

High Speed Interconnect Design at Tektronix

John Dunn

Tektronix
Beaverton, Oregon

Ansoft HF Users Workshop

January 25, 2002



In this talk, I will give some examples of how we are using Ansoft's simulation tools at Tektronix. At Tektronix, we make test and measurement equipment, especially oscilloscopes and logic analyzers. Our products span the range from entry level oscilloscopes, to some of the highest performance test equipment in the world. For example, we just announced the world's fastest real time, storage scope: 6 GHz, and a sampling rate of 20 GHz when used in 2 channel mode. For more details on the range of products we offer, please go to our web site at: <http://www.tektronix.com>.

I am going to briefly discuss the challenges of interconnect design at Tektronix, and then give four examples of how we have used Ansoft tools to assist in some of our current design challenges. Two of these examples will be issues in probe design; one will be in BGA package design; and one will be in chip design.

The Challenges of High Speed Mixed Signal Interconnect Design


- ▶ Difficult in the Time Domain or the Frequency Domain
 - High Speed (20 ps risetimes)
 - Broadband (40 GHz)
- ▶ Complexity
 - Chips have 10K transistors at speeds in GHz range
- ▶ Analog
 - Probe Design
 - Amplifier Design

The proper design of high speed test equipment presents some interesting and difficult challenges for the interconnect engineer. First, we are always pushing the envelope in performance. After all, our test equipment has to be faster than the circuits it is testing, which are themselves state-of-the-art. We have to design systems that are fast, with rise times as short as 20 ps. At the same time they are very broadband, with frequencies up to 40 GHz being considered. Many of the traditional microwave solutions are not available to us, as they are inherently narrow band solutions.

We design both digital and mixed signal ASICs for our equipment. The job of our interconnect group is to support this ASIC development. Our mixed signal ASICs are high speed and complex. We use a silicon germanium process to attain the necessary speeds, and require 10K transistors to achieve the necessary functionality. Critical sections of our chips are necessarily analog. For example, our amplifiers in the front end of our equipment or active probes has to be analog.

Probes

- ▶ Compensation Structure for the Input
 - Must be trimmable
- ▶ Matching to the Cable for the Output
- ▶ Single ended – High Speed Signals through Power/Ground
- ▶ Differential Probes – output – common mode rejection issues.
- ▶ Solutions must be DC to highest frequency of operation
- ▶ Typically – use hybrid package technology

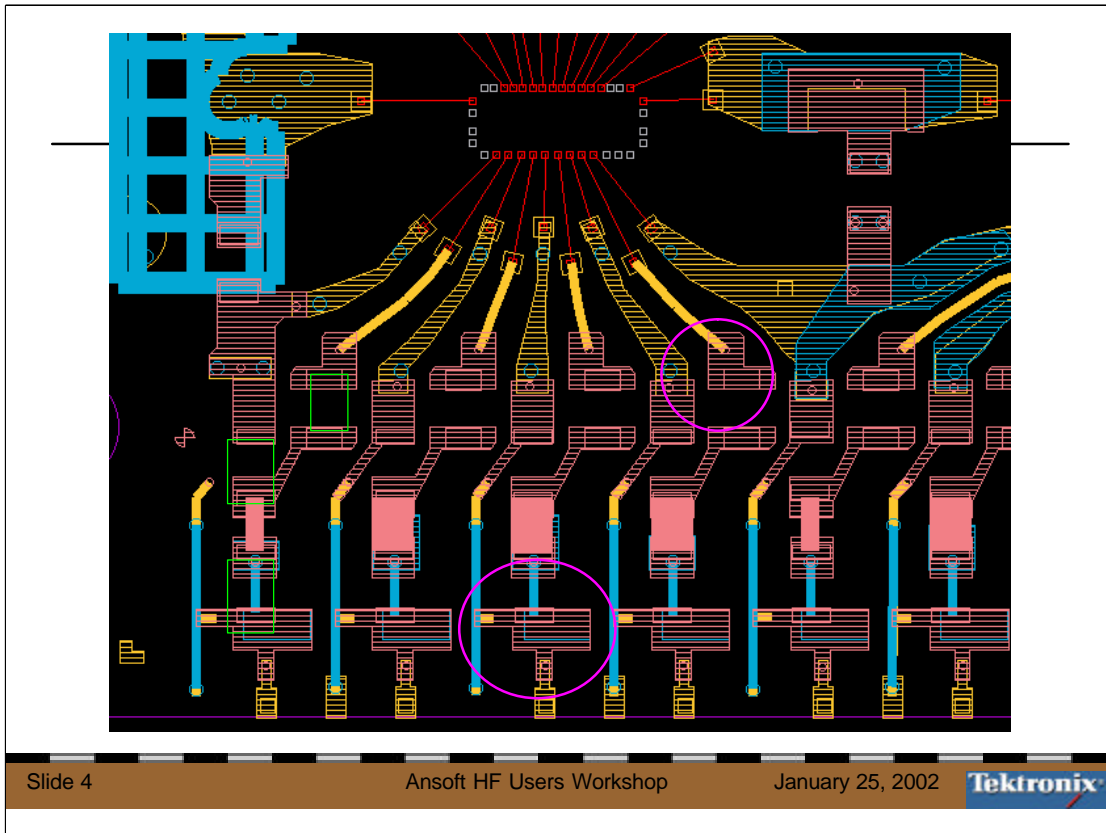


Slide 3

Ansoft HF Users Workshop

January 25, 2002

The first example I want to show occurred in the course of designing a high speed, high impedance, probe. Probe design has a number of interconnect challenges. Probes are broad band. In this case the probe has to work from DC to multi GHz. Let's examine some of the critical issues. This particular probe is single ended. Therefore, we are always worried about high speed signals on the power and ground. In the case of a differential probe, we have to worry about common mode rejection issues. We have to worry about compensating for cable losses. Of course, we have to worry about matching the output to the cable. And we have to understand the input structure well. An analog probe always has an input capacitance (from the input buffer of the amplifier). For broadband operation, we use a filter that keeps the same impedance over the entire band of operation. It is critical that we model the capacitance of the input accurately.




Here we see the input section of one of our digital probes. There are four channels going to the amplifier chip. I have circled the trimmable capacitor that is designed into these probes. The technician laser trims this capacitor to get the proper probe response. The problem is you have to make sure you have a reasonable estimate of the capacitance, or you won't be able to trim the probe correctly.

How do we get the capacitance of this fairly complicated structure? I used Ansoft's Si2D to model the structures, and verified a few sections of it using HFSS. I didn't want to use a 3D simulator for the whole problem, because of time constraints. I will show that the agreement between the two tools was good for the trimmable capacitor. It did not do as well in other sections of the circuit - for example the second circled section on the slide.

Use Si2D to Get Capacitance

- ▶ Transmission Line Structures
- ▶ Overlay Capacitors



Slide 5

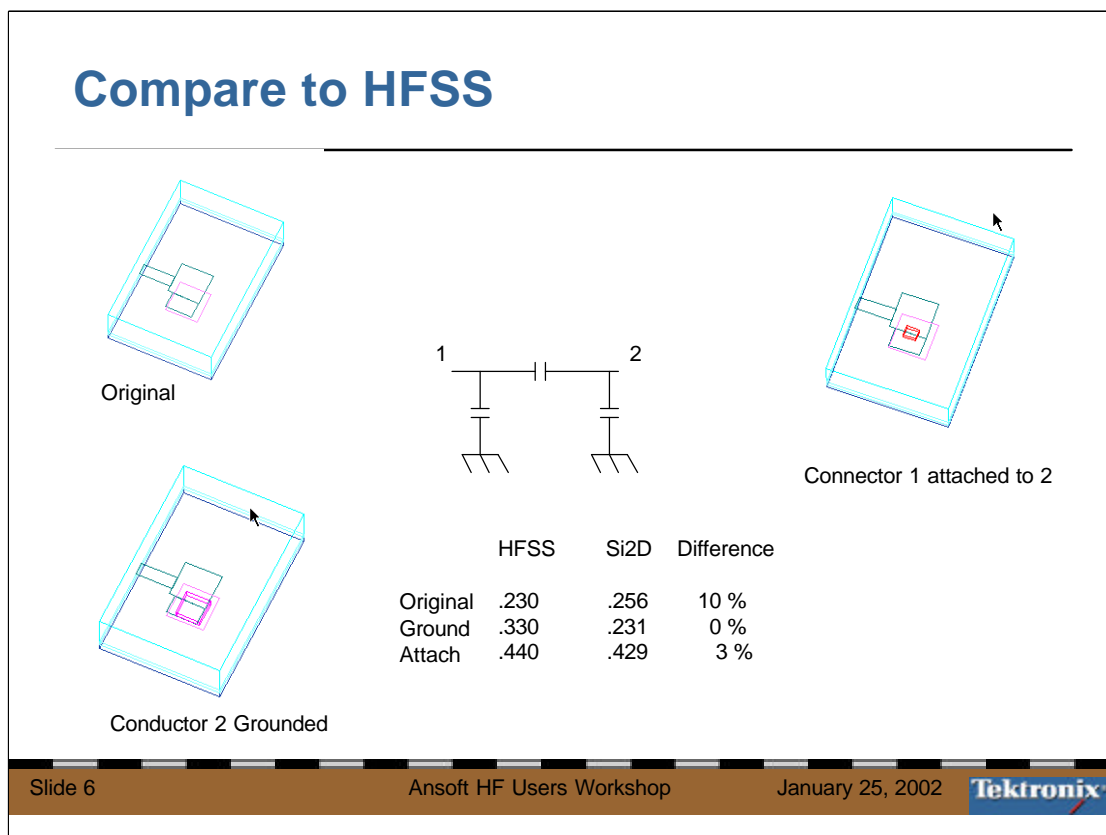
Ansoft HF Users Workshop

January 25, 2002

The trimmable capacitor is partitioned into 2D cross sections. The layer stackup is shown, with two signal layers, and a ground layer. The trimmable capacitor is on the top layer, called layer one. The second layer of the capacitor is on layer two. There are therefore three capacitances to worry about: C11, C12, C22. C11 is the self capacitance of the top to ground; C22 is the self capacitance of layer 2 to ground; and C12 is the capacitance between the two plates. The blue outline shows the top plate; the red rectangle is the bottom plate. I start by slicing horizontally. The cross section is drawn in Si2D. The length is 21 mils. I get the capacitance values shown. I then make a second cross section, as shown by the teal blue line; it is 6 mils long. Notice in this section, there is only the top plate (and ground). So, it only makes sense to talk about C11. This process has neglected the fringing capacitance at the top and the bottom. To include this, I vertically slice using Si2D. But, I subtract the parallel plate contributions from the total capacitances, as otherwise I am double counting.

Please note that these numbers are only intended to be representative.

Compare to HFSS



Slide 6


Ansoft HF Users Workshop

January 25, 2002

How well does this method work? Here is shown a comparison with HFSS. I get a difference of at most 10 %, with most the values being better than that. I regard this as an acceptable result for this type of application. Notice that in HFSS I had to run three cases, as there are three capacitances to extract. The original situation is just a one port attached to the top plate. I ran HFSS through 10 GHz, deembedded the feed line, and found the capacitance on a Smith chart, in the standard way. The second case has the lower conductor grounded. I am therefore looking at C11 in parallel with C12. The final case has conductors 1 and 2 connected. I am therefore looking at C11 in parallel with C22.

Please note that this is a different case than the previous viewgraph; don't expect the numbers between the viewgraphs to agree.

Another Example – Not as Successful



Hole is cut out under line

Ground plane is layer 5


Si2D: .332 pf

HFSS: .460 pf

Here is a case that doesn't work out so well. This case was shown by the second circle on slide 5. In this situation we have five layers, with a hole cut out of the ground plane on layer 3 we had before. This was tried to reduce the capacitance to ground. The hole is shown as the pink outline; the top layer is the gold. I tried Si2D, and then HFSS. I got a large difference in capacitance. I feel the HFSS is more accurate. This structure is inherently 3D, with nearby structures. It isn't realistic to think a 2D simulator will work.

HFSS Tip

Question: What value of impedance should you make the gap source impedance?



Answer: You want to be about equal to the magnitude of the load impedance where you are measuring.

Because: $\text{Change in } z_{load} = \frac{(z_{load} + 1)^2}{2 z_{load}}$ $\text{Change in } S_{11}$.
Where z_{load} is the normalized load impedance.

If $z_{load} \gg 1$: Slightest error in S₁₁ gives big change in z_{load} . Very unstable.
(Far right part of the Smith Chart.)

If $z_{load} \ll 1$: S₁₁ is not sensitive to z_{load} . S₁₁ doesn't change. So, no accuracy
in the mesh/solution. (Far left part of the Smith chart.)

Slide 8

Ansoft HF Users Workshop

January 25, 2002

Tektronix

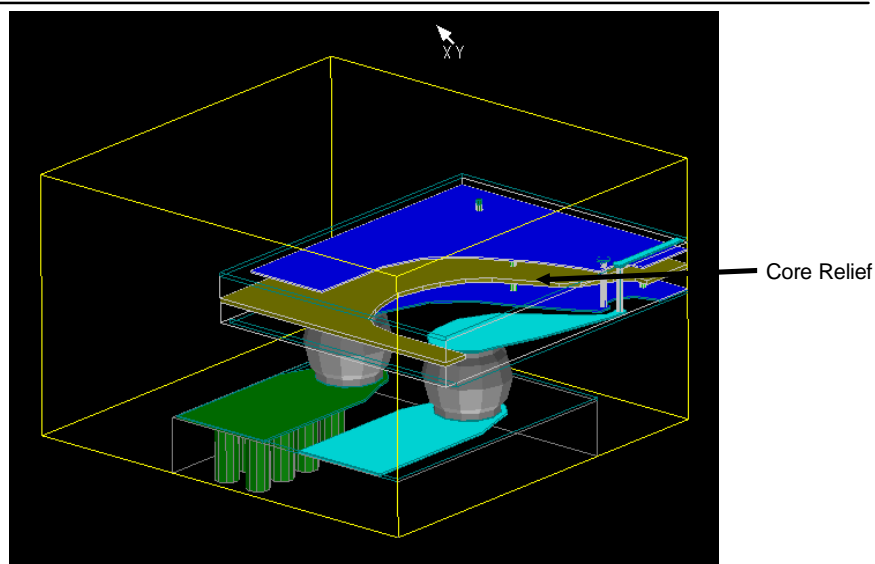
Here is a tip for users of HFSS that I discovered while running these simulations. The question is what impedance to make the gap source. If you're like me, the temptation is to use 50 Ohms. Doesn't everybody? The problem is that this can lead to big errors. Examine the value of S₁₁ from a load Z_{load} . Simple sensitivity analysis shows that the change in Z_{load} is related to the change in S₁₁ by the formula. The interesting thing is that if your normalized Z_{load} is too high (way on the right part of the Smith chart), you will be in a very unstable part of the solution space. This means the slightest change in S₁₁ will give a very large change in the value of Z_{load} . If Z_{load} is too small, you have the opposite problem. S₁₁ won't depend on the value of Z_{load} . This is bad because HFSS is refining its mesh according to S₁₁. So, it will run a few iterations, and say it's converged, leading to large errors. The correct thing to do is make Z_{source} about the same as Z_{load} .

Ball Grid Array (BGA) Package Design

- ▶ Balls introduce discontinuity.
 - Can be capacitive or inductive depending on design.
- ▶ We use both single ended and differential lines.
 - We try to use differential lines for signal integrity reasons.
 - Sometimes tightly coupled; sometimes not.
 - Input of scope or probe is single ended.
- ▶ Can modify:
 - Line widths
 - Metal layer cutaway.
 - Pad sizes on board and BGA.
 - Remove (depopulate) balls.
- ▶ Good BGA design involves a lot of tweaking. There is no silver bullet.

The second example is the simulation of a ball grid array (BGA) package. We use BGA packaging extensively for our ASICs, and it is critical that we model the signal path properly. The balls introduce a discontinuity that is sometimes capacitive and sometimes inductive. The balls are normally spaced on an uniform pitch; for the example I will show the pitch is 50 mils. However, one is allowed to modify line widths, pad sizes, and cut away planes to improve the signal integrity. Balls can also be removed. There is no magic answer to improving the design of a BGA package. Subtle changes can lead to vastly different performance characteristics.

An Example of a BGA Simulation in HFSS




Slide 10

Ansoft HF Users Workshop

January 25, 2002

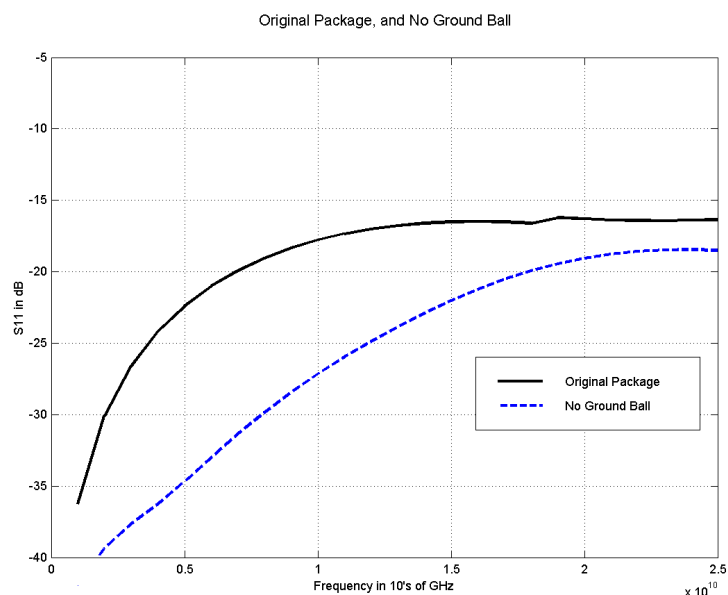
The signal path is simulated in HFSS. The light blue trace and ball is the signal trace. The green trace and ball is the ground trace. This package uses a ground-signal-signal-ground configuration. The other signal and trace are not shown here; the outer box is metal. By symmetry, the other signal line is on the other side of the vertical wall, which is acting as a virtual ground. The dark blue planes are ground planes internal to the BGA. The gold plane is a core, copper ground plane, required for this process. Notice the ground planes have been cut away. We did this in order to better match the ball discontinuity.

A Top View of the Same Package



A top view of the same package.


How to Improve the Performance



Slide 12


Ansoft HF Users Workshop

January 25, 2002



The black line shows the reflection coefficient magnitude in dB through 25 GHz. The blue line shows the package without the ground line and ball present. Notice that the response is 10 dB better. This is somewhat surprising, because the lines going up to the package have dimensions such that the outer ground lines are not seen much by the signal lines. Taking the ground lines away only changes the impedance by an Ohm or so, out of 50 Ohms. So, why would the ground ball removal matter if the ground lines aren't influencing the signal anyway?

Why it Helps to Remove the Ground Ball

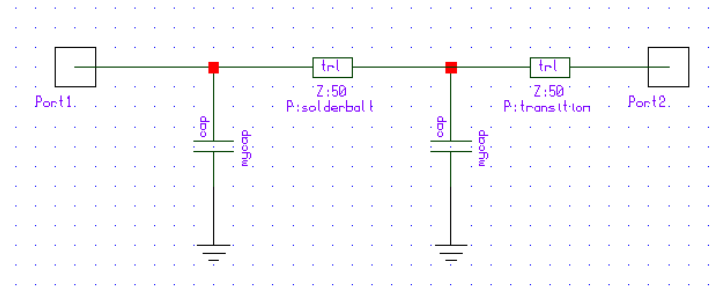


No Ground Ball

With Ground Ball

The reason the ground ball degrades the performance of the package is shown above. The pictures show the magnitude of the magnetic field midway between the solder balls. Notice the high field reaching the ground ball in the left figure. The right figure does not have the ground ball present; a much lower field is seen where the ground ball used to be. The ground ball is forcing the current to come around farther on the signal ball, which results in more coupling.

Circuit Model of Final Design

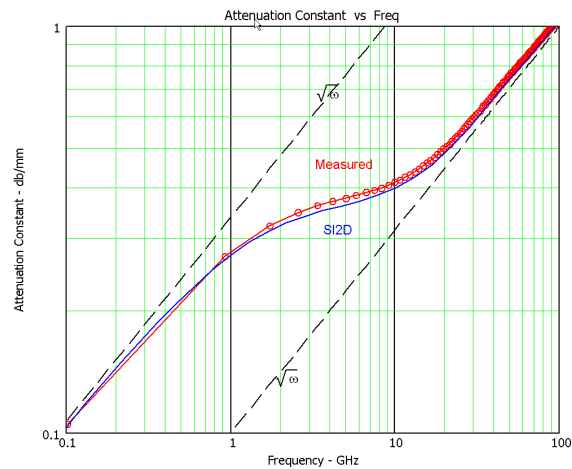


Capacitors = 10 ff
Line length = 46 mils

The finished model for the interconnect structure when no ground ball is present. This model works adequately through 20 GHz, and can be inserted in Spice. There are 2 capacitors of 10ff. They are connected together by a transmission line section of 50 Ohms. The length of the line is 46 mils; this is approximately the height plus the width of the ball. A second transmission line section models the signal path in the BGA package, including the via.

Interconnect Modeling on Chip

- ▶ IBM's SiGe Process
- ▶ Critical Lines – Line Loss is important
- ▶ Use Si2D.
- ▶ The attenuation curve is complicated.
- ▶ The W element model has problems.



Slide 15


Ansoft HF Users Workshop

January 25, 2002

The third example is at the chip level. It is important that we can model loss accurately in our high speed lines. Here are measured data for the loss (α) of a transmission line compared to the Si2D simulation. The agreement is considered to be excellent for our modeling needs. Notice that the high frequency response asymptotes to the expected square root of frequency behavior, where the skin depth loss dominates. The low frequency loss is also the expected square root behavior. Of course, this is not due to skin loss, but rather the domination of line loss over inductance at low frequencies. The transition region extends over a large frequency range, up to 10 GHz. Note that the W element model often used in Spice will not work well. The W element models loss as a constant term, and a frequency dependence like a square root. It cannot model the transition region correctly. The W element model works much better in modeling lines on circuit boards where the transition region extends over a relatively small frequency range.

Bond Wire Modeling

Back end of an analog probe.



Slide 16


Ansoft HF Users Workshop

January 25, 2002

Tektronix

The final example is a model of the bond wires from the output of the amplifier chip of an analog probe to the hybrid package in which it is mounted. Modeling of the bond wires is important in order to insure proper performance. The output is single ended with two ground wire coming out along side the signal wire. An effort has been made to keep the ground return as close to the high speed signal as possible.

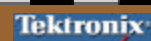
Gap Ports Used on the Chip



Slide 17

Ansoft HF Users Workshop

January 25, 2002



The chip is modeled as a block of silicon, with three pads on top: the signal pad, and two ground pads. The three bond wires attach to the three pads. There are two gap ports between the three pads. A wave port is used for the line going to the edge of the package from the chip.

Models

- ▶ **HFSS Generates S Parameters**
 - Can't put into internal Spice easily.
 - Doesn't tell you how to fix a problem.
- ▶ **Develop a Physical Based Model**
 - Takes time to do this.
 - Can tell you how to fix a problem.

It was not enough to simulate the structure with HFSS. We also wanted to develop a model that could be put into Spice. This was needed so that we had an understanding of the effects of the bond wires.

Model Development

- ▶ **Three Bond Wires**
 - Three series Inductances
 - Possible Mutual Inductance between Middle and Outer Wires
 - Pad Capacitance? End of Microstrip Line Capacitance?
- ▶ **Physical Based Modeling**
 - Start Simple – only add elements when needed
 - Each element has some physical significance
 - Use optimizer sparingly

It was decided that the first model would consist of three inductors, one for each of the bond wires. There could be mutual inductance between the bond wires. There also could be pad capacitance and capacitance at the end of the microstrip. We wanted a model that had some physical relevance; i.e., each element in the model could be related to a physical effect in the structure. We also did not want to start by running an optimizer; rather we wanted to slowly build the model up to have a better chance of getting physically meaningful values.

Simulations

- ▶ Isolate Signal Bond Wire
 - Ground the wire by turning Silicon to a Perfect conductor.
 - Back of chip connects to outer conducting box.
 - Sets inductance of signal bond wire.
- ▶ Isolate Ground Bond Wires
 - Drive Gap ports without microstrip line and signal bond wire attached.
 - Bring edge of box over.
 - Sets inductance of ground bond wire.
- ▶ Combine signal and Ground Wires
 - Get mutual inductance
- ▶ Final Simulation with Finished Model

The model was built up in stages. First, we looked at the signal bond wire inductance in isolation. This was accomplished by making the chip metal, and thereby grounding the signal bond wire. The next step is to isolate the ground bond wires by modifying the geometry to use a gap port exciting the ground bond wire. The other side of the gap port is attached to the side of the box. This gets the inductance of the ground bond wires. Next, the mutual inductance is found. Finally, the entire model is assembled and compared to the simulation of the entire structure.

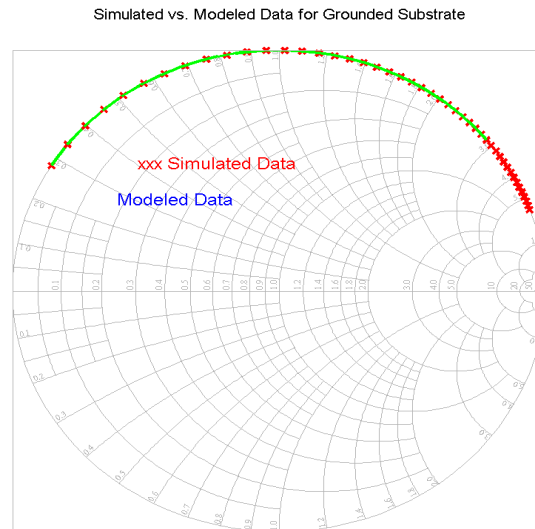
Isolate Signal Bond Wire

- ▶ Only Simulated Half of Geometry
 - Magnetic Symmetry wall put through middle of package
 - Impedance of Microstrip Line Multiplied by Two
 - Inductance Multiplied by Two
 - Capacitance cut in half
- ▶ Inductance = 1.15nH (= 2.30nH/2)
- ▶ Capacitance = 90ff (=45ff X 2)
 - Improves High End Performance (6 – 10 GHz)

The geometry was cut in half using a magnetic symmetry wall. This will double the impedance of the microstrip line, as it is cut in half. Therefore, the inductance we get for the signal bond wire will be twice as large, and the capacitance half as large as the original. After making these corrections, we end up with 1.15 nH inductance for the bond wire, with a fringing capacitance to ground of 50 ff.

Isolate Bond Wire Results


1-10 GHz
Capacitance
Not Included



Slide 22

Ansoft HF Users Workshop

January 25, 2002




Here are the simulated data from 1 to 10 GHz, versus the inductor model. The capacitor was not added in this model. When the capacitor is added, the curves lie on top of each other even for the high frequencies.

Isolate the Ground Bond Wire(s)

Signal Pad
Shorted
To Side of Box.

$L=1.75\text{nH}$



Slide 23

Ansoft HF Users Workshop

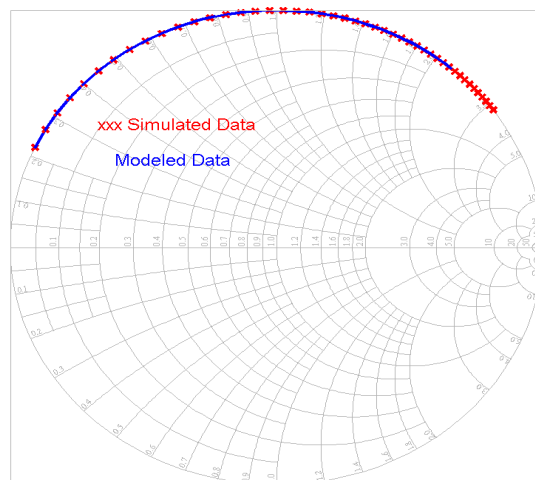
January 25, 2002

Tektronix

The ground bond wire inductance is found by simulating the wire in isolation with a gap port. The simulation gives a result of 1.75 nH.

Isolated Ground Wire

Simulated vs. Model data for Gap Port Connected to Ground Bond Wire



Slide 24

Ansoft HF Users Workshop

January 25, 2002


Tektronix

The simulated data and model are compared from 1 to 10 GHz for the ground bond wire.

Combine Signal and Bond Wires

Simulated vs. Modeled Data for Shorted Gap Port

Without
Mutual inductance



Slide 25

Ansoft HF Users Workshop

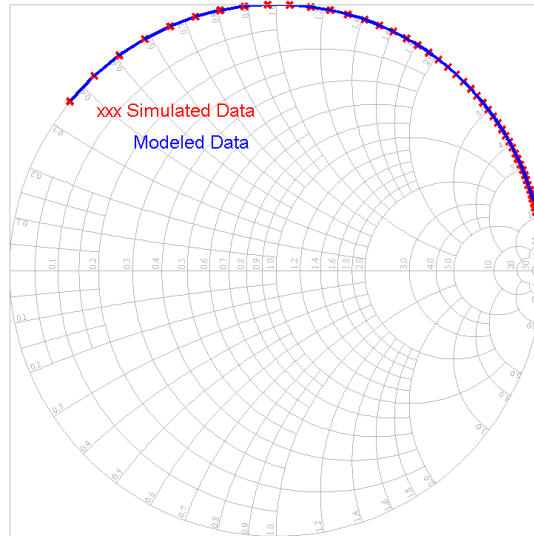
January 25, 2002

Tektronix

Now, the original structure is simulated, and compared to a model with the signal inductor, and the ground inductor. The mutual inductance between the two wires has not been included yet.

With Mutual Inductance

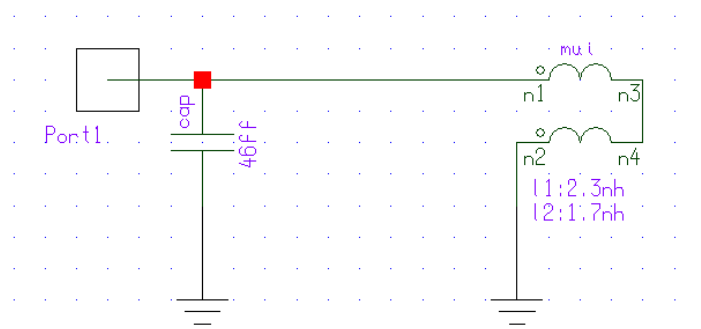
Simulated vs. Modeled Data for Mutual inductance



Slide 26



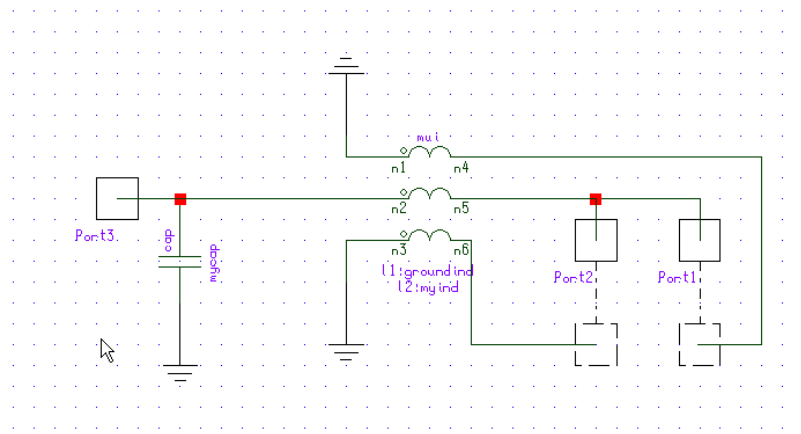
Model in Serenade



Mutual Inductance = 0.57nH

Here is the model in Serenade of the final model (with the magnetic symmetry wall).

Final Model

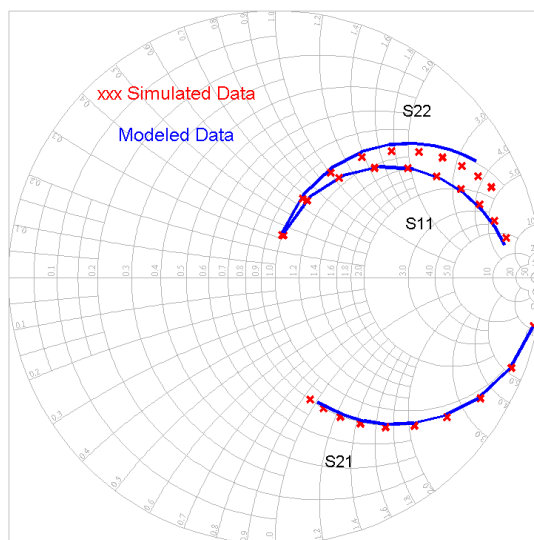


The final model with three ports is shown here.

Model vs. Simulated

Simulated vs. Modeled Data for Full 2 Port

All Ports
Renormalized
To 50 Ohms.



Slide 29

Arbitrarily Users Workshop

January 29, 2002

tektronix

Comparison of modeled versus simulated data for the original, three port circuit. The agreement is considered to be reasonable for this type of problem.

Where is it all Going

- ▶ Speeds – Moore's Law is holding
- ▶ Analog Design will be needed at high end
- ▶ Complexity – 100K transistors

- ▶ Problems:
 - Chip Level
 - ▶ 100K transistors – How do you handle parasitics/interconnect?
 - ▶ Must design with layout in mind.
 - Package Level
 - ▶ BGA's will have troubles if pushed above 20 – 40 GHz.
 - ▶ To many inputs for microwave style packaging.
 - Board Level
 - ▶ Control of parameters?
 - ▶ Performance of connectors, relays, ...

Speculations on the future. The only thing we know for sure is that the performance will be higher, which means things will be faster and more complex. At the chip level, we will be seeing order 100K transistors for mixed signal circuits. It will be more and more important to design with layout in mind from the beginning. At the package level, BGA technology will be forced to get smaller and smaller, until it will become impractical. Above 20 - 30 GHz, BGAs may not offer sufficiently good performance. This presents problems, as no other technology accommodates so many inputs. At the board and system level, it will be more difficult to find the necessary hardware that meets the required specifications.