DesignCon 2021

A Case Study in the Development of 112 Gbps-PAM4 Silicon and Connector Test Platform

Jean-Remy Bonnefoy, Samtec jean-remy.bonnefoy@samtec.com

Ted Ballou, Samtec ted.ballou@samtec.com

Istvan Novak, Samtec istvan.novak@samtec.com

Gustavo Blando, Samtec gustavo.blando@samtec.com

Scott McMorrow, Samtec scott.mcmorrow@samtec.com

Raj Mahadevan, Alphawave IP raj@awaveip.com

Abstract

The continued progression to higher data rates puts increasing demands on the design of practical SerDes channels. At 112G-PAM4, the UI is only 17.86 ps, and signal transmission in the PCB must be highly optimized for loss, reflections, crosstalk and power integrity. This paper will describe the signal-integrity and power-integrity design process, show simulated SI and PI performance correlated to measured data as well as measured eye diagrams of a test board that uses a 112G-capable silicon and high-speed compression-mount cable connectors. The resulting test channel meets the IEEE803 100G and OIF 112G-PAM4 channel operation margin specification.

Authors Biographies

Jean-Remy Bonnefoy is a Systems Engineer for Samtec's Signal Integrity Group. He is involved in the design of high-speed test systems and he leads the hardware development of evaluation and demonstration platform for high data rate interconnects.

Ted Ballou is a Signal Integrity Engineer at Samtec and previously at Intel, working with high speed PCB interconnects and serial interfaces.

Istvan Novak is a Principal Signal and Power Integrity Engineer at Samtec, working on advanced signal and power integrity designs. Prior to 2018 he was a Distinguished Engineer at SUN Microsystems, later Oracle. He worked on new technology development, advanced power distribution, and signal integrity design and validation methodologies for SUN's successful workgroup server families. He introduced the industry's first 25µm power-ground laminates for large rigid computer boards and worked with component vendors to create a series of low inductance and controlled-ESR bypass capacitors. He also served as SUN's representative on the Copper Cable and Connector Workgroup of InfiniBand, and was engaged in the methodologies, designs and characterization of power-distribution networks from silicon to DC-DC converters. He is a Life Fellow of the IEEE with twenty-nine patents to his name, author of two books on power integrity, teaches signal and power integrity courses, and maintains a popular SI/PI website. Istvan was named Engineer of the Year at DesignCon 2020.

Gustavo Blando is a Senior Principal SI 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. Gustavo has twenty-five years of practical experience in Signal Integrity, high speed circuits design and have participated in numerous conference publications.

Scott McMorrow currently serves as CTO for Samtec's Signal Integrity Group, 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.

Raj Mahadevan is co-founder and COO at Alphawave IP, a leading DSP based connectivity Intellectual Property (IP) company. He has over 20 years of experience in high performance mixed-signal IP business. He brings a unique combination of a wealth of superior technical expertise, project management, design methodology, and customer relationship skills. He has an outstanding track record of numerous first-pass IP launch successes. Prior to Alphawave, he cofounded and was a Director of V Semiconductor Inc, a successful serial interface IP company. He led a multidisciplinary team from concept formation to production of the highest speed multistandard serial interface IP portfolio on the market. Prior to this, he was Director of Engineering at Snowbush Microelectronics where he was instrumental in leading this startup to become a leading IP provider. Mr. Mahadevan has a Bachelor of Science and a Master of Applied Science degrees in Electrical Engineering from University of Toronto.

I. Introduction

With higher data rates, careful consideration must be given to the design of practical SerDes channels. At 112G-PAM4, signal transmission in the PCB must be highly optimized for loss, reflections, crosstalk and power integrity. This paper describes the signal-integrity and power-integrity design process for a test board that uses a 112G-capable silicon chip in a 16x16 BGA package [I.1] and high-speed compression-mount cable connectors [I.2]. This platform enables interfacing to additional test loads with coax cables, supporting a range of use cases where a separable connection is defined with specified limits for IL, RL, ERL, COM or other signaling metric.

For this 112G test channel, the total loss from die attach to the RF cable end needed to meet an aggressive 3dB limit at 28GHz. An additional design goal was to use PCB materials and stackup configuration that would support reasonable cost and routing flexibility to demonstrate the practical feasibility of the implementation. A 10-layer stackup with I-Tera MT-40 dielectrics and VLP copper was chosen, which included four routing layers, of which only two were required for the present evaluation topology. Four routing layers in a 10-layer stack is an aggressive cost point that posed real design challenges for via breakouts and for clean power distribution.

To meet these challenges, the design was optimized to minimize total PCB routing while enabling required mechanical features. A key design choice was to place the pressuremount connector on the bottom side of the PCB to allow closer positioning to the device package. This eliminated conflict with keep-outs for mounting holes and resulted in PCB net lengths below 16mm for the critical TX channels, each feeding an attached 6" coax cable.

With the overall topology decided, the package and PCB interconnects were modeled in iterative design cycles to achieve optimal signal integrity for loss, reflection and crosstalk. Power delivery became a pressing issue because of the minimal layer count in the chosen stack-up. With four separate power rails for the BGA alone, there weren't enough power/ground layers in the stack to support fully isolated power fills. The BGA balls for these rails were located in a tight region, and power routing had to be distributed opportunistically across several stack-up layers. To manage this complexity, iterative power integrity analyses was carried out to optimize per-pin current density and ensure the lowest possible impedance transfer for each rail. While the stackup was chosen to mimic real-life designs, the active power sources were placed further away from the critical components and were chosen to be the quietest possible, regardless of size and efficiency. As such, linear regulators were used, on some rails in a cascaded fashion. The selection of bypass capacitors ensures good stability of linear regulators and sufficiently low impedance to keep the self-generated noise by the silicon at a very low level.

The paper documents the major steps of the design flow, showing simulated SI and PI performance correlated to measured data as well as measured eye diagrams.

II. Optimizing a 112G test channel

Design challenges and goals

The evaluation vehicle, consisting of the BGA package, PCB interconnect and 6" of lowloss coax, was targeted to meet an insertion loss budget of 3dB (on the shortest TX net) at 28Ghz. Given the BGA package was already designed by a third party, the layout of the PCB was highly constrained. Our validated models for the cable and cable end showed 0.6dB loss at 28Ghz, but we had only simulated loss for the package (no measurements). Subtracting the estimated 1.5dB package and the cable from a 3dB total loss budget, this left about 1dB for the PCB. This must account for the BGA vias, routing, test point vias and the coaxial test connector body. The resulting layout is shown in *Figure II.1*.



Figure II.1: TX channels (short, Layer2) and RX channels (longer, layer9)

The figure shows the package pinout assigning the TX signals to the outside row of the 1mm pitch BGA balls. Two rows directly above the TX pins are the RX pins – all adjacent pins are GND. The signal pin pairs are aligned with the package boundary, enabling a symmetric differential via breakout. Ball pitch of 1mm allows ample room for tightly coupled diff pair escape routing. For the landlocked RX pins, via-in-pad was a natural design choice, and this approach was used also for the TX pins.

Stackup

A ten-layer PCB stackup (*Figure II.2*) was chosen using Isola I-Tera MT40 dielectric material. This features low Dk and very good if not world-beating Df. I-Tera MT40 is cost-effective, comes with VLP copper and is readily available at fab houses. In our case the logistics outweighed any slight IL penalty relative to using super-low loss materials with long lead times. For the short net lengths we were deploying (11mm for the TX net) this IL opportunity cost amounted to only 0.1dB.

The stackup is a core-out construction that allows very good layer-layer alignment for the CDD (Controlled Depth Drilling) vias we deployed. All dielectrics have the 1067 fiber weave, preferred for consistent warp and weft spacing, and the PCB was rotated in the panel to mitigate periodic fiber weave effects. At 1.7mm total thickness (67mils), the board is thin enough to easily support the 0.2mm drills used for the thru vias.

As a rule, we use anisotropic dielectric models with Dk scaled up a bit for the heterogenous composition of the dielectric the signal sees during via traversal. This improves correlation to the characteristic lower impedance discontinuities seen with most vias.

	Calc			
Layer	Thickness	Primary Stack	Description	Dk / Df
Layer - 1	0.0010 0.0026	 11	Taiyo 4000-MP 1/2oz Mix (Std Pit)	3.60 / 0.0190
	0.0050	0.0050	I-Tera MT40	3.17 / 0.0023
Layer - 2	0.0006		1/2oz Sig	
	0.0052	1067 - 76% 1067 - 76%	I-Tera MT40	3.08 / 0.0020
Layer - 3	0.0006		1/2oz P/G	
	0.0050	0.0050	I-Tera MT40	3.17 / 0.0023
Layer - 4	0.0006		1/2oz Sig	
	0.0052	1067 - 76%	I-Tera MT40	3.08 / 0.0020
Layer - 5	0.0006		1/2oz P/G	
	0.0140	0.0140	I-Tera MT40	3.56 / 0.0033
Layer - 6	0.0006		1/2oz P/G	
	0.0052	1067 - 76%	I-Tera MT40	3.08 / 0.0020
Layer - 7	0.0006		1/2oz Sig	
	0.0050	0.0050	I-Tera MT40	3.17 / 0.0023
Layer - 8	0.0006	(2.001)	1/2oz P/G	
	0.0052	1067 - 76%	I-Tera MT40	3.08 / 0.0020
Layer - 9	0.0006		1/2oz Sig	
	0.0050	(2-1067)	I-Tera MT40	3.17 / 0.0023
Layer - 10	0.0026		1/2oz Mix (Std Pit) Taivo 4000-MP	3.60 / 0.0190

Figure II.2: PCB layer stack

PCB design and via breakout optimization

~ .

Transitioning TX from top-mounted BGA to bottom-mounted test connector involves one CDD via to layer 2 and one thru via from layer 2 to the bottom (or vice-versa for RX layer 9 routing). This complementary via topology (*Figure II.3*) is good for minimizing via impedance discontinuities since stubs are inherently minimal.



Figure II.3: Complementary vias minimize stubs

With a 1dB loss budget for the PCB, via design was critical to minimize loss of signal energy to reflection, crosstalk and mode conversion. The CDD and thru vias were optimized using 3D simulation tools, checking with an ERL metric for each sweep. *Figure II.4* shows details for the CDD via at the BGA. To promote impedance control, we added five auxiliary ground vias inside the ground pin extents, impedance-matching teardrops at the layer 2 signal pads and offset ground relief compensation on the lower reference plane. Similar steps were taken with the thru via to layer 9 seen in *Figure II.5* which peeks down into the antipad stack to show the grounded antipad rings added on each layer.



Figure II.4: CDD via

Figure II.5: Thru via

On the test connector side of these nets we have single-ended vias (*Figure II.6*) that deploy dual rings of ground return vias to kill stray fields from mode conversion. In between the terminal vias we route differentially in three zones: 10mil differential escape pitch at the BGA, 14mil open field pitch to ensure that the 1067 fiber weave mitigation rotation is predictably effective, and uncoupled differential to the coax connectors.



Figure II.6: Via to coaxial test connector

One problem was defining the channel impedance, with a 900hm package needing to talk eventually to 500hm coax (1000hm differentially). Since even optimized via transitions present low-impedance discontinuities, we used 920hms at the BGA escape and the open field routing. When the routing opened up for the widely spaced coax test points, we maintained the original differential trace pitch and allowed the uncoupled impedance to rise toward 1000hms. This enables reasonable trace widths for lower copper loss.

The pre-tapeout signal integrity modeling used a pieces/parts approach, relying on concatenation of simulation s-parameters to assemble a full 4-port channel. We extracted the package layout into our modeling tool [II.1] and simulated using published Dk/Df values for the GL102 dielectrics and a lossy copper model for the package build-up traces. Breakout vias, PCB routing and cabling were all separately simulated, and the design was taped out. When prototypes were available, a bare-die BGA package was reflowed to the PCB and sent off to a test facility for 4-port measurement. (Note that crosstalk measurements were not made for this test channel).

Initial correlation

Correlation to the simulation model was poor, revealing grossly under-modeled impedance discontinuities and missing resonances in the channel IL. The concatenated channel IL and TDR are plotted in red in *Figure II.7*, shown in overlay against the end-to-end channel measurement of the dieless package, PCB and cable assembly.



Figure II.7: Initial simulation correlation for TX channel (measured is in red) – second TDR is an annotated detail

Examining root cause

The disappointing correlation called for a deep dive to root cause this. A post-mortem review unearthed several issues with the initial simulation:

- The coax test connector model had been truncated to speed up via optimization sims final channel simulations required the full connector model.
- Unused via pads were suppressed in the package model this is incorrect, as package microvias will always have pads on both ends.
- Package extraction modeling built the microvias at the nominal 0.05mm diameter; real package microvias have a v-shaped cross section with the top being 0.6mm due to inherent fab process. For better correlation we increased microvia diameter to an average 0.055mm.
- Package copper roughness was originally a Groisse model with 0.35um factor we revised this to a Huray roughness model with surface ratio of 1.5. This would put the package buildup copper somewhere between RTF and VLP for a PCB, which we think is reasonable and did improve the correlation IL and TDR.
- Various routing segment trace widths needed adjusting to match measured impedances.
- PCB thru vias had more capacitance in the measurement we increased drill size from nominal 0.2mm to 0.22mm to account for size increase due to tolerance and small drill size wander.

The changes listed above were straightforward in review, apart from package copper roughness which was an engineering judgment call guided by the correlation work. At no point was an isolated package measurement made – prototype hardware was in short supply and only one bare die package was allocated to an assembled measurement vehicle. But there is more to overall package IL than just routing loss, as we were to discover. Every BGA has a solder ball attach, which must also be considered. In this case the solder ball model was based on assumptions about reflowed solder ball dimensions rather than on hard data. And for this key discontinuity, the concatenated simulation approach wasn't optimal for modeling of the interaction between the BGA package and PCB.

Another barrier to correlation root cause was the difficulty of getting high quality measurements of the PCB itself to reevaluate our understanding of board loss. While we did have full channel measurements from the test house, we had to make our own PCB measurements. Our internal test lab only had single-ended microprobe capability for the BGA pads, and the measurements were made without terminating the other legs of the differential nets.

Validating PCB loss models

The first question to answer was how well our PCB routing model predicted the measured loss. Fortunately, the layout designed ensured that the only differences among TX nets (or among RX nets) was purely in terms of open-field routing length. Single-ended loss is not quite the same as differential loss, and unterminated legs from a TX pair cause marked periodic IL suck-outs from coupled intrapair reflections. This was not an

ideal measurement regime, but it was what we had available. To tease out "pure" trace loss from a series of noisy, single-ended BGA-to-cable measurements, we tabulated the measured IL data at 20Ghz, which was a reasonably observable data point along the IL characteristic. Plotting the various net lengths against the measured IL, we were able to find a straight slope of 0.0435dB/mm (with offset for "fixed-cost" routing elements such as vias and interconnects). To match the problematic measurement method, we then resimulated a length sweep of the open-field differential routing using single ended wave ports with the other leg unterminated. Plotting length vs. loss here was easier with no extraneous channel features to offset. The simulated loss slope at 20Ghz was the same as the measured loss at 0.0435dB/mm. This finding thoroughly validated our PCB transmission-line models.

Rethinking package/PCB interaction

For the 6" cable portion of the channel, the net loss was only 0.6dB at 28Ghz, confirmed by cable measurements. With package loss knowable from simulation only, the PCB routing loss model confirmed by correlation, and the cable loss well understood, the miscorrelation root cause pointed strongly to the impedance discontinuities. Chief among these by far was the BGA attach.

There were two problems to solve in root causing this: first, we needed to improve the simulation methodology itself by solving an attached BGA package and PCB in a unified, co-meshed simulation model. Once we deployed this approach, we saw improvement in the correlation resonances. *Figure II.8* shows a unified hierarchical simulation model with BGA package on PCB and coax connector bodies mounted on the underside. The simulation tool did remarkably well to capture boundary interactions using carefully managed pieces/parts concatenation, but co-meshing gave better fidelity, not to mention being easier to manage. *Figure II.9* shows the difference in TDR correlation between two separate, concatenated simulations vs. a unified hierarchical simulation (this comparison holds the solder ball dimensions constant for both methods).



Figure II.8: Combined, co-meshed simulation model, pkg + *pcb* + *connectors*



Figure II.9: Impact of concatenation vs combined, co-meshed simulation

Package/PCB solder ball interface

With a unified co-meshed simulation model, we were able to better study the impacts of reflowed solder ball dimensions.

Solder balls are specified in nominal diameters with a diameter tolerance. For 0.6mm presoldered balls, the tolerance is quite significant at +/-0.1mm. In addition to the raw ball, metal in the solder paste slightly increases reflowed size. Because assembly houses are concerned primarily with reliable solder joints (up to a point more solder is better), the mechanical objectives here may conspire to result in ball size distributions on the plus tolerance side. We certainly saw that to be the case in our assemblies. But for signal integrity, larger ball size (and larger tolerances) means nothing but trouble.

Separate from ball size itself is the shape of the reflowed ball joint. This is determined by multiple factors including dimensions of BGA and PCB pads and solder mask openings, and by the reflow process temperature profile. At the solder melting point, surface tension drives the solder to the solder mask boundary and reaches an equilibrium with the force from package weight. This is perturbed by flux outgassing which can leave voids buried in the re-solidifying balls. Process variations, coplanarity tolerances and the design of pads and solder masks can result in a wide range of reflowed ball shapes (*Figure II.10*).

Modeling the reflowed solder ball is a challenge. Our simulation tool supports only a symmetric ball (*Figure II.11*) which can be shaped using three parameters. But for reflowed balls with a variety of cross sections, this is an approximation that will introduce some high frequency miscorrelation.





Figure II.11 Simulation solder ball model

For better correlation, we needed to understand two issues: first, how does size and shape of the reflowed ball impact the simulation? Secondly, what reflowed dimensions did we have in our measured test assembly? To answer the first question, we ran longitudinal sweeps of reflowed ball size and again of reflowed ball shape, with the original ball volume held constant (*see Figures II.12-13*). The somewhat surprising result of these studies was that reflow shape is more predictive of deep TDR discontinuities than simply ball size itself. As effective height of the joint is reduced, capacitance and field interactions grow at a nonlinear pace.



Figure II.12: Swept reflowed ball size (0.55mm to 0.65mm pre-reflowed diameter) with fixed height; legends indicate reflowed base, width and height in microns



Figure II.13: Fixed pre-reflowed diameter (0.625mm), swept reflowed ball height (0.52mm to 0.36mm); legends indicate reflowed base, width and height in microns

The question of our actual ball size could only be proximally determined since the measured assembly was not in our possession and hence never imaged. Furthermore, it was assembled by a third party. Using their profile and mask screening recommendations, we did our own assembly of a nonfunctional package/board set and had this scanned by our imaging lab. Additionally, we borrowed another functional package/board assembly and had this imaged as well. The lab was able to take automated measurements of all reflowed balls in the BGA instance (*Figure II.14*). The data from CT-scan shows that reflowed solder ball dimensions can vary widely even for the same BGA device. These scans don't measure reflowed ball height, however. To get this, the lab used precision feeler guages to determine 0.35mm joint height for the functional board we imaged.



Figure II.14: Reflowed solder ball dimension variance over a single BGA instance

Additionally, our silicon partner provided images from their pre/post reflow testing showing dimensions in line with our data (*Figure II.15*).



Figure II.15: Pre-soldered and reflowed solder ball dimensions

Re-simulation for improved correlation

With this knowledge in hand, correlation simulations were re-done with a few clearly needed trace impedance tweaks for this particular case. We settled on a reflowed ball model with 48um base, 760um width and 34um height based on correlation sensitivity. One cited source [II.2] had a closed form equation that we used to calculate reflowed height at 30um, but this was for a complex reflow shape not supported by the simulation tool. Even with improved models, we are still missing some IL resonance at higher frequencies (*Figure II.16*). We believe most of the miscorrelation is due to the reflowed ball model approximation not capturing the more complex shape of the real thing. Here, the choice of ball model parameters has a strong impact on the correlation above 20Ghz. It's also possible that even with die probe calibration there is still some residual probing impact in the measured data, along with a small uncertainty in VNA accuracy.



Figure II.16: Improved sim correlation for TX, RX channel (measured is in red)

Learnings

The correlation effort pushed our simulation methodology to more comprehensive model integration and a better appreciation of the fine points of physical modeling for high frequency design. The perspective gained allowed a more accurate estimate of the package loss contribution and a realization that the original 3dB total loss target was too optimistic. *Figure II.17* plots the relative loss budget contributions of package, PCB and cable assembly to total modeled loss.



Figure II.17: Learnings regarding Loss budget allocation

III. PDN Design

The goal was to create a universal power distribution network for the evaluation board that allows the user to power the board from a single 5V supply or alternately feed each supply rail from a dedicated bench supply. The model of the bench supply and its connection to the evaluation board was not known, that branch of the power tree was not included in the analysis and the validation. On the evaluation board the cost and size of the PDN is of secondary importance and we may want to intentionally 'over-design' the PDN such that its contribution to any degradation of performance is minimized.

Architecture

The relatively low current consumption allowed us to use a set of cascaded linear regulators and jumper-selectable internal or external supply. The power tree for the internal supply is shown in *Figure III.1*. The main supply rail parameters and requirements are summarized in *Table III.1*.



Figure III.1: The internal power tree of the evaluation board

Net	DC voltage [V]	Max current [A]	Max transient current [A]	Allowed deviation [mV]	Target impedance [mOhm]
P_VDD	0.75	3	1.5	3	2
PA_VDDL	0.75	2	1	2	2
PA_VDDH	1.2	2	1	2	2
P_VDDH	1.8	1	0.5	1	2

Table III.1: Target numbers for the main supply rails

The unusually tight noise target on the supply rails is in line with the intention to make sure that the impact of the noise on the board PDN can be neglected.

Since the maximum current and target impedance for the four rails are similar or identical, a uniform capacitor selection was decided upon. To help the regulator to maintain the low impedance, ten pieces of 470 uF polymer bulk capacitors were added to each rail. To lower sensitivity, the 'Big-V' impedance profile was implemented [III.1]. The ceramic capacitors were chosen to maximize capacitance. 10 uF 0402 capacitors were added to the back side of the chip, bridging adjacent power-ground pads. At the linear regulator output 22 uF ceramic capacitors complemented the PDN. The capacitor selection and the lumped impedance profile for the P_VDD rail is shown in *Figure III.2*.

Four parallel capacitor banks	C1	tol. [%]	C2	tol. [%]	C3	tol. [%]	C4	tol. [%]	
Capacitance C [F]:	4.70E-04	20	1.00E-04	20	2.20E-05	20	1.00E-05	20	Fmin[Hz]
		-20		-20		-20		-20	1.E+03
Ser. resistance ESR [ohms]:	0.015	0	0.9	0	0.004	20	0.005	20	Fmax[Hz]
		-50		-50		-20		-20	1.E+09
Ser. inductance ESL [H]:	3.00E-09	20	5.00E-09	20	1.00E-09	20	1.00E-09	20	1
		-20		-20		-20		-20	Total:
Number of parts in bank:	10		1		3		36		50

Impedance magnitude of parallel capacitors [ohm]



Figure III.2: Lumped impedance profile of the P_VDD rail

Analysis

The DC drop and AC impedance on all rails were checked with a hybrid solver [III.2]. The stackup was updated during the design process to include one-ounce copper layers to lower the DC drop. Though for an evaluation board the static portion of the DC drop could easily be compensated by adjusting the DC source, the DC drop target was also set aggressively to a maximum of 1% of the nominal voltage. This also helps the component placement by ensuring that the low target impedance is not compromised by series DC resistance across the planes. The plane shapes and layer allocations were optimized in multiple simulation passes to produce similar DC drop on the different rails. Since the board is relatively small compared to the size of BGA pinfield, the BGA pins were not grouped together; instead, the load current was distributed equally across independent current sinks connected to each power pin and its nearest ground pin. Also, to minimize rail-to-rail crosstalk, any vertical overlap between power planes belonging to different nets was avoided. Since the maximum current is relatively low, current density and temperature rise was not a concern; this optimization targeted DC drop only. Figure III.3 shows the simulated end-to-end DC drop on the four rails as well as the potential surface under the BGA pinfield.



Figure III.3: Potential surface in the pinfield and DC drop table

The AC impedance was simulated at the dedicated test points and in the chip's pinfield; the same locations where later correlation measurements were taken. For the sake of easy comparison, the simulations were also done mimicking the Two-port shunt-through impedance measurements: top-bottom at the same via pair at the test points and second-adjacent power-ground via pairs from the top in the BGA pinfield. *Figure III.4* shows the locations for the AC impedance simulations.



Figure III.4: Locations for the impedance simulations

Starting from DC, the low-frequency AC performance is dominated by the dynamic performance of the DC source. The vendor-supplied models and circuit simulator [III.4] were used to simulate the startup behavior and making sure that the cascaded regulators were stable with the chosen external components. The output impedance of the downstream regulator was simulated with lumped external components. With its original model and the lumped assumption of attached input/output components, the simulated output impedance is shown in *Figure III.5*. Only one of the rails is shown. The simulation suggested that the impedance target can be met with the selected parts. The impedance profile of the full board was simulated with the hybrid solver in two scenarios: by leaving the voltage regulator outputs open and by placing an (arbitrary) series R-L element with 10mohm and 3uH values to crudely mimic the regulator output impedance, intentionally overstating its equivalent inductance so that we see the transition between the regulator impedance and bulk capacitor impedance (*Figure III.6*). This compares to the lumped impedance profiles shown in *Figure III.2* and *III.5*.



Figure III.5: Simulated output impedance magnitude of power rails with original regulator model, powered and unpowered

The difference is the inclusion of plane and via interconnect structure and potentially taking into account the DC bias dependence of ceramic capacitors. It can be done by manually altering the capacitance values or by using dynamic models [III.3]. Figure *III.7* shows the full-board impedance profiles at the four test points and BGA pinfield with the regulator outputs left open.



Figure III.6: Impedance magnitudes at the test points with the crude regulator models



Figure III.7: Two-port self-impedance port locations (left) and magnitudes (right) in the BGA pinfield

Measurements and correlation

The PDN validation was done with a VNA [III.5] with a common-mode toroid to suppress the cable-braid ground-loop error. Connections to the test points and pinfield were done with home-made semirigid probes. To check the regulators and bulk capacitors, the measurements were taken in the 100 Hz – 10 MHz frequency range with and without input power applied. Measurements in the BGA pinfield were also taken in the 1 MHz – 1 GHz frequency range. The measurement setup photo is shown in *Figure III.8*.



Figure III.8: Measurement setup. Thru calibration on the left, powered impedance measurement on the right

After taking reference measurements to confirm the dynamic range of the setup, first the impedance was measured at the test points with no input power, followed by powered regulators with minimal load current, and later sweeping the load current in the full range. It was noticed that the powered impedance profile showed two discrepancies: the low-frequency impedance seemed to settle above 10 mohm and with minimum load current there was a large impedance peak at 12 kHz approaching 100 mohms. A quick debug confirmed that the measured impedance values were correct and the high DC resistance was caused by the series jumper that selected the on-board regulators, which accidentally were simulated with practically zero resistance.



Figure III.9: *Reference measurements on the left, unpowered and powered impedance with minimum DC load at TP34 on the right*

By measuring a number of jumper pieces with different headers, it was concluded that each jumper with a size of 25-mil square posts represents approximately 12 mohms of DC resistance. To remove this error, the rest of the validation measurements were taken with the jumpers shorted with soldered shorting bars. *Figure III.9* shows one of these measurements at TP34.

To find out the source of the impedance peaking at 12 kHz, the following one-by-one debug steps were taken: the upstream regulators were disabled and bypassed to rule out the interaction of cascaded regulator loops, lossy bulk capacitors were added to the input rail to rule out the impact of bench supply that was feeding the setup, the SET resistor landing point on the ground plane was altered and the associated loop was further minimized to rule out layout mistakes. As none of these steps had any noticeable influence on the peaking, next an evaluation board for the regulator was obtained and measured without any modification. The evaluation board showed similar behavior, though due to the much less output capacitance, the peaking occurred at higher frequencies. It should also be noted that the impedance peak went down significantly with heavier loads as it is shown below in *Figure III.10*.

The regulator's data sheet does not include information about the output impedance, but it does include a plot for closed-loop gain. When the closed-loop gain of the regulator's evaluation board was measured and simulated, it showed -though to a lesser degree- a similar difference: the simulated loop bandwidth appeared to be too optimistic. After the data was presented to the manufacturer of the regulator, an updated simulation model was created and posted. On the evaluation board the updated model correctly captured the bandwidth of the closed-loop gain and did also improve the impedance correlation. With the full load-current sweep data, *Figure III.10* shows the measured output impedance at different DC load current values together with the simulated regulator impedance with lumped output capacitor models. Comparing the OFF impedance curves we see very good agreement at low frequencies in the capacitive region, and see an increasing deviations at higher frequencies where the lumped models push the simulated impedance down. With power applied, the simulated model showed little dependence on the load current and therefore a single trace referring to 1A load current is included in the figure.

Note the strong load-current dependence of the measured low-frequency peak magnitude, starting around 100 mohm with a few mA DC load (presented by the VNA ports) and then quickly settling to around 10 mohm. The plot also illustrates that the updated regulator model improves the correlation, but still does not capture the load current dependence. The various test results confirmed that the originally set 2-mohm target impedance was unrealistic and too optimistic for the selected regulator model. After realizing this, the question remained: can we improve the performance of the PDN with a BOM change? By giving up the 2-mohm target, we may be able to make the PDN simpler by using fewer and higher-ESR bulk capacitors and at the same time potentially also reducing the peak. *Figure III.11* shows the measured result of one such possible

approach by replacing the ten 470 uF 10 mohm bulk capacitors with a single 2700 uF 7 mohm capacitor. All other components were left unchanged. Note that while the mid-frequency impedance increased, at and above 0.5A DC load current the impedance stays below 10 mohm.



Figure III.10: Output impedance measured at the TP34 test point as a function of load current



Figure III.11: Output impedance measured at the TP34 test point with modified BOM, by replacing the ten 470 uF capacitors with a single 2700 uF bulk capacitor

Figure III.12 shows the high-frequency correlation in the BGA pinfield for the P_Vdd rail, corresponding to the TP34 test point.



Figure III.12: Simulated and measured self-impedance of the P_Vdd rail in the BGA pinfield at the pins identified in Figure III.4 (on the left) and comparison of extracted inductance (on the right)

The primary purpose of such validation in the pinfield is to check the high-frequency inductance of the PDN structure under the package. High-frequency in this context means the frequency range that approaches the die-package resonance frequency. The plots show one measured and two simulated curves. Measurement was done from the top side of the board with both probes and the self-impedance was approximated by the transfer impedance between the two pad pairs, which -in such a case- is the only useable term of the .s2p file. Simulated data on the other hand is valid for both the transfer (solid line) and self (dashed line) terms and therefore both are shown. Note that the measured inductance contains an error term due to the loop coupling between the hand-held probes open pins (the semirigid probe pins are approximately 40-mil long with 50-mil spacing).

IV. Measured channel performance

The signal integrity and power integrity designs were validated further with the SerDes silicon chip itself. The evaluation board was powered from a single 5V supply such that the four supply rails to the silicon were delivered by the PDN architecture exactly as described in *Section III*. The fully populated evaluation board is shown is *Figure IV.1*. The compression mount cable connector is mounted opposite the silicon on the back of the board.



Figure IV.1: Pictures of the powered evaluation board with the connector system mounted opposite the test chip.

Transmitter Eye diagram

The SerDes transmitter's Built-In Self-Test (BIST) was configured to generate PRBS-31 PAM4 signals at a data rate of 106.25Gbps. A sampling oscilloscope with a precision waveform analyzer module [IV.1] was connected directly to the end of the cable connector system. It was shown in *Section II* that the insertion loss of the channel, from the die to the RF cable end, was measured at 4 dB at 26 GHz. These measured s-parameters were loaded into the de-embedding tool of the oscilloscope and post processed to generate the transmitter PAM4 eye shown in *Figure IV.2*.



Figure IV.2: 106.25 Gpbs PAM4 Tx measured eye diagram

Note that the transmitter generates a clean eye diagram at 106.25Gbps, with good margins for all three PAM4 eyes, exhibiting excellent linearity and signal to noise ratio.

Next, a real-life channel representing a long reach (LR), mid-board to cable backplane application, was connected to the end of the cable-connector system. The overall insertion loss of this test channel, including the package and the evaluation board, is shown in *Figure IV.3*.



The test channel, before returning to the receiver, included a total of eight different interconnect transitions, 4" of PCB traces on I-Tera MT40 material, and 64" of twinax cable. The details of the test channel are shown in *Figure IV.4*.



Figure IV.4: Details of test channel

Two adjacent SerDes lanes were used simultaneously in this setup, lane1 and lane2, for the purpose of including crosstalk in the measurements. The results were similar on both lanes. For simplicity, the results of lane2 alone are detailed below. The eye diagram captured at the receiver and the histogram plots of the raw data are shown in *Figure IV.5*. Confirming the oscilloscope measurements, the results show that the transmitter generates a clean eye diagram.



Figure IV.5: Eye diagram capture at the receiver and histogram plot

Bit error rate

The receiver BIST was run to check for errors. PRBS31 data is accumulated at the receiver for 10 seconds. After the timer has elapsed, the error counter stops accumulating and the BER is extrapolated past the target BER of 1e-6. The vertical bathtub plot and BER for each eye of the PAM4 signaling are shown in *Figure IV.6*.



Figure IV.6: Vertical sample bathtub plots showing BER for all three PAM4 eyes

The pre-FEC BER, before correction algorithms, was measured around 1e-8 for this LR channel. These results meet the IEEE P802.3ck suggested BER < 1e-4 by several orders of magnitude.

V. Conclusions

Overall, the design was an unqualified success in its key function to support the silicon with low insertion loss and clean power distribution. This is evident from the low BER demonstrated. There were some important learnings, however, from the signal integrity correlation effort. The under-modeled loss contribution from the package was an eye opener that will inform future designs in the value of accurate physical modeling and checking the design against real data. The BGA package attach was similarly notable: with many designs the reflowed solder ball can be a negligible issue, but for very high data rates it starts to matter. For some use cases the designer might consider specifying a solder ball smaller than 0.6mm to leverage the better tolerances and smaller scale of the parasitic impact.

On the PI side, the regulator's SPICE models were updated after the measurements showed impedance differences at the low frequencies. This underlines the importance of testing and validating the power converters before freezing the design, which in this case -for various practical reasons- we failed to do in time.

With the SerDes test silicon populated, the overall SI and PI design was functionally validated with eye diagrams and BER measurements. The delivered board design meets the requirements for a silicon evaluation platform, enabling clean transmission of 112G-PAM4 traffic with wide eye margins and low BER.

VI. Acknowledgements

The authors wish to thank Stephen Ruscak and the modeling team of Analog Devices, Mark Crain, Lance Cooper and Jason Sia of Samtec, and Martin Wang of Alphawave for their support.

VII. References

- [I.1] AlphaCORE LR MSS, https://www.awaveip.com/en/products/connectivity-ip-cores/alphacore-lr-mss/
 [I.2] DE704 70 CH D ll E Hill D for Tart S at the second s
- [I.2] BE70A 70 GHz Bulls Eye High Performance Test System, https://www.samtec.com/products/be70a
- [II.1] Ansys HFSS 3D Layout https://www.ansys.com/products/electronics/ansys-hfss
- [II.2] As-Shipped vs. Mounted Height for BGA and LGA Packages, https://www.skyworksinc.com/-/media/SkyWorks/Documents/Articles/SMTA 102018.pdf
- [III.1] TecForum TF-MP3 "Comparison of Power Distribution Network Methods," DesignCon2006, Santa Clara, CA, February 6-9, 2006
- [III.2] Ansys SiWave, https://www.ansys.com/products/electronics/ansys-siwave

- [III.3] "MLCC Dynamic Model Supports Circuit Simulations," https://article.murata.com/en-us/article/mlcc-dynamic-model-supports-circuitsimulations
- [III.4] LTspice, https://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html
- [III.5] E5061B ENA Vector Network Analyzer, https://www.keysight.com/us/en/product/E5061B/e5061b-ena-vector-networkanalyzer.html
- [IV.1] N1000A DCA-X oscilloscope with an N1060A precision waveform analyzer, https://www.keysight.com/us/en/product/N1060A/precision-waveformanalyzer.html