

RFIC Inductor Measured vs. Simulated

Purpose of this document:

The purpose of this document is to assist you with resolving possible differences between simulated RFIC inductor results and values obtained from the design kit or from measurement.

This document is based on extensive practical experience. The author has worked for many years in RFIC inductor modelling, and worked closely with the modelling and measurement groups of different foundries.

Table of contents:

RFIC Inductor Measured vs. Simulated	1
1. Introduction	2
2. Low frequency response	2
2.1 DC Resistance	2
2.2 Low frequency inductance	6
3. The Q factor	8
3.1 Differential Q factor vs. single ended Q factor	8
3.2 Calculate differential Q factor from 2-port simulations	10
3.3 Special equations for center tapped inductors	11
4. Technology related differences	11
4.1 Substrate conductivity	12
4.2 EPI and pwell block	13
4.3 Looking closer at the substrate shunt path	14
4.4 Influence of the Sonnet box size on shunt path results	18
5. Summary	20
6. Change log	21

Document revised: 31. January 2012

Document revision: 1.1

1. Introduction


When I started to work with Sonnet in 1998, an experienced user told me *"If the Sonnet result does not agree with measurement, this means that the Sonnet model is different from what we have measured."*

Sounds trivial, but indeed most differences between simulated and measured are caused by a simulation model that is somehow different from the measured hardware. It's not random. There is always a reason for differences. Our goal is to identify and remove that difference and obtain accurate, reliable simulation results.

When we start with RFIC models for a new technology, we often compare Sonnet results to data from the design kit. Sometimes, we see differences in results. With this application note, we will try to identify what is causing the different results, and which results we should trust.

2. Low frequency response

From a two port simulation, we can separate the effects in RFIC inductors into **series** elements and **shunt** elements, as shown in the simple inductor model below. For very low frequencies, only the series path (R_s , L_s) is relevant. We can use that for checking the conductor properties.



2.1 DC Resistance

At DC, the inductor behaves like a series resistor. In case of doubt, checking the DC and low frequency resistance is a good starting point for troubleshooting.

For a Sonnet two port simulation of an inductor, or for measured s2p files, the series resistance can be plotted with the "resistance2" equation in the emGraph result viewer. For an inductor with center tap at port 3, the equation "resistance3" must be used, which calculates the resistance between ports 1 and 2 with all other ports open.

Now, we can check the Sonnet resistance for a sufficiently low frequency against the DC resistance that we have calculated manually, and against the measured or foundry models.

If there is a difference between the Sonnet low frequency resistance and the measured or foundry data, this can have multiple reasons:


- difference in conductivity or conductor cross section
- error in the Sonnet simulation, such as short circuit between the turns
- data that we compare to belongs to a different layout (yes, this really happens)
- measurement error from contact resistance¹

Now if we have these different results, but do not know which is the correct value, we can manually calculate the conductor cross section² and path length, and check that manually calculated DC resistance against the simulation and measurement results.

This is really low frequency with no RF or transmission line effects, so the math is simple. In most cases, the via resistance is small and the total resistance is determined by the conductivity of the planar metals. If we want to be very accurate, we can also include the resistance of series vias, which is specified in the foundry's process datasheet³.

If there is a large difference between the Sonnet model and the manually calculated value, with the same conductor properties, something might be wrong with the Sonnet model.

If the Sonnet resistance is larger, check if you have properly de-embedded the feedlines. If the Sonnet resistance is smaller, check for possible short circuit caused by coarse meshing, where mesh subsections from different turns cause an unintended connection. This can be easily identified if you run the analysis with current density enabled, and then look at the current distribution. You will also notice a difference in inductance if there is some short circuit.



¹ We will discuss this later, and see what options we have to double check the measurement data

² For stacked metals, with via arrays connecting the metals, the vias will carry only vertical current between the metals, and must not be included in the conductor cross section. Only for "via bar" where a solid block of metal is used as a via, that via bar can carry horizontal current and add to the effective cross section.

³ Usually specified per single via, so that we have to divide the data sheet resistance by the total amount of vias in the via array.


Differences could be caused by the Sonnet **cell size**. The cell size determines what line width Sonnet will use for analysis. If the cell size is very coarse, there might be some error from snapping the line width to a multiple of the Sonnet cell size. If the effective line width in Sonnet is much smaller or much bigger than the real width, this will of course change the effective cross section and change the DC and low frequency resistance. At higher frequencies, skin effect will push the current to the edges of the conductor, and an error in line width becomes less visible in series resistance.

To check for this error, you can reduce the cell size and do a comparison between both simulation results. It is sufficient to simulate at a single frequency, so that you can compare that one data point.

If the DC resistance calculated by Sonnet agrees with the value that was calculated manually from the geometry, but the measured or foundry data is different, this could be measurement error from **contact resistance**. In this case, you should try to get hands on the non-deembedded measurement data for the inductor and for the de-embedding standards, and check them for consistency and symmetry.

Working closely with different foundries on modelling and measurement, we have seen cases where perfectly symmetric calibration standards had a difference of >2 Ohm measured into the "short" standard on both sides, which is obviously wrong and was caused by contact resistance. Such a measurement error will have large impact on Q factor. The higher the inductor Q factor is, the more error we will see from small contact resistance in measurement.


The plot below shows the effect of a 1 ohm contact resistance on the Q factor.⁴




You might argue that contact resistance will be removed by measurement de-embedding, if it is always the same. Unfortunately, in real life, contact resistance varies.

⁴ This is for a 1 ohm contact resistance in the measurement, which causes a low Q value. But if the high contact resistance was in the de-embedding standard, instead of the measurement, that removes too much resistance from the raw data, and the measured Q factor would be too high. In real life, we have seen both cases.

The plot below shows the measured values for three samples of the same inductor, obtained from three different locations on the chip, during the same measurement session on an automated wafer prober station. This could be sample variation or contact resistance, we don't really know.



If we consider measured data to be the accuracy benchmark, now we have three different measurement results for the same inductor, and we don't know which of these curves is most accurate. In this case, it's a good idea to manually calculate the expected value, and compare it to the extracted series resistance as seen below.⁵



Unrelated to the measurement issues described above, here are more reasons why we could see differences in DC resistance between the simulation and the design kit data:

In some cases, the foundry might have **tweaked the process** without adjusting the inductor models, or without re-measuring the inductors.

And finally, we have also seen cases where someone made a mistake and the supplied data was obviously for a **different device**. This should not happen, but it does happen. In this case, we will see mismatch in multiple parameters, such as resistance and inductance and shunt elements to the substrate.


⁵ Surprised to see negative series resistance at high frequency? This is not real, of course. The negative effective resistance shows the limits of our simple equivalent circuit, where lossy substrate below results in a seemingly negative series resistance.

2.2 Low frequency inductance


The low frequency inductance is quite robust, and insensitive to most technology parameters. At higher frequencies, we will see parasitic capacitance as well as skin- and proximity effect change the effective series inductance. But at low frequency, the inductance can be estimated with reasonable accuracy (typically 2-3% error)⁶ with simple text book equations⁷, independent of the technology details. This means that in case of doubt, we can calculate an additional data point with these equations.

If inductance is so robust, why would we ever have any discrepancy in inductance? The reason is that **measurement** of small inductance values is a challenge, especially when it comes to sub-nH inductors where the feed structure has much higher inductance than the inductor itself. Small errors in the measurement will result in big errors in the extracted inductance.

Another reason for difference in inductance are closed metal loops that are near the inductor. We often see that in measurement. The image below shows an inductor with feedlines, as used for measurement, and the "pure" inductor as provided in the design kit.




Due to the closed ground loop that is so close around the inductor, the measured inductance is ~6% smaller than the stand alone inductor. It can be very misleading if the design kit shows the stand alone inductor, but provides measured data that still include the coupling to the ground frame.




⁶ The only exception are very conductive layers that are directly below the inductor, and change the fields.

⁷ **Simple Accurate Expressions for Planar Spiral Inductances**, S. Mohan, M. Hershenson, S. Boyd, and T. Lee, *IEEE Journal of Solid-State Circuits*, 34(10):1419-1424, October 1999.
http://www.stanford.edu/~boyd/papers/inductance_expressions.html

Note that a similar "closed metal loop around" also affects the Sonnet simulation, because the Sonnet simulation box can be considered a closed loop (or shorted secondary). We need to have a sufficiently large simulation box to minimize this ground loop effect. The example below shows the same inductor with different size of the Sonnet analysis box.



In the effective series inductance, you can see that the smallest box (1.7 coil diameters) creates a significant (4%) change in inductance. For more accurate results, it is recommended to use a box size that is ~3 inductor diameters.⁸



Here, we have assumed that your actual hardware design (the circuit that you want to build) does not include a ground loop. But if it does include such ground loop, you can include that ground in the simulation, or use the Sonnet box to represent the ground loop.

⁸ For more details, see application note "Cell size and box size in Sonnet RFIC inductor analysis" <http://muehlhaus.com/support/sonnet-application-notes>

3. The Q factor

Obviously, calculating the Q factor is of great interest. This tells us about the loss in the inductor.

If there are differences between the Sonnet Q factor result and measured or design kit data, we should check a few things:

- Do both curves use the same definition/test conditions of Q factor (differential vs. single ended) ?
- Do results agree at low frequency?
- Can the difference be explained by series resistance?

3.1 Differential Q factor vs. single ended Q factor

Yes, we can have some misunderstanding here.

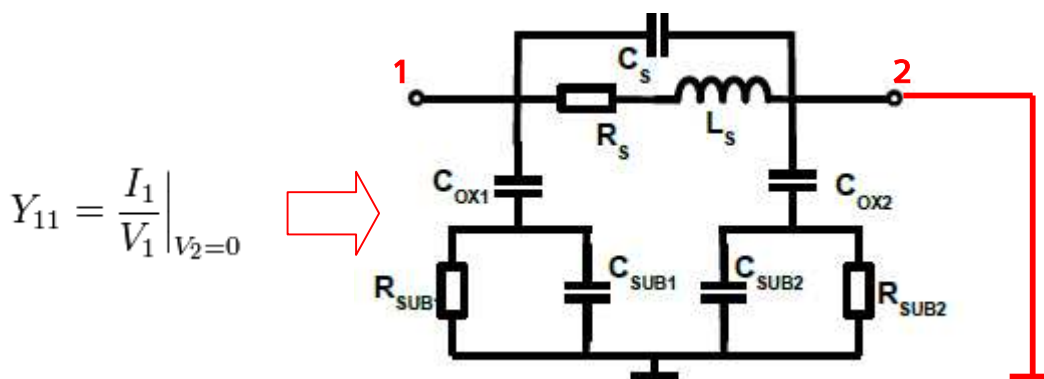
The Q factor value depends on how you **connect/measure** the inductor!

When you look at the Sonnet equations, you will find the "Q factor" equation for 2-port simulations and the "Q factor3" equation⁹ for 3-port inductors where port 3 is the center tap.

It is important to understand what the underlying definition/equation is:

$$Q = - \text{Imag}\{Y_{11}\} / \text{Real}\{Y_{11}\}$$

The Q factor is calculated based on the Y-parameters at port 1, which means that we look into port 1 with all other ports shorted¹⁰.




That Q factor from Y11 is a standard definition, but it is important to understand that **we get different results between two-port and one-port configurations**. The two port configuration is popular for inductor simulation because it includes all possible modes, no matter if we have differential or common mode signals at the ports, or a combination of both. **However, the above Q factor equation applied to this two port model will not give the differential Q factor between the ports!**

⁹ This equation was introduced in Sonnet 13.52


¹⁰ All other ports shorted because the Y11 definition enforces all other ports to have 0V.

One way to get the Q factor for the differential mode is to use a differential port configuration. In the Sonnet model, we now have ports 1/-1 instead of 1/2. This will also result in different inductance at higher frequencies, because the effective shunt elements are different now.



For a two port setup, the single ended Q factor is lower than the differential Q factor.


When you compare your Q factor values to data that others have calculated, make sure that the inductor is connected/measured the same way.



Now the question is: can we also calculate the differential Q factor from a two port simulation? The answer is yes. We need to measure the Q factor between the ports. The calculation and results are shown below.

3.2 Calculate differential Q factor from 2-port simulations

To calculate **differential** Q factor from 2-port simulations, we need a different equation which measures differentially between port 1 and 2.



$$Z_{diff} = Z_{11} + Z_{22} - Z_{12} - Z_{21}$$

$$L_{diff} = \text{Imag} \{ Z_{diff} \} / 2\pi f$$

$$Q_{diff} = \text{Imag} \{ Z_{diff} \} / \text{Real} \{ Z_{diff} \}$$

The Sonnet equation file with Q_{diff} and L_{diff} is available for download.¹¹

These equations can be used for 2-port data, but also for 3-port data where port 3 is the center tap.

Finally, when we discuss "differential Q factor" here, don't expect that everyone will use that term in the same way. Always double check how that Q factor was measured/simulated. The best option is to have Sonnet results and measured/design kit data available as a file, with both data tested with the same single ended/differential port configuration, and use the exact same setup for plotting the Q factor.

¹¹ http://muehlhaus.com/wp-content/uploads/2012/01/differential_inductor.zip

3.3 Special equations for center tapped inductors

The equations for differential L and differential Q are based on Z-parameters, which imply that any other ports are left open. This means that the equations can also be used for center tapped inductors, where we have an additional port (port 3) at the center tap.

The problem with the inductance2, resistance2 and Q factor equations is that they are based on Y parameters, and thus defined for V=0 at all other ports. In other words, a center tap at port 3 would be grounded when using these equations. This is not what we want.


The Sonnet equations inductance3, resistance3 and Q factor3 are supplied for single ended evaluation of center tapped inductors, with an open circuit at the center tap¹² at port 3. The results are the same as using the "normal" equations (inductance2, resistance2, Q factor) and leaving the center tap unconnected.

4. Technology related differences

To simulate in Sonnet, we use some description of the technology, which contains metal conductors, vias, dielectrics and the substrate.

For a discussion of how to build an EM stackup for Sonnet, please refer to the application note "Defining a stackup for Sonnet RFIC analysis"¹³.

For the moment, let's assume that you have designed the stackup with care, used the proper thick metal models with NS=2 sheets¹⁴ and taken into account that Sonnet thick metal penetrates into the dielectric above, so that the dielectric height must be defined as the sum of metal thickness and spacing to the metal above,



Hardware stackup


Thick metal stackup for simulation

¹² Yes, the center tap must be at port 3, this is hard coded into the equations.

¹³ <http://muehlhaus.com/support/sonnet-application-notes>

¹⁴ Yes, two sheets. As long as the metal height is not (much) bigger than the gap between the lines, this is perfectly fine. If you use more than 2 sheets, than strange things can happen to loss, depending on the frequency. Often, the two sheet model is more accurate for loss than a 3 sheet model. If you insist to use more than two sheets, you should always do some convergence analysis for loss with more sheets.


In the previous chapter, we discussed the low frequency response that allows to verify the conductor properties. Now, it's time to look at the full frequency range, where the shunt path to the substrate limits the Q factor.



For an inductor over an ideal lossless substrate, the Q factor would be determined by the series resistance and the self resonance frequency. The self resonance frequency is determined by the inductance and the parallel capacitance. The parallel capacitance consists of inter-winding capacitance + parallel capacitance through substrate.

4.1 Substrate conductivity

If we compare that ideal case over lossless substrate¹⁵ (blue curve) with the real case over conducting substrate¹⁶ (pink curve), we see that the curves agree at low frequency. In this frequency range, the Q factor is only limited by the conductor losses.



At very high frequency, the self resonance will cause $Q=0$ even with a lossless substrate.

In between, the substrate loss will cause a reduction in Q factor, because some energy is coupled into the conducting substrate.

¹⁵ In this example $\epsilon_r=11.9$ and $\sigma=0$ S/m

¹⁶ In this example $\epsilon_r=11.9$ and $\sigma=5$ S/m

The substrate conductivity for the bulk silicon is defined the technology data sheet. This is the value that we use for the Sonnet substrate.


If you import your stackup from the Assura pro file, please note that you have to set the substrate properties manually, because they are not included in the Assura files! The **default substrate value** that you will see during import is only that - a hard coded default which the Sonnet developers have created. You **always** have to adjust the substrate properties to the actual values for your technology.

4.2 EPI and pwell block

On the surface of the substrate, the conductivity will change, depending on the doping. This is also affected by some masks, such as pwell blocking layer used in the inductor area.

Over the years, we have learned "the hard way" that some technologies have a thin layer with higher conductivity at the surface of the substrate, underneath the inductor, which is not documented by the foundry. This will change the Q factor.


The plot below shows the reduction in Q factor if we add a 2µm EPI layer with 50 S/m over the 5 S/m bulk substrate. The effect depends on frequency and technology.




Next, we will look closer at the substrate path and try to identify such effects.

4.3 Looking closer at the substrate shunt path

For troubleshooting substrate effects, it is useful to separate the shunt path into a capacitive and a resistive part. This can be done if we have analyzed the inductor as a two-port. We can then calculate some effective R and C for the shunt path.



The effective values can be displayed using the equations RFIC_Substrate_Cap and RFIC_Substrate_Res. These equations are not included in the standard Sonnet equations, but available for download.¹⁷



Unfortunately, the values are not directly displayed as Cox, Csub and Rsub, but we can derive the values from the calculated Cshunt and Rshunt:

- At low frequency, we have Rshunt = Rsub. At higher frequencies, we see the effective resistance of Rsub in parallel to Csub, so that the effective Rshunt decreases.
- At very low frequency, we have Cshunt = Cox.
At higher frequency, we see Cox in series with Csub || Rsub.
At very high frequency, Csub dominates with $1/(j\omega Csub) \ll Rsub$, so that we see the series connection of Cox and Csub.


Let's look at some examples to see how that evaluation works.

¹⁷ http://muehlhaus.com/wp-content/uploads/2011/11/rfic_inductor.zip


To use this equation file in the Sonnet plotting tool emGraph, go to Equation > Manage Equations

Below is a comparison of Cshunt for the "normal" case (blue) and for the ideal, lossless substrate case (pink).

For the shunt capacitance, you can see that we have a flat curve with the ideal lossless substrate. For the real cases with conducting substrate, we can see how the shunt capacitance changes with frequency. From the curve, we could also calculate the values of $C_{ox}=76\text{fF}$ and $C_{sub} = 1/(1/22\text{fF} - 1/76\text{fF})$.




For the shunt resistance plotted below, you can see that we have ~ 0 Ohm for the lossless substrate, with some numerical precision issues at very low frequency. This is because the shunt impedance is determined by the capacitance and the real part is very small, thus sensitive to numerical precision error,




For the real, lossy substrate, you can see that we have $R_{sub} = 880 \text{ Ohm}$. At higher frequencies, R_{shunt} is $\text{Real}\{R_{sub}||C_{sub}\}$ which allows to double check the C_{sub} value calculated from the C_{shunt} curve. Both values must be consistent.

By calculating C_{sub} from the C_{shunt} curve and from the R_{shunt} curve, we also have an interesting way to test the data consistency. The author does this by fitting the curve with a lumped model. If the data is consistent, Q factor curve and C_{shunt} and R_{shunt} curve all match between data and lumped model.¹⁸

Now let's check the effect of different substrate conductivity on Q factor and on the extracted R_{shunt} .




Despite the numerical noise at low frequency, you can see how the substrate conductivity changes the shunt resistance.



Looking at this parameter allows to separate substrate effects from other effects. For example, the height of the inductor above the substrate will change the Q factor and the shunt capacitance, because C_{ox} changes, but R_{shunt} is unchanged.

¹⁸ If we can not match measured data with such a lumped model, this might indicate some mistake in de-embedding, where wrong shunt capacitance was removed. In this case, there is some offset between the C_{shunt} values extracted from the measurement and the C_{shunt} values that we need to match the Q factor curve. We have seen some of these cases, which could be resolved by repeating the measurement.



In the next example, we see how a change in oxide thickness = change in distance to the substrate¹⁹ will show in Q factor, Cshunt and Rshunt.



We see that Qmax has changed and f(Qmax) has changed. The curve is identical at low frequency, so we know that it's not conductor losses. The difference must be related to losses in the shunt path.


In this test case, we already **know** that we changed the distance to the substrate, and not the substrate itself. If this was real data, looking at the Q factor plot, we would not know what the reason for the difference is.

But if we now look at the Cshunt plot now, we immediately see that Cox must have changed. The Cox value has decreased from ~76fF to ~48fF.



¹⁹ Original model: oxides 2µm + 2µm, modified model: oxides 2µm + 5µm

Nice! From that plot, we know that the difference is in the capacitance from the inductor to the substrate, and not in the substrate itself. Let's double check by plotting the shunt resistance.



Indeed, the shunt resistance is almost identical. We can conclude that the difference in Q factor is caused by a change in the oxides above the substrate. From the Q factor, we would not have seen that, but by looking at the shunt path we found out.


4.4 Influence of the Sonnet box size on shunt path results

Wasn't that exciting? We can exactly check the substrate properties by extracting the shunt path.

But wait a second, is it really that exact? The inductor couples to the conducting substrate, and the substrate is then connected to the Sonnet box, which is our global ground. What about the Sonnet box size? Shouldn't the substrate path change when we change the box size?


Yes, true. The length of the path through the substrate depends on the size of the Sonnet simulation box. And in measurement, it depends on the path length to the substrate contacts that might connect the probe ground pads with the substrate. What does that do to Q factor, Cshunt and Rshunt?

Let's do an experiment, and compare our baseline inductor (left) with the results for a much larger box size (right).




The small box size is 384 μ m, the large box size is 768 μ m. The substrate thickness is 300 μ m. The substrate touches the box side walls and the bottom of the analysis box.


For the Q factor, fortunately, we don't see much difference between the two box sizes.



For the shunt capacitance, we see identical Cox and some small difference at higher frequencies caused by differences in Csub, due to the ground path length.




For the shunt resistance, the bigger box shows a somewhat higher Rsub.



It is important to understand that these differences are caused by the different length of the substrate path to ground, with identical substrate properties.

So don't immediately question your foundry's substrate data if measured and simulated Rshunt have some offset. This might be some difference in ground path length. For a difference in substrate properties, or an additional conductive EPI layer, we would see not only an offset but also a change in shape over frequency.



5. Summary

What is the essence of all this?

Double check your Sonnet models and your stackup, but also take measured or foundry design kit data with a grain of salt. Measured models might be significantly off, due to process variation, contact resistance, de-embedding issues or other reasons that are not immediately visible. By using the systematic approach presented here, you will be able to locate the origin of differences, in the series or shunt path, in the oxide or substrate. Once you have tracked down some difference, you might now get back to the foundry to discuss possible reasons and error sources. It might be some EPI layer that they did not mention in the documentation, some process changes since the data was created, or some other reason. But there is always a reason. Let's go and find it, so that we all enjoy more accurate and reliable data in the future.

For suggestions and questions related to this document, please contact:
Volker Mühlhaus, volker@muehlhaus.com

6. Change log

31. January 2012:

Added a description about differential calculation from 2/3 port inductor simulations, including a download link to the Sonnet equation file for Ldiff/Qdiff.