



Cascaded vs End-to-End Multi-Pin Interconnect Simulation Models

Robert Branson, Samtec
Robert.Branson@samtec.com

Steve Krooswyk, Samtec
Steve.Krooswyk@samtec.com

Scott McMorrow, Samtec
Scott.McMorrow@samtec.com

Gustavo Blando, Samtec
Gustavo.Blando@samtec.com

Abstract

When performing bulk simulations on a multi-pin connector for design, optimization, or correlation, it is important to do everything possible to ensure the simulations are expedient. To combat increased simulation times, it is common practice to divide the board and multi-pin connector models into separate simulations that are then cascaded. However, this splice must be done on the BGA/Press-Fit/SMT region between the connector and PCB, where the TEM assumptions necessary for cascading do not necessarily hold true. This practice often introduces impedance and crosstalk inaccuracies between simulation and measurements. Comparatively, end-to-end multi-pin connector simulations introduce none of these issues, but lose the benefits associated with cascaded multi-pin connector simulations.

This paper demonstrates the differences between cascaded and end-to-end interconnect models through correlated simulation/measurement models. Field plots are used to show the inaccuracy created by the artificial simulation boundary where the model was cut to perform the cascade. Then we study how to devise mitigation strategies to maintain the decrease to simulation time generated by running cascaded multi-pin connector models while avoiding simulation inaccuracy that is generated through the same process.

Authors Biography

Robert Branson has been a Signal Integrity Engineer at Samtec for the past 3 years. He got his undergraduate in Electrical Engineering from the University of South Carolina in 2020 and is set to complete his masters in EE in 2022. He has worked in connector/PCB design, development, and correlation for numerous Samtec products, and has experience in using 3D modelling software for the analysis of connectors.

Steve Krooswyk is a Sr SI design engineer at Samtec. His 18 years of experience includes contributions to the design, simulation, and correlation of interconnect and I/O, as well as an influential role in PCIE standards. Prior, Steve was the PCI Express technical lead for Intel Data Center Group, co-authored the book High Speed Digital Design: Design of High-Speed Interconnects and Signaling, and received an MS in EE from U. of South Carolina.

Scott McMorrow currently serves as a Strategic Technologist for Samtec, Inc. As a consultant for years too numerous to mention, Scott has helped many companies develop high performance products, while training signal integrity engineers. Today he works for "the man," where he continues being a problem solver, a change agent and "betting his job" every day.

Gustavo Blando is a Senior Principle Engineer and leading Principal SI/PI Architect at Samtec Inc. In addition to his leadership roles, he's charged with the development of new SI/PI methodologies, high speed characterization, tools, and modeling in general. Gustavo has twenty plus years of experience in Signal Integrity and high-speed circuits.

Introduction

The basis for a large amount of Signal Integrity work, and indeed the work of many fields within Electrical Engineering, is Linear Network Theory. Linear Network Theory describes the ability to represent a complex system, i.e. an electrical circuit, in terms of a system of linear components. It simplifies a complex problem down to something that can be solved numerically.

For the purposes of Signal Integrity work, Linear Network Theory is used to cascade a series of s-parameters, representing the individual performances of sections of a larger system, together to form that complete system. This is used at all levels of Signal Integrity. At the most basic, transmission line models can be cascaded together to get the overall performance of a line. At the most complex, most spec lines (such as GenZ, PCIe, ect...) rely on cascading to estimate channel performance of large, multi-PCB systems. These systems compromise the silicon generating a signal, through PCBs and connectors, and the destination specific to that signal. All of this is predicated on the idea that Linear Network Theory can be used to estimate overall performance for these channels.

Linear Network Theory is a large and encompassing idea, which necessitates a few key assumptions before it can be used. Each component being cascaded must accurately reflect the performance of its related section for the complete channel to be accurate. Of note to this discussion is the assumption of isolation and independence between the components. This means that every component exists separately from each other and can be considered a separate piece. In terms of electromagnetics, this necessitates that a TEM (Transverse Electro-Magnetic) boundary exists between all separate components. Each component must have a clear and delineated start/end that can be reviewed independently.

Problems start to arise when the TEM assumption cannot be so easily made. This issue increases at higher frequencies, where it is difficult to delineate where one component ends and another begins in the complex and far-reaching electromagnetic fields. Complications can arise in considering two systems as independent when electromagnetic fields interfere between them.

The simple and obvious solution to this problem is to combine any systems where this is occurring into one system with combined electromagnetic fields. This is fine on paper, but much of Signal Integrity work is done using 3D electromagnetic field solvers, and this solution becomes a problem when that is taken into account. The solve time for a complex model increases exponentially with the number of tetrahedra within the mesh for that model. An internal and informal study performed within the author's present employer found a roughly 3x increase in solve time for a 2x increase in tetrahedra. Taking this into account, it becomes clear that it is in the best interest of resource usage to keep model geometry as small as possible.

Given that there is a direct conflict between solve time and geometry size, the distinction on where the line can be drawn between models that can be linearly cascaded is extremely important. Increased solve time not only costs more computing resources, but also the engineering hours it necessitates to deal with a longer wait between submitting a model and receiving of the resultant data. Wherever possible, it is optimal to cut models down so that one

large model becomes two separate models that are then cascaded together. This has the additional benefit of making it easier to swap components out without having to re-run the large, combined model. However, this must be balanced with the inaccuracies this can introduce when done incorrectly.

This paper is going to explore a particular situation where this conflict arises while modelling multi-pin connectors. Specifically, the example of open pin-field, multi-pin connectors which sit on PCBs will be considered. This can be edge-mount, surface-mount, Edgecard, or any situation in which a connector is making direct contact with a PCB or PCB-like structure. This paper will explore the potential problems/benefits that can arise when designating the PCB and connector as two separate components to be cascaded together. This paper will also explore some solutions to the problems that can arise.

What is Cascading in Signal Integrity?

Relative to this paper, it is important to define what is meant by cascading. Cascading refers to the methodology of combining s-parameters together to get a combined s-parameter which contains the information of both s-parameters in one. The process is shown below, and in Figure 1, where you can see an example combined network.

$$\begin{aligned}
 S_{11} &= N_{11} + \frac{(N_{21}N_{12}N\#_{11})}{(1 - N_{22}N\#_{11})} \\
 S_{12} &= \frac{(N_{12}N\#_{12})}{(1 - N_{22}N\#_{11})} \\
 S_{21} &= \frac{(N_{21}N\#_{21})}{(1 - N_{22}N\#_{11})} \\
 S_{22} &= N\#_{22} + \frac{(N\#_{21}N\#_{12}N\#_{22})}{(1 - N_{22}N\#_{11})}
 \end{aligned}$$

In this example, N is the first s-parameter and N# is the second. S represents the combined s-parameter.




Figure 1: Example of a cascaded network of s-parameters

Cascaded vs End-to-End Interconnect Models

Since this paper will be an exploration of the differences between cascaded and end-to-end interconnect models, it is also worthwhile to have a rigid definition of what the two model types mean within this paper. A complete channel simulation, generalized below in Figure 2, represents the whole network which contains the connector. This generally consists of the silicon generating the signal, through a PCB, through the connector, and onwards until the destination. In the case of 3D electromagnetic modelling, this would never be done as a single simulation due to the massive complexity and time required to build/run such a model.




Figure 2: Complete interconnect channel model

Common practice in modelling multi-pin connectors which land on a PCB structure falls into two categories: cascaded modelling and end-to-end modelling. For the rest of this paper, the cascaded model will refer to the model on the right in Figure 3. This model has a connector run as a separate 3D electromagnetic simulation to the PCB break-out region. The cut is done in the BGA ball, press-fit, or SMT region separating the two.

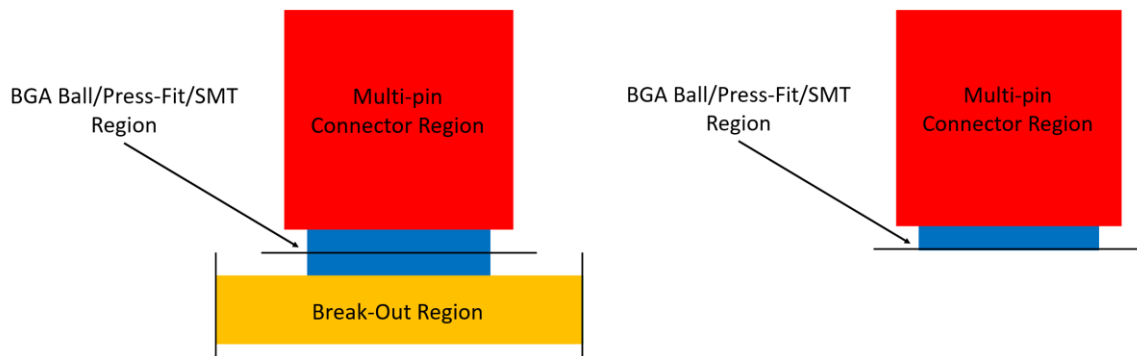


Figure 3: End-to-end interconnect model (left); cascaded interconnect model (right)

Alternatively, the end-to-end model refers to the model on the left in Figure 3. This model simulates the connector and PCB break-out region concurrently in a single design. In both cases, the cut in the PCB is made at the end of the BOR, where the trace exits into a TEM region.

Example of Differences in Model Types

To validate the claim that there is a difference in performance that can be seen between end-to-end and cascaded interconnect models, an example will be used from a real situation that occurred while the authors were working at their present employer. An open pin-field, mezzanine, multi-pin connector was designed and built using the cascaded modelling approach, due to its benefits in the design process. It allowed quick and smooth iteration on the design. When the time came to test the connector, it was assembled, placed onto a PCB, and measured.

Figure 4 shows a simulation of the differences between end-to-end and cascaded models of the connector mentioned above. Note that while the data shown is simulated, the end-to-end simulation model correlates to the actual measurement of this connector when it was built. The plots contrast the cascaded model and the end-to-end model with each other, demonstrating how large of a difference it made for this particular connector.

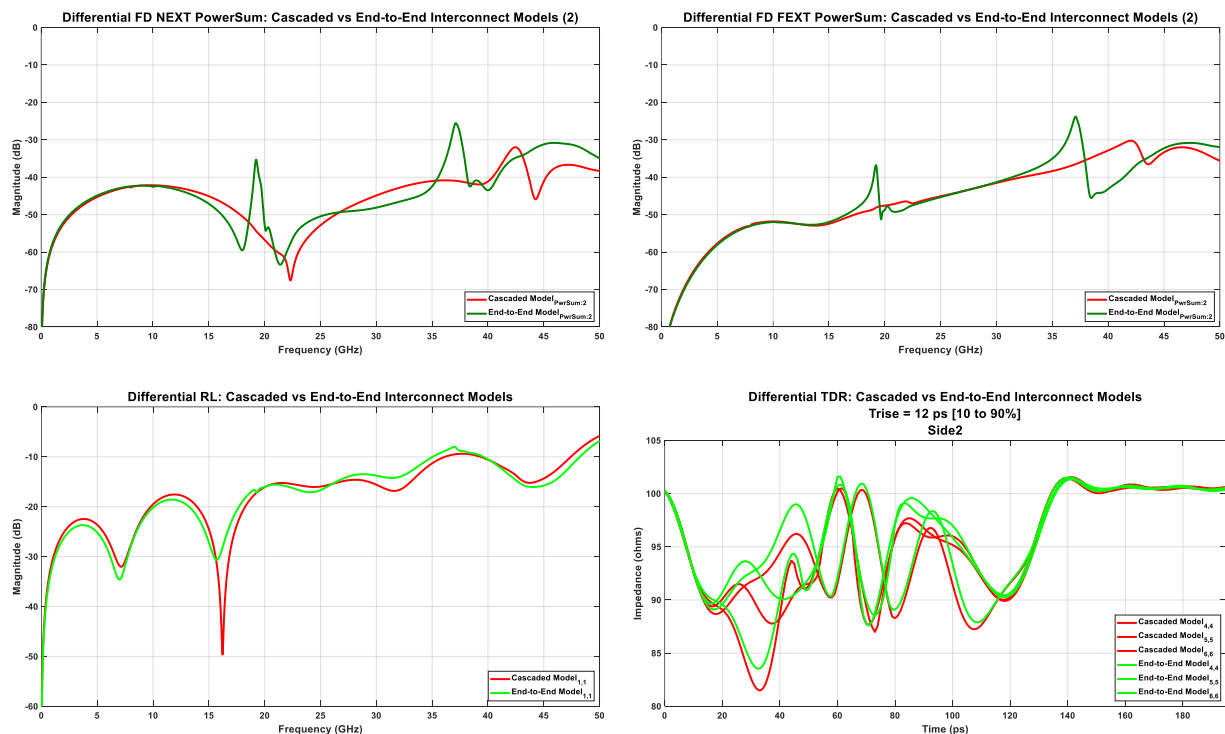


Figure 4: Cascaded (red) vs end-to-end (green) connector models for an open pin-field, multi-pin, mezzanine connector

The plots in Figure 4 show the cascaded model in red and the end-to-end model in green. The first and most prominent observation to note is that there was a resonance that occurred in the end-to-end model which did not occur in the cascaded model at 19.22GHz. This can be seen in both the near-end and far-end crosstalk PowerSums. A small ripple of the impact can also be seen in the return loss plot. This difference was verified in actual measurements to be exactly the same, emphasizing this as a real issue which can occur for connectors.

This resonance comes from a length-based resonance that is caught only by the end-to-end model and not the cascaded model. This length-based resonance will be discussed more later, when the electromagnetic field plots are shown, but for now it is sufficient to say that it is a geometry dependent resonance that can only be seen when the board is included in the simulation model. It was not seen until the measurements were taken because all models had previously been run independent of the board.

The second observation is the difference in impedance between the two models. This is a secondary concern to the resonance issue, but still accounts for around a 3ohm reduction in impedance at the boundary between the models. For this particular model, a cut was made at the BGA balls to separate the PCB and connector (with each model containing half of the solder ball geometry). An incorrect impedance is observed for two reasons: (1) PEC boundaries of the artificially introduced ports create a coupling path that does not exist otherwise and (2) because this is not a perfectly TEM region, like linear network theory requires for complete accuracy, non-TEM fields are lost as an error.

Both of these issues demonstrate variance in results between the two different modelling methodologies. Depending upon which approach is used, there is information that can be lost in a 3D electromagnetic field solver. As such, possible ramifications should be considered when deciding which modelling approach to use.

In the case being documented, once the problem was realized, the connector model was fixed and the design was changed to mitigate the realized resonance. However, the damage was already done. This mistake cost both money and engineering hours to rebuild the parts and boards to fix the resonance issue. These losses could have been avoided completely by constantly using an end-to-end model during the entire design process, however this is not a perfect solution either due to the additional time taken when using end-to-end models. The alternative methodologies section later in this paper will look at some other approaches that can be used to address this issue.

Building a Simplified Model

The next step in identifying the disconnect between the cascaded and the end-to-end interconnect models is to try to understand what caused the differing results to occur. Later in the paper, when the results are discussed, this understanding will lead to an attempt to prevent that disconnect from happening.

In order to begin to understand what caused the difference, it helps to start with a simplified model instead of a finished connector model. Modern high-speed connectors are a careful balance between electrical and mechanical requirements. This produces, as a necessary byproduct, a connector which is highly complex. These modern connectors have pins with numerous bends, ridges, and varying widths. By simplifying this connector back down to its most basic electrical essence, it becomes much easier to run tests to try and understand what is fundamentally causing an issue to occur.

Figure 5 shows the simplified model geometry which will be used for testing in the rest of this paper. It is designed to mimic the original connector's performance shown in Figure 4 as closely as possible. This means that it is also an open pin-field, multi-pin mezzanine connector of the same height as the original. The pins, however, are constant rectangles. Additionally, its BGA balls are perfect cylinders. The PCB region was kept the same as the original connect to minimize differences other than the connect itself.




Figure 5: Simplified connector model

As with the original cascaded and end-to-end models, the cut to make a cascaded model was done halfway through the BGA balls. To verify a similar performance, the NEXT performance is shown in Figure 6. As evident, the behavior is very similar to what was demonstrated in Figure 4. There is a slight difference in resonance frequency and magnitude in the simplified model due to the different electrical path through the simplified connector, but its behavior is similar enough to the original connector to use as a test case.




Figure 6: Simplified cascaded vs end-to-end interconnect models NEXT

Analyzing the Source of the Problem

This section will use the simplified model that was developed as a test vehicle to analyze the difference between the two modelling methodologies. This is done with the goal of understanding how those differences occurred. It will end with a discussion of what caused the resonance in this connector and how the cascaded model hid that resonance while the end-to-end model did not.

The first step that in analyzing the resonance in this connector was to cut the model at different locations and observe when the resonance occurs. For these simulations, the results themselves will not be shown, but instead it will be stated whether or not a resonance was observed in the results, thus indicating whether the model correctly captured the resonant behavior contained within the end-to-end model.

These slices are shown in Figures 7, 8, and 9. The first slice, shown in Figure 7, shows a normal cascaded model. The model is cut halfway through the BGA balls. As observed previously in Figure 6, this model does not demonstrate the resonance that is being sought.

Figure 8 shows the slice being done at the bottom of the L1 ground plane. This is cutting through the signal via to determine whether or not the resonance can be seen if ONLY the L1 ground plane is included in the simulation, and not the rest of the PCB. While this will not produce a good impedance result (the region being cut is clearly not TEM) the goal at this moment is only to locate what is necessary to recreate the resonance. As in the previous case, Figure 8 did not produce a resonance.

Figure 9 shows the slice being done at the bottom of the L2 ground plane. Once again, it is cutting through the signal via and placing a port at that sliced off location within the PCB. This time, the resonance was observed in the results.




Figure 7: Simplified model cut at the BGA ball




Figure 8: Simplified model cut at the bottom of L1




Figure 9: Simplified model cut at the bottom of L2

There are some observations that can be taken away from these results, but that discussion will come at in the next section.

To better understand why some configurations do not resonate, the next item to investigate before looking at the results as a whole are electromagnetic field plots. For these plots, the simplified model was expanded out to two rows to get row-to-row crosstalk data. These plots only show the cascaded and end-to-end models, and not the layer-by-layer cuts that were done in Figures 7-9.

The most basic plots are shown in Figure 10 and 11, where the H-fields are shown for the cascaded and end-to-end models. These H-fields are measured in air at the halfway point in the BGA balls, where the ports are placed for the cascaded model. These H-fields were measured at the resonant frequency to capture what happens differently when the connector resonates.

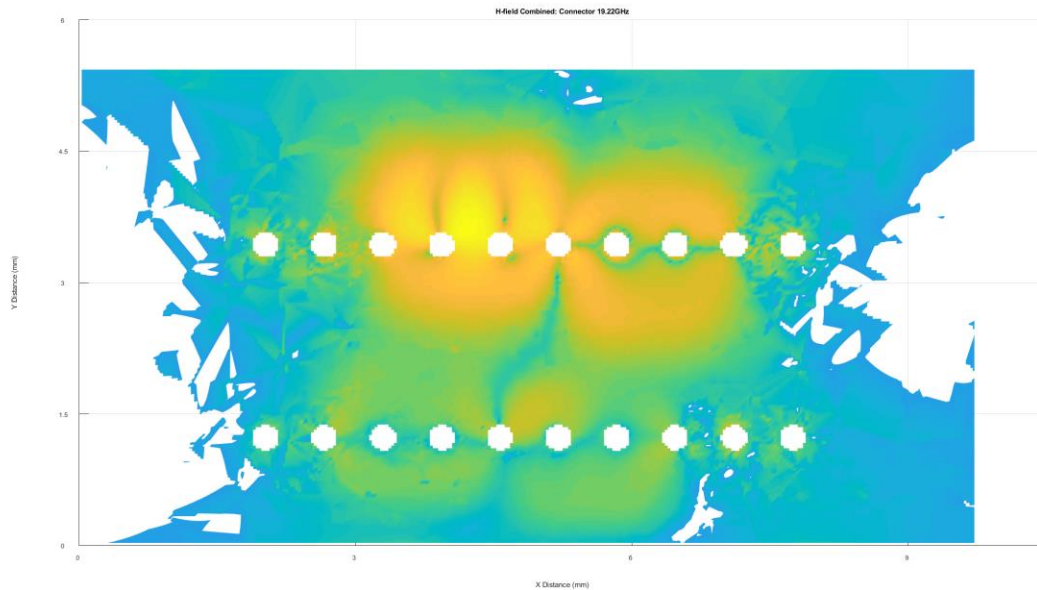


Figure 10: End-to-end model 19.22GHz H-fields at BGA ball




Figure 11: Cascaded model 19.22GHz H-fields at BGA ball

The main observation that should be made here is how destructive the ports are in the cascaded model as compared to the end-to-end model. The edge of ports represent a solid line of PEC surrounding the fields in that region. By necessity, these ports therefore interrupt the smooth flow of electromagnetic field lines in the region they are placed, creating the straight lines that can be seen in Figure 11.

Figure 10, the end-to-end model, shows the opposite. Fields are correctly allowed to radiate along the BGA region as would be expected. It can also be very clearly seen in the field lines that there is crosstalk occurring on the neighboring pairs (both in-row and cross-row) which is completely gone in the cascaded model. It should therefore come as no surprise that the cascaded model sees less crosstalk at the resonance than the end-to-end model.

The final data that will be looked at are the E-fields in the region. Like the H-fields, they were taken at the halfway point down the BGA balls and at the resonant frequency.




Figure 12: End-to-end model 19.22GHz E-fields at BGA ball




Figure 13: Cascaded model 19.22GHz E-fields at BGA ball

Figure 12 shows the end-to-end model and Figure 13 shows the cascaded model. A similar pattern can be seen between the two, suggesting that the electrical behavior is similar between them. This is to be expected, as the same signal is passing down both lines with the same geometries. The difference, however, can be clearly seen in Figure 14.




Figure 14: Difference between Figures 12 and 13

Figure 14 subtracts the difference between the E-fields of the two previous graphs. While it was difficult to see in those figures, it now becomes clear that the E-fields are varying quite significantly all across the board. For most of the simulation space, this is around -50dB, which could largely be considered negligible. However, right at the excited differential pair and across to the neighbors it can send crosstalk to, the difference is much more significant. Here, the difference comes as high as 5dB, showing a much larger impact in E-fields.

These numbers don't directly translate to crosstalk in a numerical sense, but the E-fields and H-fields are what cause crosstalk to accumulate as the signal passes down the pins. When both are so warped and distorted in the BGA region by the presence or absence of ports, it clearly causes a major difference in the amount of crosstalk that is seen between the cascaded and end-to-end models.

Discussion of Results

With the results gathered from the previous section, some conclusions can be drawn about the resonance that occurred in the end-to-end model but not in the cascaded model.

As mentioned earlier in the paper, the resonance in question here is a length-based resonance. This can be calculated by determining the speed of electrical propagation in the dielectric, determining the length of the pins, and then finding the half-wave resonant frequency for that particular length/DK value. An example is shown in Figure 15.




Figure 15: Example geometry for a length-based resonance

With the type of resonance known, it should be possible to understand the resonance from the details that were already gathered. The resonance is based on the length between the two ground planes of the PCBs the connector is attached to (since it is a mezzanine connector). This makes it clear why the resonance did not occur when the connector was cascaded. The two ground planes act as a roughly 1ohm impedance for the ground return path. This is a result of the ground return path having a 50ohm impedance for the length of the pins. This massive discontinuity causes the reflections that result in a resonance at the frequency related to the length of the connector. Because there are two massive discontinuities at either end of the 50ohm pins, the energy bounces back and forth without major attenuation.

However, this is not the whole story. Figure 8 included the L1 ground plane in the simulation but did not see a resonance. This points to the answer being more than just the existence of the ground impedance discontinuity causing the resonance. A second factor is the nature of the energy being dispersed by the port within the simulation structure. By placing a port at the bottom of the L1 ground (inside the anti-pads), a single return path is created from the BGA balls to the antipad edge. In reality, the return current will follow many paths into the depth of the PCB. The PEC of the port allowed that energy to escape instead of building up and reflecting around L1 of the PCB.

This can be observed in Figure 9, whereby placing the port inside the antipad of L2 ground, the resonance was seen to occur. By including the vias between L1 and L2, current on L1 is now dispersed to multiple paths that better represent the ground impedance discontinuity of the end-to-end model.

Figure 10 and 11 showed how destructive the port was to the surrounding electromagnetic field structures. The port therefore needed to be outside of the resonant structure in order to not have such a destructively interfering effect. Figure 14 showed how much of a difference this can cause to crosstalk, especially in the areas around each differential pair.

All of this work for this particular problem can be put together to show the exact conditions under which the length-based resonance for an open pin-field, multi-pin, mezzanine connector

will be hidden by the simulation. From here, it would be possible to run simulations that include only a part of the PCB as a sanity-check to ensure that there is no hidden resonant behavior. However, that doesn't necessarily mean that this is the perfect solution either. Simulation time will still be increased because a part of the PCB has to be included in the simulation. Additionally, impedance results will be highly inaccurate due to the cut in the PCB being placed in a non-TEM region.

This process was shown to provide an example of how to understand a particular resonance in a connector and design a unique approach to that resonance which allows it to be simulated in as simple a manner without necessitating an end-to-end simulation that captures that resonance. There could be specific situations in which this is the best solution, but many will be unsatisfied with incorrect impedance results coming from the way that the solution had to be implemented. The conclusion will go into more details on what the final takeaways of this process should be, but there are other alternatives to this approach to consider first.

Alternate Methodologies

Until this point, this paper has presented cascaded and end-to-end simulations as the only two options. However, there are different strategies that can also be used. Given the pros and cons of each cascaded and end-to-end modelling that have been shown up until this point and with neither being a perfect solution, it is worthwhile to examine the alternatives.

The first alternative is known as a multi-mode simulation. This is a cascaded model in which only one port is used instead of a separate port for each differential pair. An example of what that geometry would look like is shown in Figure 16, with the green box at the top and bottom representing the wave ports that are placed. The red pins are the differential pairs and the yellow are the grounds, while the grey are the actual ground references used for the ports.




Figure 16: Multi-mode cascaded model simulation (purple boxes are 7-terminal wave ports)

Since only one port is being created on both the top and bottom, with the pins at the end as the ground reference for that port, the end result is that the simulation treats all pins as signal pins. Post-processing is necessary to take the resultant s-parameter file and convert it into the necessary format. The ground ports need to be grounded so that they do not add to the crosstalk. In the example above, the 3D electromagnetic field solver would generate a s14p, which would then be cut down to an s8p by removing the grounds.

A simulation was run using this method in Figure 17 to show a normal cascaded model vs a multi-mode cascaded model. As can be seen, the resonance is present in the multi-mode model.




Figure 17: Cascaded vs multi-mode simulation models FEXT

This methodology has the advantage of still being a cascaded model. That means that it runs faster (due to the smaller geometry), and the PCB used can be swapped without re-running the connector model. However, there are some disadvantages compared to the normal cascaded model. The primary disadvantage is ease of use. These models are often not as stable as a normal simulation model. It can be prone to errors and crashing, as well as not converging. This is due to the increased RAM necessary to support more complex field calculations. Secondly, since some ground pins are assigned as excitations in the port, they are not acting as ground pins. This not only means additional wasted time for the solver to consider excitations from the ground pins, but also that the ground pins don't work to dampen E-fields like they normally would. Finally, cascading of models with a different number of ports (as is created by this methodology) is a special case and may not be easily supported by some tools.

A second alternative solution is simply called intermittent solving. It means to run cascaded models for the majority of design work, but to intersperse end-to-end models during the process to sanity-check and make sure nothing is being missed. This method is beneficial for bulk design work, wherein many changes and iterations have to be made to a connector. It saves time and complexity, while making sure that nothing is missed before the final product rolls out. The end-to-end runs can be done every certain number of cascaded runs, or at major milestones. The disadvantage here is that a large amount of design work could go to waste if an issue that only appears in the end-to-end model is ignored in the cascaded model.

These alternative techniques, like the original cascaded and end-to-end modelling techniques, are also not perfect. There is no ubiquitous solution to this problem which fixes everything, but instead different methods that can be implemented after considerations of each project's needs and parameters.

Conclusions

In conclusion, this exploration of cascaded and end-to-end modelling, as well as their alternatives, has shown the advantages and pitfalls of each. No method could solve every problem, but the methods outlined provide a toolkit for approaching each unique situation and finding the best approach.

A large portion of this paper was devoted to analyzing the specific problem that occurred in a connector and trying to understand that problem. This was originally done with the intent of looking for a perfect solution that could allow for cascaded modelling with none of the downsides, but rarely is there a methodology that has no downsides. Instead, this now demonstrates how to go about the process of analyzing a cascaded vs end-to-end modelling disconnect. Should a situation arise wherein there is a difference between the two and one of the alternate approaches is not viable, it is possible to create a unique solution tailored to the specific problem at hand.

The solution shown is not meant to be the end-all, but instead the specific solution for one class of connectors with one type of resonance (length-based resonance on an open pin-field, multi-

pin, mezzanine connector). There are numerous types of resonance and numerous types of connectors. Each connector will have its own challenges and information obscured by the cascaded methodology, and so each must be approached differently. Hopefully, the steps laid out in this paper will aid in the decision process for resolving issues that arise with any particular connector.

In the end, if a final suggestion from the authors were to be given, it would be that the most basic solution that is useful in most situations would be intermittent solving. It deals with a lot of the pros and cons of each methodology by simply doing both. It will take more engineering hours upfront than sticking to one method, but that doesn't necessarily mean that it won't pay off in the long run. Of course, this won't be true in all cases. Learning to judge each situation appropriately and responding with the best fitting solution is the never-ending journey of growth for a Signal Integrity Engineer.

Method	Pro	Con	Recommendation
End-to-End Model	Most Accurate	Solve Time	Final Design Proof
Cascaded Model	Fastest	Incorrect Resonance, Impedance Near Ports	Common Iterative Designs
Separated Model within PCB (Figure 9)	Correct Resonance	Not Interchangeable, Impedance Near Ports	Specialized Iterative Designs
Multi-Mode Model	Best Accuracy for Separated Model	Complexity and Tool Limitations	Specialized Iterative Designs
Intermittent Solving	Balance of Accuracy and Speed	Doesn't Match Best of Alternative Methods	Bulk Iterative Design Work

Table: Methodology Recommendations Based on Explored Techniques

References

- [1] Balasubramanian, V., Smith, S. B., & Agili, S. S. (2007). Comparison of S-parameter concatenation to full-wave simulation for high-speed interconnect analysis. In International Engineering Consortium - DesignCon 2007 (pp. 1629-1643). (International Engineering Consortium - DesignCon 2007; Vol. 3).
- [2] Liao, M., Wei, L., Ling, F., & Wan, C. Optimizing BGA Ball Pattern for Signal Integrity. DesignCon 2021
- [3] Zhang, M., Liu, M., Xiao, K., Zhang, Z., Wang, W., & Wu, Y. Minimally Specification Compliant Behavioral Modeling Methodology & Its Applications in HSIO. DesignCon 2020.
- [4] Manukovsky, A., Bloch, R., & Sayfan-Altman, S. The Hidden Challenges in 112-Gb Channel Design & Modeling. DesignCon 2020.
- [5] Morrison, C., Chobanyan, E., Alavi, P., Zhang, Z., Wang, W., & Wu, Y. End-to-End System-Level Simulations with Retimers for PCIe Gen5 & CXL: A How-To Guide. DesignCon 2020.
- [6] Samtec gEEk spEEk webinar, 18 August 2022.
<https://www2.samtec.com/l/271452/2022-08-19/35krkg>