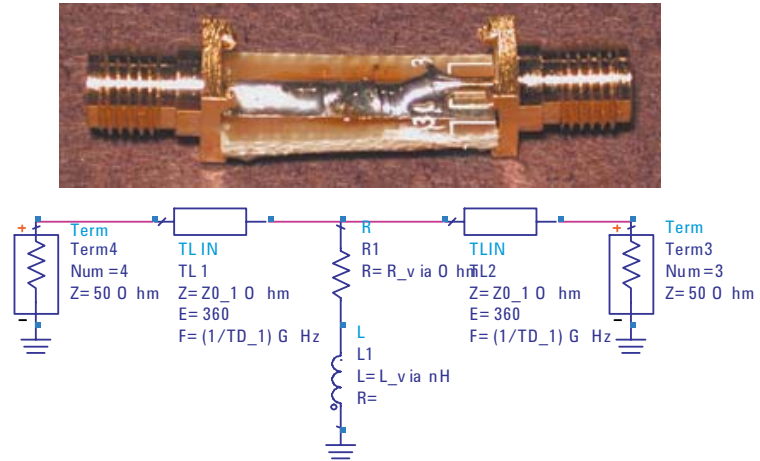


Figure 16. De-embedding the fixture

This parameterized model is used to simulate the 2-port S-parameters. Compare them to the measured results and optimize the parameters in the model to get agreement with the measurement. There are only four parameters, the impedance of the lines, their time delay and the R and L of the via model.

Figure 17 compares the measured and modeled S21 response with no transmission line in the model with the measured response. Using a value of 0.54 nH, the simulated response closely matches the measured response from 1 MHz to about 200 MHz. Beyond 200 MHz, there is noticeable deviation between the phase of measured and modeled S21. This is what gives rise to the apparent frequency variation in the inductance.

If we do not take into account the fixturing, the phase of S21 causes the phase of the impedance to vary and the phase of the impedance causes the imaginary part of the impedance to deviate from linear increase with frequency, which is the behavior of an ideal inductor. It is only possible to get good agreement for the ideal inductor model and the measured response, at frequencies below when the phase length of the fixture is a small fraction of a wavelength.

However, by including this phase length in the model, excellent agreement can be obtained all the way up to the bandwidth of the measurement. In this case, the measured and modeled S21 match from 1 MHz to 3 GHz, using the simple model of an ideal transmission line and an ideal inductor. The value of the series resistance was less than 1 mOhm, to fit the data, and can be ignored.

The measurement is consistent with a via looking like an ideal inductor with an inductance of 0.54 nH. This is very close to the earlier estimate of 0.5 nH for this 064 mil long, 25 mil diameter rod.




Figure 17. De-embed fixture by modeling

The Power Distribution Network (PDN)

The most important application where low impedance measurements play a critical role is in the power distribution network, or PDN. This is often also referred to as the power delivery network. The PDN is composed of all the interconnects from the voltage regulating module (VRM) which supplies a tightly regulated low voltage source through all the wires, vias, planes and packages, to the pads on the chip. The PDN includes the interconnects which transport the power and ground return. Its purpose is to keep the voltage across the power pads of the chip constant.

If the voltage across the power pads of a chip drops below the specification, typically, tighter than 5% variation is tolerated, the gates will switch more slowly and may cause timing problems which result in bit failures. If the voltage is too high, it may cause reliability problems, rapid aging of the devices or breakdown of the gate oxide.

Due to variations in the microcode that determines which gates switch and when, the current draw by the chip can vary from DC or a few Hz, all the way up to a few times the clock switching frequency. The impedance of the PDN must be below the target value from DC to above the clock frequency.

Unlike signal paths, which usually have a single-ended impedance of around 50 Ohms and a specific target impedance, PDN impedance is very low, often below 10 mOhms, with the goal to be as low as possible or at least no greater than a target value. This is why the ability to measure low impedances is so important.

Impedance design goals for the PDN

You can estimate the impedance design goal for the PDN with a simple analysis. The impedance of the chip and the interconnect of the PDN are in series. The VRM voltage is across both of them. For most of the voltage drop to be across the chip, the PDN should have a very low impedance compared to the chip. As currents through the chip change, due to different microcode running, and the impedance of the chip fluctuates, you want to keep 95% of the voltage drop across the chip, which means less than 5% of the voltage is dropped across the PDN interconnect. For this to be the case, the impedance of the PDN needs to be less than 5% the impedance of the chip. And, this needs to be the case from DC up to the full bandwidth of the switching currents in the chip.

The specification to keep the maximum impedance of the PDN below 5% of the chip, is easy to estimate. The target impedance is the rail voltage times the percent of ripple allowed, divided by the amplitude of the current. Since the impedance of the PDN should be less than 5% the impedance of the chip to the full bandwidth, you need to know the amplitude of the current at the highest frequency components.

Of course, the spectrum of the current through the chip is going to vary with different applications. In principle, you need to know the spectrum in order to create a realistic target impedance. This is only possible from measuring the performance of the chip under all operating conditions, or from simulating the chip performance using a good SPICE model. Neither of these are routinely possible. In practice, to determine the maximum current amplitude at the highest frequency, you have to make some guesses based on experience.

A typical example might be a Vdd supply of 1v and a maximum current through the chip of 2 A. You really need to know what the transient current is, the amplitude of the current, at the highest frequency component. As a rough estimate, the maximum transient current is at most half the worst case DC current. Translated into the frequency domain, the current amplitude at any frequency is roughly about half the DC worst case value.

This assumption allows you to estimate the target impedance as about 10% the V_{dd} supply divided by the worst case current, or in this example, about 50 mOhms. This is the target impedance, you want to keep the impedance of the power distribution network below this target impedance.

The key assumption here is that the worst case transient current is 50% the max current. In a logic device that has a low quiescent current, and suddenly turns on with all of its gates switching all the time, like in special control devices, the worst case transient current is probably going to be closer to the max current. However, in more typical cases, the worst case transient current draw is probably much less than 50% the max current, so this estimate of the maximum amplitude of the current at any possible frequency component probably has a bit of margin factored into it.

Typical case:

Keep 95% voltage drop across chip, < 5% voltage drop across PDN:

Keep $Z_{PDN} < 5\% Z_{chip}$,
from DC to BW of currents

$V_{dd} \sim 1\text{ v}$

$I_{max} \sim 2\text{ A}$

Assume transient current is 50% of I_{max}

$$Z_{target}(f) = \frac{\text{Voltage rail x ripple}}{I(f)}$$

$$Z_{target} = \frac{V_{dd} \times 0.05}{0.5 \times I} = 10\% \times \frac{V_{dd}}{I}$$

$$Z_{target} = 10\% \times \frac{V_{dd}}{I} = 10\% \times \frac{1}{2} = 50\text{ m}$$




Figure 18. Impedance design goals for the PDN

Measured impedance of a voltage regulating module (VRM)

All the components that make up the VRM must have an impedance that is below this target specification, starting with the output of the VRM. If the output impedance of the VRM is above the target specification, it almost doesn't matter what the rest of the PDN looks like. This is why measuring the impedance of the VRM is so important. It also sets the limit to how high a frequency the VRM is able to provide the low impedance and where the board-level PDN begins to play a role.

The output of a VRM is typically very low, on the order of a few milliOhms. To measure this low an impedance requires the two-port technique. In addition, because the VRM plays its dominate role in the frequency range from DC to about 100 kHz, it is below the normal frequency range of a VNA. To perform these low frequency and low impedance measurements, a customized, low frequency network analyzer from Ultimetrix can be used.

In this example, the Ultimetrix P4800 power delivery network analyzer was used to measure the impedance of a VRM while it was under a 1 A load, from 10 Hz to 40 MHz. The measured values show a low impedance below about 10 kHz of 0.03 Ohms. Above about 100 kHz, the impedance begins to increase, and looks very much like an ideal inductor. Along with the measured impedance of this VRM, is the simulated impedance of an ideal 27 nH inductor. This illustrates that a model for this VRM might be a series resistor of about 0.03 Ohms and a series inductor of 27 nH. This model would be used in the SPICE simulation of the PDN as the starting place at the lowest frequency.

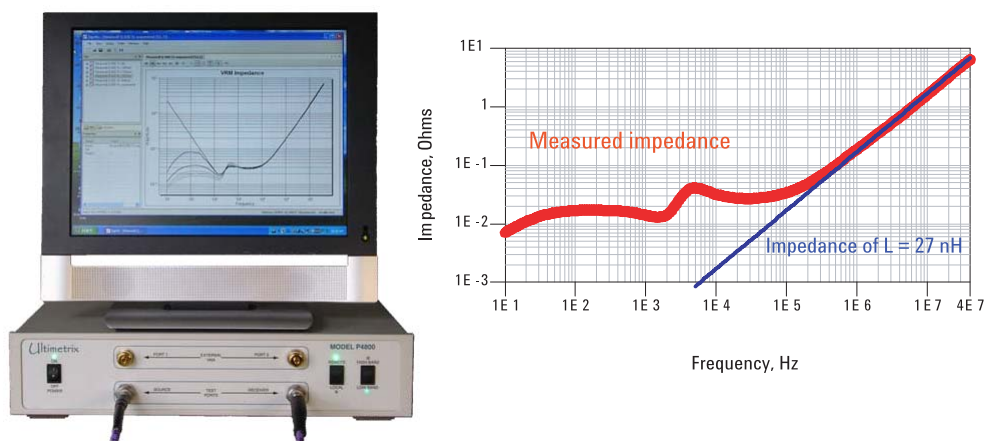


Figure 19. The Ultimetrix P4800 power delivery network analyzer shows measured impedance of a VRM

The output impedance of this particular VRM varies based on the load it sees. In general, the higher the load, the more current it draws and the lower the output impedance. Since the measurement is an AC measurement, it can be performed while the device is driving a DC current.

In these four examples, the output impedance was measured as the DC current load varied. The VRM powered off, the VRM turned on, but with no current draw to observe the passive impedance, a 1 A load and a 2 A load.




Figure 20. The VRM powered off

Figure 20 shows that at a low frequency the unpowered impedance profile looks like a capacitor. For the impedance to be 2 Ohms at 10 Hz just from capacitance, means the capacitance must be about 10,000 uF. This is the bulk capacitance on board the regulator. The equivalent series resistance of the capacitors, being electrolytic, is about 30 mOhms. Above about 100 KHz, the impedance begins to increase linearly in frequency, matching the profile of an ideal inductor with an inductance of about 27 nH.

The unpowered impedance profile, when compared to the various powered profiles, shows that the active part of the VRM, the regulation part, can only control the output impedance, for frequencies below about 1 kHz. Above this frequency, it doesn't matter whether the VRM is turned on or not, it's the passive elements of the VRM that are keeping the impedance low.

As the current draw increases, the impedance of the VRM drops. From open, to 2A, the output impedance drops from about 80 mOhms to as low as 18 mOhms.




Figure 21. The VRM without current draw shows the passive impedance

The phase of the impedance profile also tells a story. With power off, the phase of the impedance is nearly -90 degrees at low frequency, showing capacitive behavior, while at higher frequency it is inductive, with a phase of nearly 90 degrees. At the mid range, the phase is nearly 0 , as expected from a resistor.

When powered on, the phase at low frequency is slightly inductive. This is all related to the regulator design and the feedback circuit used to maintain constant voltage. The peak in phase at about 3 kHz is a measure of the highest response frequency of the regulation circuit. Beyond 3 kHz, it is the passive components in the regulator, the electrolytic capacitors, which are providing the low impedance.

Extracted R, L of 2A loaded VRM

When the VRM is drawing a 2 A load, a simple model for its output impedance is a series resistor and inductor. Using the measured impedance, the value of the series resistance can be read directly as the real part of the measured impedance, while the imaginary part of the impedance can be read as the reactance, from which the inductance can be extracted.

Without going through any more complicated model fitting, the resistance and inductance can be read directly off the front screen of the P4800 from the real and imaginary components of the impedance. In this way, you can read the equivalent values as 30 mOhms and 27 nH directly off the plots. Of course, their values are not perfectly constant because the model for the VRM is more complicated than a simple series RL circuit, but, this is a very good approximation. These are the terms that can be used in a SPICE simulation for the entire VRM.




Figure 22. 1 amp load of a VRM




Figure 23. 2 amp load of a VRM




Figure 24. Another example of VRM used in digital applications

Here is an example of another VRM used in digital applications. The impedance was measured using the P4800 for the two regulator states of off and on. When off, the RLC of the passive elements that make up the regulator system is measured. When powered on, the low frequency impedance is pulled dramatically lower by the regulation feedback circuitry.

The impedance, when on, is kept below about 5 milliOhms up to about 1 kHz. This is a measure of the response of the regulator itself. Beyond 1 kHz, the regulator is not able to respond fast enough to keep the impedance at its low value, and it is the passive elements associated with the regulation system which is really providing the low impedance. For this particular regulator, the feedback loop bandwidth can be adjusted slightly, trading off the on value impedance.

Fitting simple RLC model to VRM

When off, the impedance of the VRM shows information about the passive elements. This can be modeled with a simple RLC circuit and the agreement is very good. The fit corresponds to 4 milliOhms of ESR, 5.4 nH of ESL and 810 uF of capacitance. These values are typical of a very good tantalum capacitor, or possibly a few tantalum capacitors in parallel.

In the on state, the impedance matches a series resistance of about 5 milliOhms with a series inductance, but corresponds to the response of the feedback circuitry in the VRM. The powered on model would have this RL in parallel with the RLC of the powered off state. This is the model that would be used in the SPICE simulation of the PDN, incorporating the VRM.

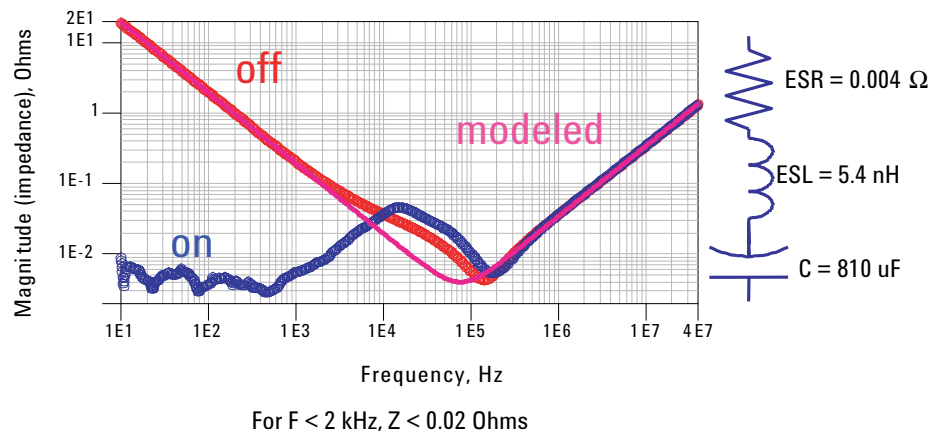


Figure 25. Fitting Simple RLC model to VRM

12 to 1.2 v supply application

The final example of a VRM is from Power-One. It is designed as a point of load regulator, able to handle a wide range of input voltages from 5 v to 14 v, and output 1.2 v. It is primarily designed for high current processor applications. This experimental unit has two large electrolytic capacitors. One is used on the high voltage input side, and the other is used for the low voltage output side, in addition to the various SMT tantalum and ceramic capacitors.

The coaxial pigtailed are mounted directly across the large, electrolytic capacitor on the output side.




Figure 26. 12 to 1.2 v supply application

Measured impedance powered on and powered off

In the powered off state, the impedance profile has the appearance of two capacitors in parallel. In the powered on state, the regulation circuitry pulls the impedance down to 10 mOhms from 10 Hz to about 1 kHz. This is the general frequency range in which DC regulator circuits have a linear response. Beyond this frequency, their output impedance begins to increase, until the impedance of the passive circuitry takes over.

The output impedance of this regulator can be adjusted, trading off feedback frequency response. Above about 8 kHz, the powered on and off impedances are identical. Figure 27 clearly shows that above 8 kHz, the impedance of the VRM is all about the passive circuits. This is why an accurate model of the passive circuits will be useful to include in the SPICE simulation of the VRM.




Figure 27. Measured impedance powered on and powered off

Modeling the VRM

In this case, the two resonance dips at 40 kHz and 300 kHz, with the parallel impedance peak at 200 kHz, indicates a good candidate model to describe the passive decoupling network might be two capacitors in parallel. The circuit model is two RLC circuits in parallel, each set of RLC elements with different parameter values.

Optimize the 6 elements to give good agreement in the powered off state between the measured impedance and the simulated impedance from the model. When the agreement is excellent, you can use the model to extract the values of the capacitor elements.

The large capacitor has values of 6 milliOhms, 14.7 nH and 910 uF. This is the large electrolytic capacitor attached to the output of the low voltage side. The second capacitor, with values of 5 milliOhms, 5.5 nH and 34.3 uF, is the ceramic capacitor mounted to the board. This is very typical of a mounted capacitor that is not well optimized for low impedance.

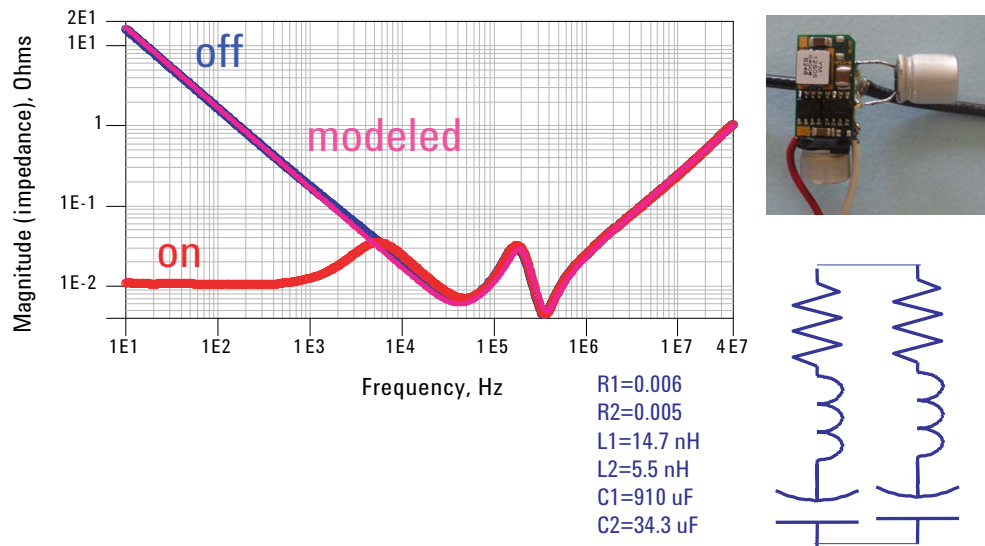


Figure 28. Modeling the VRM

Summary of VRM Behavior

These examples of VRM performance illustrate a few common performance features. The regulation circuitry can reduce the output impedance to the milliOhm range at frequencies below about 1 kHz, depending on the feedback loop. Beyond 1 kHz, it is the passive network on the VRM which keeps the impedance low.

This suggests, to really minimize the VRM output impedance at high frequency, all the techniques used at the board level for low inductance capacitor mounting should be applied to the VRM. This includes: power and ground planes, short surface traces to the capacitors, multiple vias.

It is the combination of the regulation circuitry for low frequency and the passive network at high frequency that the VRM impedance can be kept below a target value.

The P4800 PDNA can also interface directly with a VNA to span not only the 10 Hz to 40 MHz range of the instrument, but extend all the way to the upper range of the VNA. This connection enables a dynamic range of more than 9 decades in frequency, with an impedance dynamic range from 1 milliOhm to more than 100k Ohms, or 8 decades in impedance.




Figure 29. Ultimetrix P4800 With Optional VNA 10 Hz to 18 GHz

Impedance of a capacitor

Here is an example of the measured impedance of a small ceramic chip capacitor from 10 MHz to 10 GHz, showing the behavior of a simple RLC series circuit across this entire frequency range. The capacitance can be extracted from this measured data as about 1 μF with a equivalent series resistance of about 0.015 Ohms and a series inductance of about 200 pH.

This mounting technique for a capacitor is great for measuring, but is not related to how the capacitor is mounted when used in the PDN, so it represents only an approximation to the partial self inductance of the capacitor body. In practice it is very hard to achieve mounting inductances even close to this value.

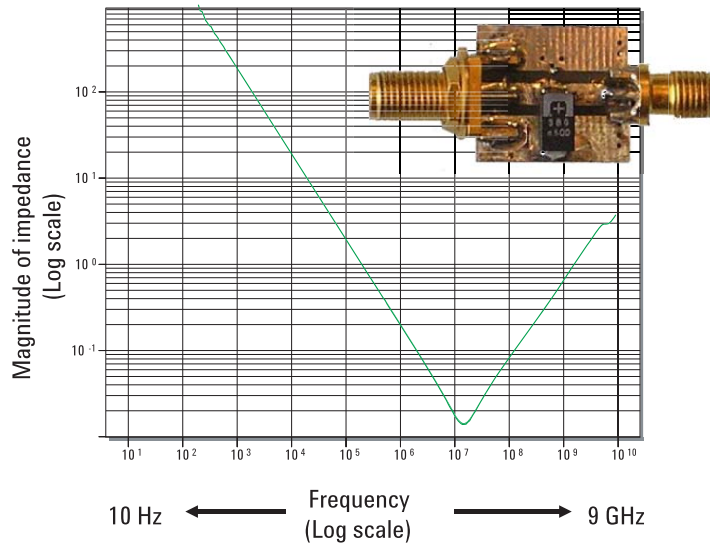



Figure 30. Impedance of a capacitor

Decoupling Capacitors

This 2-port technique is ideal for measuring the impedance of decoupling capacitors and for exploring how the mounting technique affects the equivalent series inductance. The fixture shown in Figure 31 was provided, courtesy of X2Y Corp. It consists of a 4 layer board with a pair of power and ground planes at the core. The other two layers are capture pad layers on the top and bottom surfaces.

Pads are arranged on the top surface to connect various decoupling capacitor sizes and types to the planes below and connections to the two planes are made on opposite sides of the board with SMAs. When a capacitor is mounted to the top surface, vias connect the pads to the planes. There is approximately 0.7 inches from the SMA to the capacitor location, so the maximum accurate frequency should be about 420 MHz.



Len ~ 0.7 inches, $F_{max} \sim 0.3/0.7 = 420$ MHz




Figure 31. A capacitor fixture board is a 4 layer board with a pair of power and ground planes at the core

A 1-port measurement of the impedance of the capacitor includes the series inductance of the SMA connector, the inductance of the planes, the spreading inductance of the vias, and the via inductance. With uncertainties in the SMA length, plane inductance and spreading inductance, it is difficult to extract the inductance associated with just the mounting of the capacitor.

However, in the 2-port technique, the structures that contribute to the impedance are the capacitor, the vias, and a small amount of spreading inductance in the planes. The 2-port impedance is much more sensitive to the mounting inductance of the capacitors.




Figure 32. Using a 2-port technique to measure capacitor fixture boards

Behavior of real, 220 nF, 0603 MLCC capacitor

In this first example, an 0603 multi layer ceramic capacitor, sometimes referred to as a MLCC, is mounted, as shown in Figure 33, to pads which connect to the vias to the planes below. The pads are minimal length. This is not the recommended layout for a production board, because there is no solder mask dam between the capacitor pads and the via. There is nothing to prevent solder from the pads wicking down the via hole, and starving the pads for solder. This could result in an unreliable capacitor assembly. However, for a demonstration vehicle, it is not a problem. The picture on the right shows the capacitor soldered to the pads.

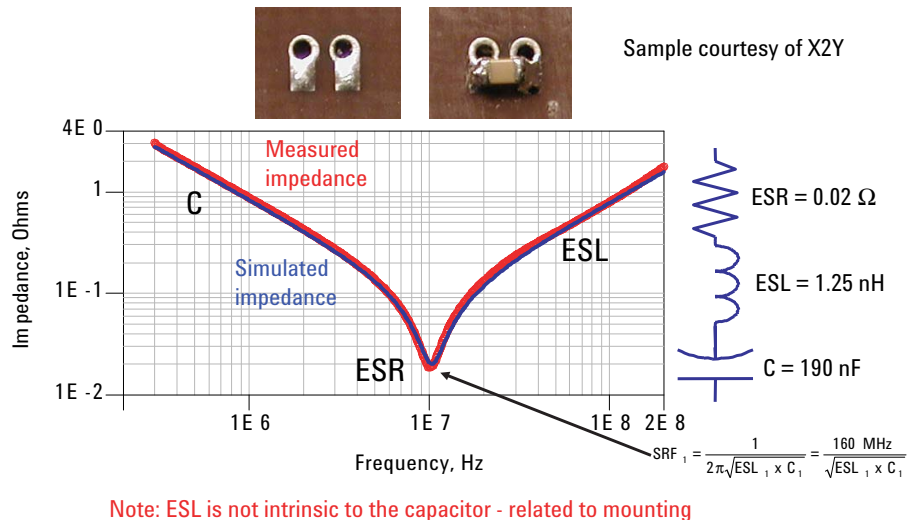


Figure 33. Multilayer ceramic capacitors

The measured impedance, using the 2-port technique is shown as red circles, measured from 300 kHz to 200 MHz. The low frequency behavior looks like an ideal capacitor, while the high frequency behavior looks like an ideal inductor.

The blue trace is the simulated impedance for a model of the capacitor based on a simple series RLC model. In this case, the Agilent Advanced Design Software (ADS) was used to perform the SPICE simulation. In this model, a value of 190 nF was used for the ideal capacitance. This influences the behavior at low frequency- the drop in impedance as frequency increases. The higher the capacitance, the lower the level of the curve, but the slope will always be the same. The impedance drops like 1/f, which on a log log scale is a slope of -1.

The equivalent series inductance, the ESL, was chosen as 1.25 nH to match the measured response. This is the mounted loop inductance of the capacitor, that includes the inductance of the capacitor, the pads, the vias and a small amount of spreading inductance in the planes. Because of the way the device was fixtured, the measurement is not so sensitive to most of the spreading inductance of the planes and the inductance of the planes themselves. It is the ESL that limits the high frequency impedance of the capacitor element. A lower ESL will bring the impedance down at high frequency.

This simple RLC model is an excellent model for the impedance behavior of this real capacitor from low frequency to very high frequency. This demonstrates that with the right parameter values, this model can be used in a SPICE simulation of the board to obtain realistic results. By fitting the three parameter values of the R, L and C to the measured data, the capacitance, equivalent series inductance and equivalent series resistance can be extracted.

It is clear that the capacitance of the capacitor has no influence on the high frequency impedance of a real capacitor. At high frequency, the impedance of the capacitor is all about the equivalent series inductance, ESL of the capacitor. The ESL is not an intrinsic measure of the capacitor itself as it is strongly dependent on how the capacitor is mounted to the planes. Do not tie an ESL value to a particular capacitor. Some capacitors can be integrated to boards with a lower ESL than others, but any capacitor can be made to have a very high ESL by inefficient attach designs.

The third term in the model is the equivalent series resistance, the ESR. This is ultimately what limits how low an impedance a capacitor will have. The impedance of a capacitor can never be lower than its ESR. This term depends on the internal metallization of the capacitor, how thick it is and how many planes are in parallel. In this case, the measured ESR is 0.02 Ohms. The combination of the capacitance and the ESL gives rise to a resonance in the impedance behavior of a real capacitor.

The self resonance frequency, SRF, is the frequency where the L and C combine together to look transparent, and cancel out. At this frequency, the capacitor will show the lowest impedance it has, the ESR. It is numerically 160 MHz divided by the square root of the L times C. In this case, an estimated 10.4 MHz, it actually measures at about 10.5 MHz.

For some capacitors, L and C are frequency dependent and decrease at higher frequency, due to skin depth effects

Behavior of real, 200 nF, X2Y capacitor

Figure 34 shows the measured impedance of an X2Y capacitor, with a total of 6 vias to the planes below. The combination of 6 vias and the internal design of the capacitors helps to decrease the mounted inductance of this capacitor compared to conventional capacitors.





Figure 34. Behavior of real, 200 nF, X2Y capacitor

However, even with this design, the measured impedance profile still fits the simple RLC model extremely well. It's just that the values of the R and L are different. In this capacitor, the ESL is 0.42 nH. This is to be compared with 1.25 nH for the 0603. This is a reduction of a factor 3 in the series inductance. This translates into a reduction in the impedance at high frequency of a factor of 3 and could enable a parts count reduction by a factor of 3 as well.

Multiple capacitor samples

Using three different sample boards, the impedance of different mounting configurations was measured and compared. In this example, 0603 capacitors were mounted with 2 vias and 4 vias, while an X2Y capacitor was mounted with 6 vias. At low frequency, the impedance of each configuration is the same, as they all have roughly the same capacitance. However, at high frequency, the different inductances of each configuration show up as having a different impedance.

Superimposed on the measured impedance of these different structures with the 2-port technique is also a measurement of the 2-via capacitor, using the 1-port technique. It can be seen that the impedance at low frequency, due to the capacitance of the capacitor, is the same as the 2-port technique. However, the inductance comes out very high, and the resonance frequency is at a lower frequency than for the 2-port method. Even the ESR comes out too high, since there is series contact resistance in the 1-port measurement.



Test vehicle, courtesy of X2Y

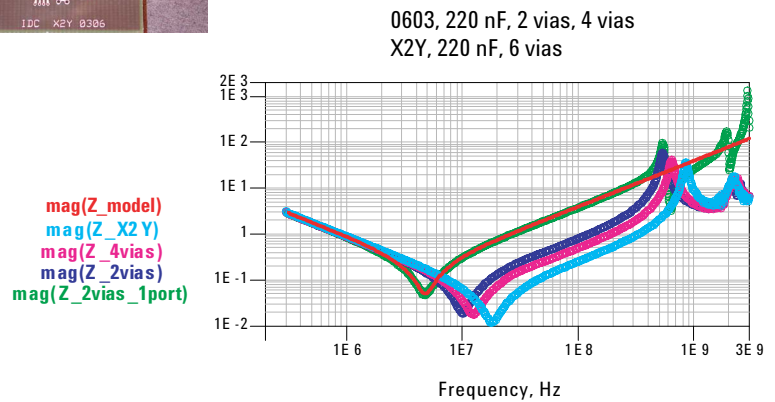



Figure 35. Multiple capacitor samples

Examples of measured total mounted inductance of capacitors

Because you already know that this simple RLC model is a good model for this data, you can take the high frequency end of the impedance, and extract the inductance directly, assuming the reactance is all inductive.

The vertical scale is 2 nH. You read the ESL of each configuration right off the plot. In addition, measurements were taken with the 2 vias and 4 vias shorted on the top surface with a small shorting block, roughly the size of the capacitor. The difference between this inductance and the capacitor inductance is a measure of the higher inductance of the capacitor due to its multiple layers and the farther the effective distance of the currents from the nearest plane below.

At a glance, this shows the relative mounted loop inductances of the decoupling capacitors and the significantly lower ESL of an X2Y capacitor over a conventional 0603 capacitor.



Test vehicle, courtesy of X2Y

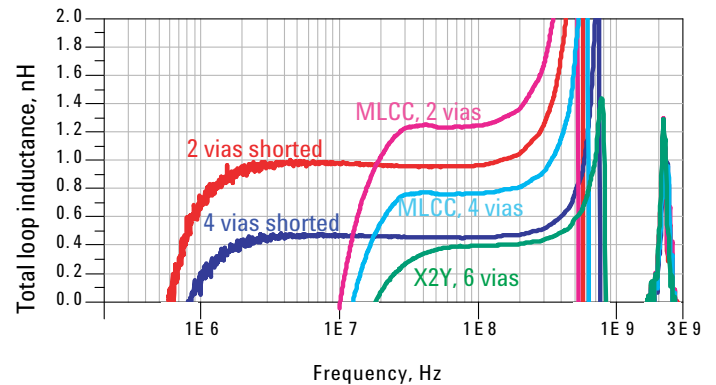


Figure 36. Examples of measured total mounted inductance of capacitors

Planes and Capacitors

The impedance of a pair of planes can be measured by attaching two SMA connections between the planes, at locations within $1/20$ th of a wavelength apart. Figure 37 shows an example of the measured impedance of a power and ground plane pair in a circuit board about 6.5 inches long. The dielectric spacing between the planes is about 64 mils. This is the impedance between the two planes, from one edge. The effective length that will influence the resonant frequency is 6.5 inches. The first resonance should be at about $1.5/6.5 = 230$ MHz.

The red trace is the impedance of the bare board with nothing on it. You can see the high impedance at low frequency, due to the intra-plane capacitance. This drops down to a minimum value at about 210 MHz, pretty close to the estimate, at the first resonance where the impedance is a minimum. Above this first resonance, you see the typical peaks and valleys of transmission line behavior. Of course, the specific peaks and valleys measured depends on where you probe the impedance of the planes. In this case, probing from the edge to measure the lowest frequency resonances.

The blue trace is the impedance of the same board with one, 85 nF capacitor on it. The 2-port technique also eliminates the artifact of the series inductance of the test leads to the board. In this way you are looking at the intrinsic impedance of the planes. At low frequency, you can see the drop in impedance when the capacitor is added. The impedance peak in the blue trace, called a parallel resonance, is the interaction between the inductance of the capacitor and the capacitance of the planes. At and above the first board resonance, you can see the impedance of the capacitor and plane together is identical to just the plane by itself. This observation gives rise to statement that above the plane's first resonance, the impedance the chip will see is independent of the capacitors on the board. The series impedance of the planes limits the impedance the chip can see looking into the planes.

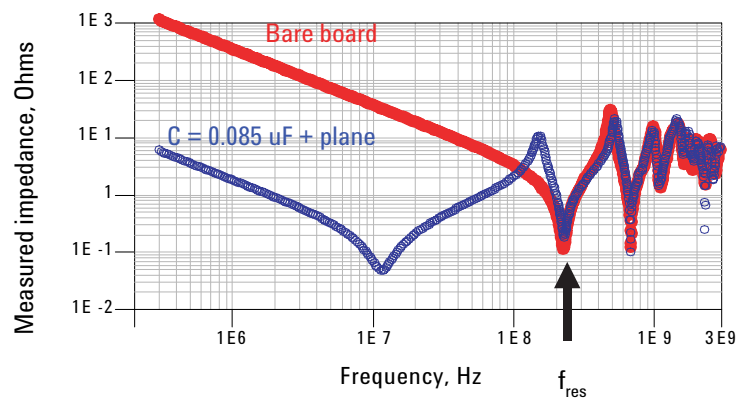


Figure 37. Above plane's first resonance, impedance is independent of the capacitor

Multilayer board example: 5 x 10 in, 5, 2 mil thick dielectric layers

The next example of a pair of planes is courtesy of Sun Microsystems™, and GigaTest labs. The board being measured is a 12 layer board, with 5 power and ground plane pairs. Each pair has a dielectric spacing of only 2 mils. The top two plane pairs are connected to the top component layer using vias. These two plane pairs are connected in parallel.



This board is designed to have 0508 Inter-Digitized Capacitors (IDC) surface mounted using via in pad, a very low loop inductance combination. Using microprobes means the electrical length of the fixturing to the DUT is measured in mils, and the impedance measurements are accurate well into the GHz regime.



Figure 38. Probing the the PCB structure and layer layout

How to Set Up a Measurement System

The GigaTest probe station enables a variety of probing methods, including probing of the same via from opposite sides. This probe station allows precision alignment of the microprobes to pads anywhere over the field of the board, and calibration right to the tips of the probes.




**GigaTest Labs
PCB probe station**

Figure 39. Measurement system: GigaTest Labs 3030 Probe Station, Agilent 8753 VNA, Agilent ADS Software.

PCB fixture layer stack up


There are three different configurations for the probes, each of which measures a slightly different response to the via.

When both ports are on the same pad, the 2-port impedance measurement is sensitive to the via loop inductance from the top pad to the planes, the spreading inductance of the planes, and the loop inductance of the planes themselves. This is usually dominated by the loop inductance of the via.




**Figure 40. Both ports on the same pads:
Via loop inductance, spreading inductance, plane inductance**

When the two ports are on the same via, but on opposite sides of the board, the impedance measurement is no longer sensitive to the loop inductance of the via, but it is sensitive to the spreading inductance and the inductance of the planes.



**Figure 41. Each port on the same via, opposite sides of the board:
Spreading inductance, plane inductance**

Finally, when the probes are on adjacent vias, the impedance measurement is sensitive to the plane impedance, and little bit of the spreading inductance, depending on the proximity to the vias.



**Figure 42. Each port on the same planes, adjacent pads:
Plane inductance**

Measured impedance: 1-port and 2-port, planes isolated

The impact of the contact impedance of the probes can be easily seen by comparing the measured impedance of the plane pairs from the top surface of the via pad using the 1-port method and the 2-port method. When 1-port is used, the measured impedance has the contact resistance to the pads, with solder on their surface, and the probe tips. In this example, it is roughly 0.2 Ohms.

Using the 2-port method, the series resistance of the loop impedance of the via pair can be measured directly as the minimum resistance at about 75 MHz as 5 milliOhms.

The general behavior of the impedance plot in the 2-port case is that of an ideal RLC circuit. The resonances corresponding to the plane resonances are just barely discernable in the impedance behavior. This is because they are severely damped by the series resistance of the planes. For a lossy transmission line, the attenuation from resistive loss is related to the series resistance divided by the characteristic impedance. With wide conductors and very thin dielectric, the characteristic impedance is very low, so the attenuation is very high.

The near-linearly increasing impedance as you go up in frequency is due to the loop inductance of the vias which is in series with this measurement. From the impedance, the loop inductance of the via pair can be estimated as about 80 pH. In this particular example, the high frequency impedance, dominated by the loop inductance of the via pairs, is the same for the 1-port as 2-port technique. This suggests that in this example, the over travel of the 1-port probe was very close to its calibrated state.






Figure 43. Above 100 MHz, the impedance is dominated by the via pair loop inductance

Measured impedance from top surface: 1-port and 2-port, planes shorted

In this next example, the planes are shorted at the locations where the capacitors will be attached. In this particular case, a 1-port measurement of the planes shorted is compared to the 2-port measurement, again with the two probes touching the same via pads.

The 1-port measurement of the shorted planes shows a contact resistance of 0.5 Ohms. This high value is due to the probe tips touching solder, which has an oxide build-up and causes high and variable contact resistance. The probes were scrubbed an extra amount to try to rub off the oxide coating, but this did not help.

In fact, the heavy scrub used is further seen in that the impedance measured by the 1-port method is actually lower than for the 2-port method. When the 1-port probe was calibrated, there was a particular distortion in the shape of the probe tip due to the over travel. This defined the zero inductance position. When this measurement of the shorted planes was performed, the probe tip was distorted even more, compressing the tip and decreasing the loop inductance of the tip. This gave an artificially low value to the measured inductance.



1. Above 100 MHz, impedance is independent of whether the planes are shorted or not
2. Resonances are weak because of the low Z_0 of the planes and the R/Z_0 attenuation

Figure 44. Measured impedance from top surface: 1-port and 2-port, planes shorted

The 1-port method gives a loop inductance at high frequency of about 60 pH, while the 2-port method gives exactly the same loop inductance in the shorted as the open case of 80 pH. The difference of 20 pH is the uncalibrated inductance of the probe tip, in the 1-port measurement.

Figure 44 also shows that at high frequency, the measured impedance of the planes is dominated by the loop inductance of the vias, and is independent of whether the planes are open or shorted, or are covered with decoupling capacitors. This is a limitation of probing from the top side of the planes through the vias.

Probing from the same and opposite sides of the via pairs

When measured from the top, either with a 1-port or 2-port method, the impedance measurement is dominated at high frequency, not by the planes, but by the via loop inductance (roughly 0.08 nH).

One way around this artifact is by using the 2-port technique and probing from opposite sides of the same via pair. In this case, port 1 forces the current through the via loop and into the planes, and port 2 measures the voltage created across just the planes. This voltage does not include the voltage drop across the via to the top surface, but only the spreading inductance in the planes and the plane inductance.

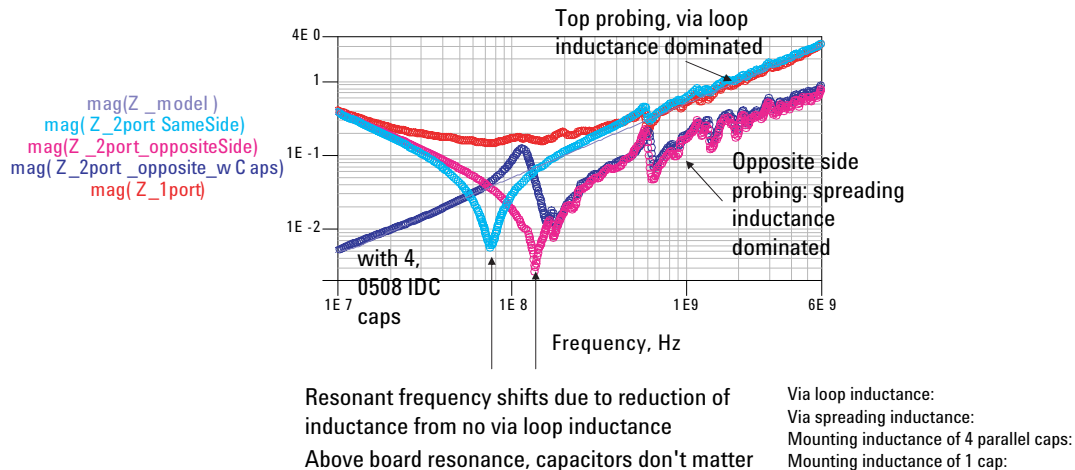


Figure 45. Probing from the same and opposite sides of the via pairs

At high frequency, the impedance from opposite sides of the board looks inductive, with a loop inductance of about 20 pH. This is the spreading inductance of the vias into the planes. At high frequency, there is some contribution from the plane resonances, at 700 MHz and beyond, but most of the impedance is dominated by the spreading inductance of the via contacts into the planes.

The measured impedance from opposite sides of the board, is compared to a bare board, and to a board with 4, 0508 IDC capacitors mounted on it. At low frequency, the impedance shows the capacitive nature of the pairs of planes. When the capacitors are added to the board, the low frequency impedance drops considerably, due to the very low impedance of the capacitance. However, the impedance behavior at low frequency is dominated by the loop inductance associated with the capacitors and the planes.

Below about 70 MHz, the impedance of the planes and the IDC capacitors matches a loop inductance of about 80 pH. This is the loop inductance, the spreading inductance of about 20 pH, and the loop inductance associated with 4 IDC capacitors in parallel. The loop inductance of just the capacitors is about 60 pH. Given four of them in parallel, the loop inductance of each one must be about 240 pH. This is the mounted inductance of the IDC capacitors, a very low value compared to conventional 0603 capacitors, which were shown earlier to be about 1.25 nH.

By using low inductance capacitors, and optimized mounting techniques, it is possible to achieve ultra low mounted inductance of capacitors. This means fewer parts are needed to achieve the same target impedance, or lower a target impedance and greater margin can be achieved with the same part count.

2-port fixturing of bare circuit board

As an alternative to using microprobes, it is possible to make two connections to the same surface of a board, but to adjacent pads, using coaxial pigtailed. In this configuration, the 2-port impedance measurement is not sensitive to the via loop inductance or the spreading inductance of the via contacts to the planes, but is sensitive to the impedance of the planes. This technique can be used to measure the impedance profile of a board with and without capacitor components.




Figure 46. 2-port fixturing of bare circuit board

In the case of a populated board, the capacitor must be removed from the pads that are measured. The signal pin of the coax goes to the + side of the rail and the shield goes to the ground side of the PDN.

These two cable ends connect to port 1 and port 2 of the VNA. To enable accurate measurements to high frequency, the phase length of the cables should be small, and the cable pigtailed kept short. Alternatively, if the cable length is well known, an S parameter file can be created for the cables and they can be de-embedded from the measurement.

Measured impedance of bare board: 1.2 v planes

This is the measured impedance of the power and ground planes for the 1.2 v rail in the board. This is for the unpopulated, bare board, showing the capacitance at the low frequency and the resonances above about 170 MHz. In this case, you'll see that the first board resonance is about 200 MHz. The impedance of the loaded board shouldn't be affected much beyond about 200 MHz in this example, and the highest frequency for the on-board PDN impedance could be selected as roughly 200 MHz based on this measurement. To this board all the decoupling capacitors are added, and then the voltage rail is re-measured.






Figure 47. Measured impedance of the power and ground planes

Example from Altera Stratix II GX signal integrity kit

The capacitor placement on this board, supplied by Altera Corp, is also optimized for low inductance. The decoupling capacitor pads are located on the back side of the package attach surface. The crossed rows of decoupling capacitors located within the power and ground pads at the center of the package can be seen in the artwork. This minimizes the spreading inductance as it gets the capacitors close to the BGA solder balls. The other capacitors are located around the periphery of the package.



Placement adjacent to pwr/gnd solder balls



**0402 capacitors
Short surface traces
Wide surface traces
Vias close together**

Figure 48. Example from Altera Stratix II GX signal integrity kit

Most of the decoupling capacitors are 0402 devices, with a reasonably small footprint. The surface traces are short and wide. The vias have been rotated around to be closer together.

Altera Stratix II GX signal integrity board: 1.2 v planes ($Z_{\text{target}} = 84 \text{ m}\Omega$)

Figure 49 shows the measured impedance of the PDN, with the bare board and the capacitor components on the board. There is no VRM nor any package or chips on the board.

The capacitors have little impact above 200 MHz due to the plane resonance. Also, the impedance is approximately below the 80 m Ω target impedance at less than 200 MHz. The impedance toward the high end frequencies matches the impedance expected from an 80 pH inductor.

This is a technique that can be used to validate any PDN design. It is capable of measuring milliOhms of impedance up to the GHz frequency range.




Figure 49. Altera Stratix II GX signal integrity board: 1.2 v planes ($Z_{\text{target}} = 84 \text{ m}\Omega$)

Conclusion

The use of 2-port VNA techniques can open up an entirely new window to observe the low impedance behavior of interconnects. This applies most importantly to the power delivery network where low impedance over a wide frequency range is the most important design feature. Using 2-port VNA techniques can completely eliminate the artifacts associated with contact impedance of the probes or fixturing to the DUT.

When the impedance of the DUT is much less than 25 Ohms, such as less than 5 Ohms, the impedance of the DUT can be read right off the front screen of the VNA as 25 Ohms times S_{21} . However, with a little post processing of the data, the impedance can be read across the entire impedance range with no approximation.

The one limitation imposed by the fixturing is to keep its electrical length less than about 1/20th a wavelength.

When measuring the impedance of planes, the area where the two contacts are made will have a dominate impact on what exactly is measured. When both probes are on the same via pad, the series inductance of the via loop can often dominate the measured impedance. When opposite sides of the same via are measured, the spreading inductance of the planes is measured. When adjacent pads are measured, the plane impedance can be measured.

This technique is a powerful tool that should be in the tool box of every signal integrity and power integrity designer.

Web Resources

For more information, visit:

www.agilent.com/find/plts



Agilent Email Updates

www.agilent.com/find/emailupdates

Get the latest information on the products and applications you select.



Agilent Direct

www.agilent.com/find/agilentdirect

Quickly choose and use your test equipment solutions with confidence.



www.agilent.com/find/open

Agilent Open simplifies the process of connecting and programming test systems to help engineers design, validate and manufacture electronic products. Agilent offers open connectivity for a broad range of system-ready instruments, open industry software, PC-standard I/O and global support, which are combined to more easily integrate test system development.

Front cover photo courtesy of Ultimatrix.

Sun, Sun Microsystems, and the Sun Logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Remove all doubt

Our repair and calibration services will get your equipment back to you, performing like new, when promised. You will get full value out of your Agilent equipment throughout its lifetime. Your equipment will be serviced by Agilent-trained technicians using the latest factory calibration procedures, automated repair diagnostics and genuine parts. You will always have the utmost confidence in your measurements.

Agilent offers a wide range of additional expert test and measurement services for your equipment, including initial start-up assistance onsite education and training, as well as design, system integration, and project management.

For more information on repair and calibration services, go to

www.agilent.com/find/removealldoubt

www.agilent.com

For more information on Agilent Technologies' products, applications or services, please contact your local Agilent office. The complete list is available at:

www.agilent.com/find/contactus

Phone or Fax

United States:

(tel) 800 829 4444
(fax) 800 829 4433

Canada:

(tel) 877 894 4414
(fax) 800 746 4866

China:

(tel) 800 810 0189
(fax) 800 820 2816

Europe:

(tel) 31 20 547 2111

Japan:

(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Korea:

(tel) (080) 769 0800
(fax) (080) 769 0900

Latin America:

(tel) (305) 269 7500

Taiwan:

(tel) 0800 047 866
(fax) 0800 286 331

Other Asia Pacific Countries:

(tel) (65) 6375 8100
(fax) (65) 6755 0042
Email: tm_ap@agilent.com

Revised: February 5, 2007

Product specifications and descriptions in this document subject to change without notice.

© Agilent Technologies, Inc. 2007
Printed in USA, February 15, 2007
5989-5935EN



Agilent Technologies