Interconnect Modelling Sensitivity Studies for High Signaling Data Rates

Adewale K. Oladeinde
Mixed Signal Development
NVIDIA Corporation
Austin, TX, USA
aoladeinde@nvidia.com

Jose A. Hejase
Mixed Signal Development
NVIDIA Corporation
Austin, TX, USA
jhejase@nvidia.com

***Abstract*—High speed interconnect modelling using 3D Electromagnetic solvers is a necessary step towards evaluating high speed channels in circuit simulations. Inaccurate modelling can have severe implications on circuit design direction, system design cost and/or future system field functionality. Modelling setup details, which may have seemed unimportant in the past, are becoming significant as data rates under consideration are increasing. This work presents three case studies on commonly utilized interconnects and how different modelling setups can yield different SI metrics for them. For each case study, recommendations are provided towards accurate modelling.**

Keywords—Electromagnetic Modelling, High Speed Interconnects, Signal Integrity, Scattering (S) Parameters.

I. Introduction

High performance computing server systems have exhibited considerable development over the last couple of decades. These advances have been happening on multiple fronts including: architectures, mechanical and thermal design, circuit design and capabilities, substrate and printed circuit board (PCB) technologies, connector & cable design, in addition to analysis methods and simulation tools. A big part of the motivation for these advances has been to support data rate and/or overall system bandwidth increase. Ongoing development for future high-speed links capable of supporting 224Gbp/s (PAM4) data rates will be seen commercially soon. Only about 10 years ago, the highest signaling data rate differential interfaces carried was less than 25Gbp/s (NRZ) for the big majority of signaling standards [1].

As data rates increased, circuit designs improved and became more capable. Channel non-idealities such as loss, reflections and crosstalk became more important to be dealt with and, usually, more challenging to mitigate. What was good enough before is not anymore and the signal integrity community finds itself in many instances searching for channel improvements wherever possible.

Furthermore, interconnect modelling tools have improved on multiple fronts over the years including solver accuracy, tool application comprehensiveness, time to solution and computational resource usage efficiency. Interconnect modelling is an extremely important step for signal integrity design, used to optimize high speed interconnect designs. This is necessary to arrive at “correct by design” systems helping minimize time to market and cost of testing iterations prior to going to market. Additionally, high speed channel interconnect models influence the circuit design as such models are used for circuit simulations to optimize the properties of equalization characteristics to handle a certain range of channel SI characteristics. As a result, it becomes essential to use these modelling tools appropriately so that the interconnect model characteristics are representative of reality and no model caused “artifacts” are created.

In this paper, three different modelling sensitivity studies are presented. These include: (1) modelling of interposer and substrate bump breakout for accurate transition representation in case the interposer and substrate sections are modelled separately and then cascaded, (2) modelling of vias with a variable number of faces to represent the via cylinder and how that influences the via SI metrics, and (3) modelling of trace cut-outs to avoid resonances generated by ground flood stubs resulting from the cut-out operation.

The rest of the paper discusses each of the case studies including: describing the interconnect being modelled, case study problem statement, simulation setup, simulation results and discussion towards providing best practice recommendations for future similar modelling.

II. Accurate representation of the transition between interposer and substrate when Cascading

Interposers have been used in chip packaging solutions for many reasons including enabling high wiring density between multiple dies with short wiring lengths [2]. High speed signals not communicating between dies on the same interposer must pass through the interposer into the substrate prior to leaving the substrate. The impact of the interposer on the high-speed signal integrity is not negligible and sometimes could result in problematic channel performance. As a result, it is necessary to model the interposer as part of the high-speed full channel modelling.

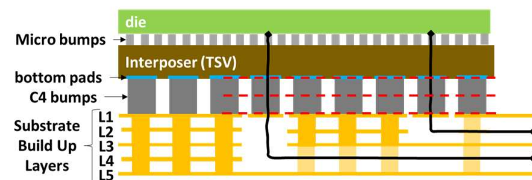

Fig. 1. High speed signal transition from die micro bumps through interposer to the substrate build up layer wiring.

As is well known, the full high-speed channel is often modelled in pieces before being cascaded to create the full channel representation. This is usually done to arrive at models faster or to abide by limited computational resources. For example, substrate and PCB wiring are modelled separately and usually themselves are divided into pieces separating breakout wiring from main route wiring. Modelling interposer with the substrate bump area breakout is possible but it comes at large time and computational resource costs, due to the challenge posed to the solver when having to deal with a wide range of feature sizes present in the interposer versus those in the substrate. As a result, it is usually preferred to model the interposer and substrate bump area breakout separately. Doing so, however, requires particular attention to the signal transition between the interposer and the substrate to ensure it is accurate after cascading the models.

979-8-3503-1798-5/23/$31.00 ©2023 IEEE

Figure 1 shows an illustration of a signal transition from the bottom of the die (die RDL) through the micro bumps to the interposer to the c4 solder bumps before getting into the substrate build up layers. Modelling the interposer and the substrate bump area breakout separately could encompass three different scenarios: (A) including the solder bump with the interposer model, (B) including the solder bump with the substrate model and (C) including part of the solder bump with the interposer model and the other part with the substrate model. The three different options are represented by the dashed red lines in Figure 1. For (C), part of the solder bump is assumed as 50% in this study. In this work, models for these three options were generated and compared against the results for the single merged model utilizing a high-speed differential interconnect designed to support signals with a 56GHz fundamental frequency in 86Ohm differential impedance channels. The modelling carried out in this case study included 8 Tx pairs and 8 Rx pairs, all near each other.

Ansys HFSS was used to generate the models for this case study using numerical solver settings defined to arrive at solutions with low error tolerances and a frequency sweep range between 0Hz and 100GHz with 25MHz steps. Similarly, this was done for the other case studies in this paper.

In the case of the merged model, wave ports were defined at the die RDL side and at the substrate trace breakout side. For the cascaded modelling option with the solder bump as part of the interposer model, ports at the bottom side of the interposer model were defined utilizing a PEC reference representative of the substrate top copper layer (L1) ground flood surface while the substrate model had ports defined directly on substrate L1 as well. The ports defined on L1 were wave ports with the same size as the existing voids around the signal c4 bump pads on L1. On the other hand, the models for the other two cascaded options defined ports at the transition between their interposer and substrate models utilizing an added PEC sheet with circular wave ports defined referenced to it.

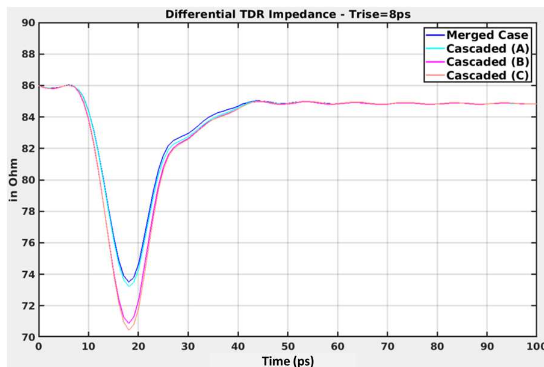

Fig. 2. Differential TDR impedance of the cascaded modelling options compared to the merged model.

Figures 2 shows the differential TDR impedance (Trise=8ps) for the worst-case TX pair for the merged model and the three cascaded model options. Cascaded option (A) is almost perfectly matched to the merged case while the other two options differ from the merged case. Similar trends are present in other SI metrics. This can be explained by the artificial impact the additional PEC puts into the modelling. Specifically, when it comes to impedance, the added PEC capacitance is causing more impedance mismatches. The results presented show that it is better to use an existing conductive plane as reference for port definition when modeling channels by cascade, as opposed to defining non-existent PEC sheets to act as a reference.

III. VIA NUMBER OF FACES

In a high-speed channel, PCB and substrate vias are signal routing structures prone to causing impedance mismatches and exciting propagating parallel plate modes resulting in crosstalk to neighboring signal traces or vias if care is not taken when designing them. To mitigate their potential adverse effects, especially at higher data rates with fast rise times, accurate modeling and simulation becomes crucial [3].

A via is typically composed of a barrel and more than one pad spanning some layers within a substrate or PCB stack up. The barrel and pad have circular cross-sections but when modelled, a discrete representation of the circular cross-section needs to be used to work with typical numerical solver meshing algorithms. The piece-wise representation of the circular cross-section becomes more critical as data rates increase. The choice of how fine the representation of the circle could lead to SI result differences. On the other hand, an extremely fine piecewise representation of the circle could lead to longer solution times while not providing significant additional accuracy.

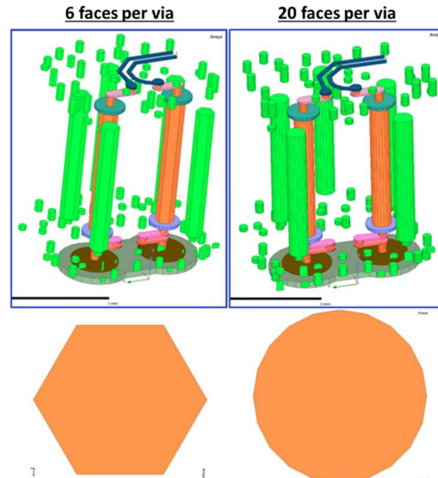

Fig. 3. High Speed differential via transitions with 6 faces and 20 faces; along with via cross-section.

In this paper, a substrate high-speed differential pair via transition from upper build layers to bottom BGA pad through a core was modelled and optimized for a high-speed channel carrying a signal having a 56GHz fundamental frequency. The substrate had a 5-2-5 stack-up with 33um thick GL102 ABF material in the build-up layers. The core PTH via had a 300um pad diameter and a 200um barrel diameter.

In the models, a differential wave port was defined under the BGA pads and a lumped port was defined at the end of short lead-in trace in the upper build up layers. The number of sides/faces for the vias within the models was varied between 6 and 20 faces. Different Anti-pad sizes were used to optimize the via transition as close as possible to a nominal differential pair impedance of 86Ω, assuming the 12 faces case. Figure 3 shows HFSS models for two cases: 6 and 20 faces. Please note that the substrate layer planes are hidden for visibility reasons.

Figure 4 shows the differential TDR impedance of the substrate via transition (looking from the BGA pad side)

using different number of sides/faces per via and a 10ps rise time. The via with 6 number of faces shows highest inductive impedance. This can be explained by the reduced via cross-section area. The impedance of the differential PTH via decreases as the number of faces increases resulting in a ~4.5Ω maximum difference in the TDR response between maximum and minimum number of faces considered. There is almost no change in the impedance between 16 and 20 via faces. This result shows a low number of faces could lead to an inaccurate representation. Furthermore, it shows that one does not need to go to 20 faces for an accurate result, 12 or 16 faces could be used instead resulting in a quite accurate representation at less computational cost.

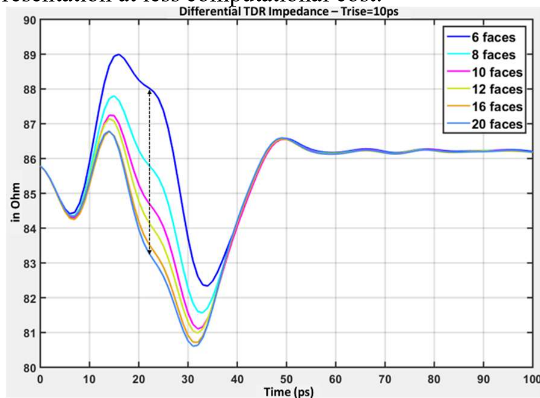

Fig. 4. Differential TDR impedance of the via transition while varying the number of via faces.

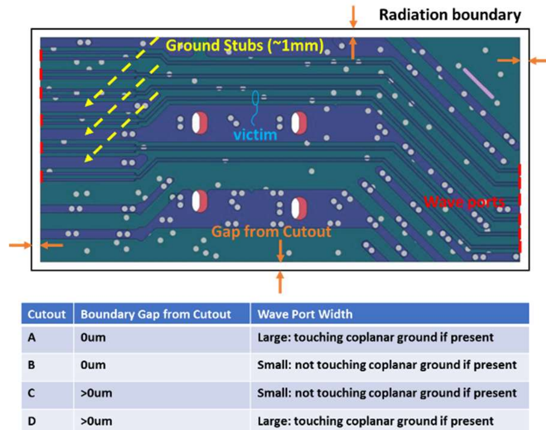

Cutout	Boundary Gap from Cutout	Wave Port Width
A	0um	Large: touching coplanar ground if present
B	0um	Small: not touching coplanar ground if present
C	>0um	Small: not touching coplanar ground if present
D	>0um	Large: touching coplanar ground if present

Fig. 5. Cut-out model top view and the different cases considered.

IV. Trace cut-out Modeling resonances

The continuous increase in the number of IOs for high-speed differential signaling has contributed to the increase in substrate and PCB sizes. To analyze the SI for a particular high-speed interface, it is sometimes necessary to model the channel in pieces to allow for faster simulation time with the use of lower memory. High-speed routing in PCB and substrate sometimes has ground floods in between signal traces to meet copper density rule and help reduce crosstalk. These ground floods are designed to have ground stitch vias and could result in signal integrity issues if they are missed from the design or when cutting out a section of the design to be considered for simulation. Excluding some stitching vias when doing a cutout can result in a ground "stub". A visual sanity check of a design cut-out model is important to ensure that unwanted artifacts do not impact the model simulation results.

To evaluate the sensitivity due to cut-out generated ground stubs and ways to mitigate their resonances, four models were created in HFSS for substrate differential wiring. Figure 5 shows the four model characteristics. Each model cutout case includes six differential pairs, four of which are neighboring a ground flood stub at least on one side. Differential wave ports were defined on both ends of the traces. Figure 6 shows the PSFEXT comparison of these four cases assuming the victim pair shown in Figure 5. The results clearly show that Cutout "C" exhibits a resonance due to the ground flood stub left over after the cutout operation. On the other hand, it is also clear that Cutouts A, B and D do not exhibit any resonances. This shows the importance of flushing the boundary against the cutout and/or extending the wave port to touch the neighboring ground to mitigate unreal ground stub resonances. That said it is also clear that Cutout A victim pair exhibits slightly less PSFEXT compared to the victims in the Cutouts B and D. This can be alluded to secondary coupling effects caused by not flushing the radiation boundary (in case of D) and not extending the wave port to touch the neighboring coplanar ground flood (in case of B).

It is recommended that any trace cutouts done from PCB or substrate use the flushed radiation boundary in addition to a wide enough wave port to touch the neighboring ground flood, if applicable.

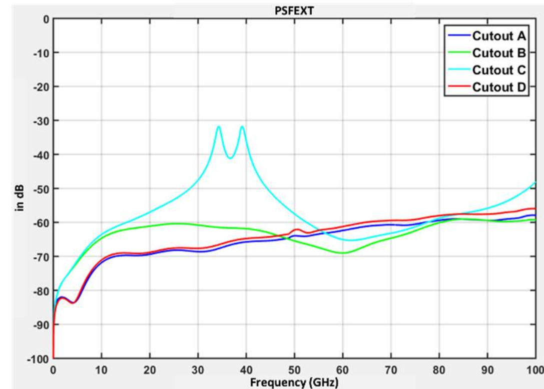

Fig. 6. PSFEXT signature on victim differential pair identified in Figure 5.

V. Conclusion

Three high speed interconnect modelling sensitivity case studies were carried out in this work. The sensitivity studies addressed common simulation setup issues and their impact on SI metrics. Results were discussed and remarks were provided for best accuracy. Future work includes evaluating more case studies and potentially carrying out measurement correlation for select case studies.

References

[1] S. Mirabbasi, L. C. Fujino and K. C. Smith, "Through the Looking Glass—The 2022 Edition: Trends in solid-state circuits from ISSCC," in *IEEE Solid-State Circuits Magazine*, vol. 14, no. 1, pp. 54-72, winter 2022.

[2] A. Usman *et al*., "Interposer Technologies for High-Performance Applications," in *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 7, no. 6, pp. 819-828, June 2017.

[3] B. Toby and K. Bois, "Rapid Design of High Performance 25+ GT/s Vias by Application of Decomposition and Image Impedance," DesignCon 2017 Proceeding