DesignCon 2019

Signal/Power Integrity Optimizations in an IoT Automotive Package

Benjamin Silva, Intel Corporation <u>benjamin.p.silva@intel.com</u>

Yan Fen Shen, Intel Corporation yan.fen.shen@intel.com

Mohamed Eldessouki, Intel Corporation <u>mohamed.s.eldessouki@intel.com</u>

Abstract

The demand for connected smart vehicles has increased exponentially in recent years and this includes the SoC package that powers the vehicle's sub-systems. The thermal, reliability, and life cycle use conditions dictate many physical design aspects of the SoC. These physical design elements prove detrimental to both signal and power integrity, limiting the SoC's speeds and capabilities. Signal and power integrity mitigation strategies directly oppose one another, so it becomes necessary to balance both disciplines to provide the best possible overall system performance. This paper investigates the challenges, tradeoffs, compromises, and results in the design of an automotive SoC package.

Author(s) Biography

Benjamin Silva is the technical lead and manager of the signal and power integrity team within the Internet of Things Group at Intel. He has been working in the field of signal and power integrity for almost 20 years with his current focus being system level signal/power integrity optimization and design. His past work focused on package level power integrity optimization, HSIO channel modeling/margin estimation, and full channel DDR modeling. His other interests include creating new signal/power integrity flows/methodologies to accelerate the product development cycle and simulation to lab correlation.

Yan Fen Shen is a Power Integrity technical and project lead at Intel Corporation. She has 16yrs+ of experience working in Power Integrity, covering client and IoT products. Her area of expertise is SoC + platform power integrity, covering Core power, I/O power, and FIVR technologies. Her recent work includes autonomous design, neural network, and LPDDR5 development.

Mohamed Eldessouki received his B.Sc. degree from Alexandria University, Alexandria, Egypt, in 1991, and his Ph.D. degree from the Illinois Institute of Technology, Chicago in 2003. He became an Assistant Professor with Port-Said University Cairo, Egypt. In 2008. Currently he is an Analog Engineer in IOTG, Intel Corporation. His research interests include wave propagation in semiconductor devices, device modeling, and signal integrity.

1.0 Introduction

A viable product in the automotive space requires the support of very harsh use case conditions. Extended temperature operation (-40°C to 110°C), high numbers of thermal cycles (7x more cycles than a standard mobile SoC), and a long lifetime/reliability expectation (up to 15 years) require substantial changes be made to a mobile SoC to qualify it for automotive applications. Design drivers like an integrated heat spreader (IHS), board level underfill (BLUF), AEC-Q200 qualified passive components, stripline-only routing, and a thick core package substrate become key enablers to meet the aggressive use case conditions. These design drivers have a non-negligible impact to the signal and power integrity (SI, PI) of the SoC and of the system. The design drivers described above end up degrading the PI to a greater degree than the SI, so focus is initially placed on improving the PI performance. However, SI and PI find themselves in direct opposition with one another; design changes that would drastically help the PI end up seriously impinging on the signal quality and timing. This creates an environment where overall system performance needs to be considered and compromises made between signal and power integrity performance.

This paper will focus on the design process, tradeoffs, and final system performance achieved as an SoC was developed to meet the automotive requirement. BLUF and the package size precluded the use of land-side decoupling capacitors (LSC), so all on-package decoupling caps were required to be placed on the surface layer of the SoC (die-side capacitors, DSC). The limited selection of automotive qualified capacitors (AEC-Q200) and high cost of those that were available limited the total number of components that could be placed. The existence of the thick package core, which was required to meet stiffness/warpage targets, meant that all DSC connectivity had to be handled on the top half of the package in order to make each DSC effective. To achieve this implementation, much of the space that would have been allocated to routing high speed signals had to be shared or given up for power routing. The final balancing of SI and PI resulted in more slightly compromised PI performance, package layers, dual referenced LPDDR4 traces, split referencing on LPDDR4, dual referenced high speed I/O (HSIO) traces, and impedance discontinuities related to route over void (ROV).

In the end, full performance parity with a consumer graded SoC design (in terms of millivolts and picoseconds) was not possible, but the design team was able to mitigate most of the SI and PI issues that came up. The final automotive design had slightly worse electrical performance than a consumer graded mobile version (henceforth referred to as the 'mobile SoC' or simply 'mobile'), but the automotive graded part was still able to meet the workload models and use cases required of automotive. Comparisons will be drawn to this mobile SoC, which will stand as a proxy for a typical consumer grade implementation of a high performance compute SoC.

2.0 Automotive Use Conditions

In order to meet the reliability, temperature, defect, and safety standards of the automotive market, platforms and SoCs must conform to more stringent specifications than consumer grade products. Table 1 [1] shows a high level overview of the design requirements compared to the consumer space. The physical manifestation of these requirements affects the SoC in a myriad of ways and creates conditions where the desired optimizations and platform considerations for SI and PI are directly opposing one another.

Design Driver	Automotive	Mobile (Consumer)
Temperature Rating	-40°C to 110°C	0°C to 85°C
Manufacturing	Integrated Heat Spreader (IHS) Stiffener	
Temperature Cycling (TC)	7x Mobile_TC	Mobile_TC
Defects Per Million (DPM) Target	0.02x Mobile_DPM	Mobile_DPM
Product Lifetime	5x Mobile_life	Mobile_life
Product Reliability	2x Mobile_rel No μ-strip trace routing	Mobile_rel µ-strip trace routing OK
Passives	AEC-Q200 [2]	Consumer Grade
Qualification tests	AEC-Q100 [3]	Consumer

Table 1: Automotive package requirements relative to mobile

Meeting the requirements from Table 1 involves specific physical changes on the SoC and platform. Manufacturing and thermal design constraints necessitate an integrated heat spreader (IHS) to be included as part of the SoC package. The IHS coupled with the thick package core (>400µm) make for a very rigid package substrate which necessitates the use of a board-level underfill (BLUF) material between the package and the motherboard to maintain solder joint reliability; this will also ensure the temperature cycling and product lifetime/reliability constraints are met [1]. If BLUF is not used, warpage would cause the product to suffer from large numbers of BGA failures, Figure 1(a). Figure 1(b) [4] shows another failure mode where BGA cracking is caused by temperature cycling. The product design team would need to specify high numbers of non-critical to function (NCTF) pins to compensate for these BGA failures. Given a rigid landing zone with a fixed pinlist requirement, large numbers of NCTF pins has the adverse effect of requiring a larger package substrate in order to fit the required number of critical to function (CTF) pins.



Figure 1: Solder joint reliability (a) warpage causing BGA disconnects (b) BGA cracking from temperature cycling

Further driving reliability, microstrip routing on the package is not allowed due to trace cracking and delamination concerns, as shown in Figure 2 [4]. Finally, automotive qualified capacitors are larger in size than comparable consumer graded parts and the selection/availability is very limited.



Figure 2: Package delamination

3.0 Signal and Power Integrity Challenges & Tradeoffs

While many of the use conditions from section 2 have implications to both signal and power integrity, the PI performance is degraded to a much greater degree. Further confounding the design is the reality that the desired design alterations made to alleviate the PI performance bottleneck then have the unfortunate side effect of negatively impacting the SI of the various I/O interfaces. This creates a dichotomy between signal/power integrity where overall system performance needs to be considered and necessitates compromises between signal and power integrity performance.

3.1 PI Challenges

A power delivery network (PDN) is the system of components used to deliver power from a main power source to the active circuits located on an ASIC. The PDN includes voltage regulators, filters, decoupling capacitors, and the internal PCB framework that delivers power from the source to the loads. For modern CPUs and SOCs, purposeful engineering of the PDN is required to provide stable power with manageable noise to ensure proper and reliable operation of the core logic and I/O interfaces [5].

The PI degradation caused by the automotive requirements is driven by a fundamental shift in strategy with regard to the on-package power delivery network. The modified strategy is comprised of four major compromises, each degrading the PI performance in different ways. The subsections below summarize these key compromises to the PDN as a function of the automotive use conditions.

3.1.1 Physical Location of Package Decoupling Capacitors

Modern processors have a voracious demand for large amounts of current to be delivered in a short period of time. The higher the frequency of this current consumption, the nearer to the CPU that reservoir of charge needs to be. If it is too far away, the parasitic resistance and inductance that delivers that power creates large amounts of noise which can reduce maximum operating frequencies. One of the most important ways to provide this power is through capacitors that are placed electrically close to the circuitry being powered. In consumer grade CPUs, capacitors are often placed on the underside of the CPU substrate (Figure 3a). These 'landside' capacitors (LSC) are then connected electrically to the CPU die through a multitude of parallel vias traversing the package substrate. This method of connectivity provides a very low inductance path between the CPU die and the LSC on package (Figure 10(a)) and is beneficial to performance. In the automotive space, LSC are not an option for various reasons. First, the use of BLUF materials could cause capacitor cracking concerns (Figure 3b [4]). Second, the automotive qualified package requires close to a full grid BGA to meet reliability and temperature cycling requirements, so there is no physical space to put LSC. Third, the physical size of AEC-Q200 qualified capacitors are too large to fit with the small, tightly packed BGA balls required.

These limitations force the design to utilize die-side capacitors (DSC), capacitors placed on the top side of the package, as the sole means of on-package decoupling.



Figure 3: (a) LSC on a standard Intel® LGA CPU (b) Capacitor cracking due to BLUF

3.1.2 Package Core Thickness

The thick package core required to meet the automotive use condition necessitates the use of larger, longer, plated through hole vias (PTH) than the thin-core or coreless packages that are typically used on consumer grade processors. The inductance of a via/PTH, approximated by (1) (where μ = magnetic permeability of free space, h = core thickness, s = spacing between power and ground vias, and r = via radius), can become a hindrance to high frequency decoupling as the core thickness scales. On the signal integrity side, these PTH vias also increase crosstalk which has a direct impact to channel margin. For the automotive use condition, the core thickness was 3 to 7x thicker than the equivalent mobile component. This put the PTH inductance squarely in the 'danger zone' from a performance perspective.

$$L = \frac{\mu}{2 * pi} * h * \ln\left(\frac{s}{r}\right) \tag{1}$$

3.1.3 NCTF Requirement

BGA balls near the package corners and around the die shadow region are susceptible to solder joint cracking due to the cumulative stress of aging and temperature cycling. Being that the automotive use condition is harsher and the reliability requirements are more stringent than the consumer space, the automotive package has a much higher NCTF load and therefore needs more pins assigned that could eventually fail. These NCTF pins take away BGA balls that could otherwise be used for signal and power. This could be mitigated by simply growing the package size, but cost, reliability, and market limitations put an upper bound on the allowable package size. Hence, the result is increased signal:ground:power ratios within the package which put an overall strain on the SI/PI return paths. Figure 4 shows an example ballmap with the NCTF pins colored in red. Note that for the automotive package discussed in this paper, the NCTF requirement was 2x that of the mobile SoC.



Figure 4: Example NCTF map (NCTF colored red)

3.1.4 Automotive Graded Passives

The final main PI compromise is the requirement for utilizing AEC-Q200 graded components. These devices are physically larger, more expensive, and a have a smaller catalog to select from than their consumer graded counterparts. In general, the current available AEC-Q200 passives on the market are limited to components that are large/tall and are not a very good fit for package level design, where the geometries are so small and the form factor constrained to meet customer needs. Table 2 shows a relative comparison between several typical package and PCB level capacitors. Note that for similar capacitance values, the AEC-Q200 graded components are one or two case sizes larger and more than double the height. Recall that an IHS was another requirement for automotive. The IHS comes with keep out zone (KOZ) requirements that limit DSC placement and reduce the total number of capacitors that can physically fit on the package. The larger physical size of the AEC-Q200 parts combined together with the IHS KOZ requirement means that much fewer capacitors can be placed relative to the mobile SoC, further degrading the PI performance.

Capacitance	Consumer grade		Counterpart AEC-Q200 components	
	Case size	Height	Case size	Height
0.1uF	01005	0.22mm	0402	0.55mm
1uF	0204	0.22mm	0402	0.55mm
10uF	0402	0.39mm	0603	0.9mm

Table 2: Consumer vs. automotive grade capacitors

3.2 SI Challenges

Considering SI alone, the automotive use conditions outlined in section 2 do not cause serious problems. The impact to the signal integrity of large numbers of package layers, thick package core, inductive PTHs, and aggressive design rules can be mostly mitigated. The challenge comes from the rearchitecting of the PDN. Left unchecked, a solution that is optimized solely for PI will give rise to platform tradeoffs that will degrade the signal integrity of the I/O buses and will greatly reduce margins. A holistic view was needed to frame the problems in the appropriate context and to enable the engineering team to make decisions that would result in the maximum total margin from a system-level perspective. Neither SI nor PI will get everything that they desire and both sides of the divide have to think in pragmatic terms rather than focus on perfection. The subsections below discuss the

fundamental aspects of the design that had the biggest contention between SI/PI. Presented are the tradeoffs that had to be made and a brief description of their impact. Detailed results of these tradeoffs will be presented in the next section.

3.2.1 Stackup Definition & Referencing Scheme

The desired stackup for PI as seen in Figure 5(b) requires top half power routing (THPR), preferably orienting the power and ground layers adjacent to one another to minimize loop inductance. This has the side effect of pushing all signal routing to the bottom half of the package. Bottom half routing is a non-starter for SI because it is physically impossible to fit the PTHs required without fanning out the signals from the very dense C4 bump field in order for them to fit. The fanout, by definition, requires topside signal routing to be accomplished.

L1

GROUND



Figure 5: (a) Ideal stackup for SI, (b) Ideal stackup for PI, (c) Balanced compromise stackup

Conversely, the desired stackup for SI as seen in Figure 5(a) pushes all power to the back half of the package. As previously discussed and shown in Figure 10 (b) and (c), pushing the power to the far side of the package PTH proves very detrimental to the PI performance, and is not a viable option. Figure 6 shows the difference in crosstalk for single ended transmission lines between the "Ideal SI" and "Ideal PI" stackups. The "Ideal PI" stackup appreciably increases the amount of crosstalk. This is primarily due to the fact that the closest signal routing layer is on a core layer, with one of the reference layers being very far away (the thickness of the package core). This asymmetric stripline configuration then behaves more like an embedded microstrip than a true stripline, as evident by the ~20dB more crosstalk in the "Ideal PI" case.

The same analysis was carried out for HSIO. Figure 7(a) shows the differential to differential NEXT and FEXT for routing on L2 with single reference vs routing on L5 above the core with single reference. Although routing on L5 causes an increase in crosstalk, it is below -40dB and it is insignificant to signal jitter. This is due to the close coupling between pairs which provide strong return path independent on the GND plane. Figure 7(b) shows the mixed mode crosstalk parameters. Common to common mode is the most significant one and it depends on the reference plane location. The way to reduce common

mode coupling is to maintain phase relation between differential pair and increase differential to differential spacing.



Figure 6: Crosstalk performance for 13mm single ended traces with SI and PI optimized stackups



Figure 7: Crosstalk performance for 13mm differential stripline traces with 85µm pair-pair spacing (a) differential crosstalk (b) mixed mode crosstalk

One obvious answer would be a stackup that somehow allowed topside routing for both signal traces and power. Simply adding multiple layers to achieve this could provide a fix, but cost, reliability, and substrate technology limitations put an upper bound on total number of package layers allowed. It becomes clear that the only true path forward would be a stackup that implemented dual referencing. In the case of the SoC being studied in this paper, this dual referencing became the accepted strategy for LPDDR4, PCIe gen3, and USB3.1. The final agreed upon stackup that balances the desires of both signal and power integrity engineers is shown in Figure 5(c). Note that the layer count in the 'balanced' stackup is increased by 2 layers over the other ideal-SI and ideal-PI stackups. This was partly a function of the 'no μ -strip routing' rule imposed by the quality & reliability ruleset, but the extra layers were also beneficial in enabling dual referencing on the LPDDR4 bus. Note also that this stackup assignment is meant to give an overall strategy for the SoC design team. The assignment was not necessarily consistent across the entire SoC; some layers featured a mix of power/ground/signal routing, depending on the region and what could be fit.

3.2.2 Dual/Split Referencing

As seen in Figure 5(c), the balanced stackup featured a situation where various I/O signals were routed as dual referenced striplines. While dual referencing in and of itself is not necessarily a show stopper, this particular design was further confounded with a split in the LPDDR4 power side reference directly

adjacent to the die edge. This automotive product relied on repurposed silicon that was already well on its way through the design pipeline. As such, the bump pattern on the silicon was fixed and was not optimized to allow for proper dual referencing. The power bumps for one of the main LPDDR4 supplies (POWER1 in Figure 8) were located on the die edge which forced the connection to happen very close to the die edge. The LPDDR4 I/O were several rows deep toward the center of the die (DQ/DQS in Figure 8). As a result, a thin slot was formed in the area between the edge power bumps and the LPDDR4 I/O. The POWER1 plane was routed on L3 and broke the continuity of the reference plane located there. The LPDDR4 traces on L2 and L4 therefore had signal routing that crossed a very long split on L3 as shown by the red arrows in Figure 8. This split referencing effected all LPDDR4 signals, including the critical DQ/DQS bytelanes.



Figure 8: Split plane on the POWER1 side of the reference for LPDDR4 stripline routing

3.2.3 Constrained Routing

For PI, it is desired to place the DSC as close to the die as possible. However, to physically break the signal lines out, very dense routing adjacent to the die edge precludes placing the necessary pads and vias as seen in Figure 9(a). Hence, for SI, a contention with PI is created due to the SI desire to place DSC as far from the die edge as possible. Pushing the DSC out degrades the PI performance, but the proximity of the DSC was dictated by the need to break out the signals. The package design rule in the breakout region was pushed to the limit, meaning that the traces were packed as tightly as possible. Figure 9(a) shows an example of this on PCIe.





Figure 9: (a) Constrained routing due to DSC placement for PCIe (b) ROV on HSIO caused by DSC pad placement

3.2.4 Route Over Void (ROV)

The layer count constraint, stackup compromise, and stripline-only routing requirement meant that regions of I/O routing were directly underneath the DSC pads. Even with one side of the stripline routed over a solid plane, the ROV on the other side caused increased crosstalk that degraded eye margins. With the number of parallel channels being broken out, the ROV was seen as being less detrimental than forcing the pair closer together or otherwise altering the routing. Figure 9(b) shows an example of an HSIO bus suffering from ROV. Multiple interfaces (single ended and differential) had to deal with ROV caused by the DSC pads. ROV is allowed only for HSIO routing but not single ended DDR. With HSIO, both pair can be routed over voids as it has less impact on crosstalk than with single ended signals.

4.0 PI Optimization & Strategies

The PI challenges set forth in Section 3.1 created an environment where the total power integrity resources were limited with respect to what was allowed on the mobile SoC product. The traditional LSC components had to be relocated to the die side, the package core was increased substantially, and NCTF requirements reduced effective pin counts. These tradeoffs created a serious contention among the different power rails on the SoC and there was great debate on where to allocate the limited numbers of DSC and how to mitigate the risk caused by the less robust PDN. The team devised a few key strategies and optimizations to maximize performance while minimizing program risk. These items are described below.

4.1 DSC Connectivity Strategy

Section 3.1.1 discussed the need for reservoirs of charge that are located electrically close to the circuitry being powered. Being a modern SoC based on state of the art Intel ® Core TM technology, there was a large demand for on package decoupling. As presented in Section 3.0, the automotive SoC had to live with a more restricted PDN that had a reduced set of resources relative to the mobile SoC. Fewer capacitors could be placed on the automotive SoC than on the mobile SoC, and these DSC capacitors had a larger loop inductance than the LSCs used previously (Figure 10). Allocating the DSC on the automotive SoC was a zero sum game, where benefiting one power rail came at the detriment of another. Careful attention was paid to this process, and the main strategy was as follows.

4.1.1 Package Layer Assignment

Section 3.1.2 described the core thickness on the automotive SoC being 3-7x thicker than on the mobile part. To mitigate the effect of the larger inductive PTH path on the high frequency power supply decoupling, the only viable option was to force all power planes (and ground returns) that connect the silicon to the DSC to be in the top half of the package. Top half power routing (THPR) is a critical enabler; without it, the DSC are ineffective and may as well be placed on the host platform. Figure 5 showed the various stackup options that were available. Figure 10 shows the general relationship between the various locations for package capacitors discussed thus far. The only stackup that enabled THPR without severe detriment to the SI domain was the balanced compromised stackup seen in Figure 5(c).



Figure 10: Relative inductance for various decoupling schemes (a) LSC in a typical mobile SoC with a thin core (b) DSC in automotive SoC with the thick core (c) DSC with THPR in automotive SoC with the thick core; all drawings to scale

4.1.2 DSC Placement

The major variable that dictates efficacy of the DSC is how far away the capacitors are placed physically from the die edge. Efforts were concentrated on pushing the design rules and KOZ requirements to the limit to allow DSC placement as close as possible. The signal breakout was another limiting factor which drove DSC away from the die in order to allow for signal break out, Figure 9(a).

4.1.3 DSC Allocation

By looking at the noise margins on all power rails, the team could see which ones had a bit of flexibility and could sacrifice DSC to be placed on other rails. In the end, the high current power rails ended up being the ones that had the largest demand for DSC and were also the rails that had the least amount of noise margin to work with. Less critical rails made do with fewer DSC and even board level decoupling.

4.2 Second Stage (Edge Cap) Decoupling Enablement

As mentioned above, some of the less critical rails had to give up DSC and live with board level decoupling. This worked for some cases, but there were some power rails where the demand for DSC was too great to be met with what could be fit on the package and standard board decoupling was not an attractive option. For these rails, the team employed a strategy to increase efficacy of board level decoupling caps. These second stage or 'edge cap' decoupling caps are placed on the top layer of the board and are further away physically and electrically than DSC, but by deliberate placement of power and ground balls at the edge of the SoC package, very short pathways can be routed on the board that allow edge caps to provide higher performance than usual. By utilizing the top two layers of the PCB to provide power and ground connectivity, the effect of the board can be reduced as much as possible. Figure 11 shows an example of how edge balls can help reduce total loop inductance to capacitors placed on the top of a PCB. Note that if the power balls were located closer to the center of the SoC

package, the current would have to stay for a longer distance on the PCB planes, which have a larger spacing and therefore a larger inductance than the package. By transitioning the current into the smaller, tighter package geometry as soon as possible, the total loop inductance to the edge caps is reduced. Per unit length, Lloop_PKG << Lloop_PCB, therefore the goal is to keep as much routing as possible on the package and as little as possible on the PCB.



Figure 11: Utilizing edge balls to reduce inductance to board level decoupling caps

4.3 Power Corridor for High Current Rail

Power corridor was utilized, where possible, to provide a robust PDN connectivity from the die to the board BGAs and onwards to the voltage regulator. Figure 12 shows an example of such implementation for the higher current consuming rail in the SoC. As described in Section 4.2, edge Power and Ground BGAs were used to capitalize this second stage of decoupling on the PCB surface layer (Figure 11).



Figure 12: Example of power corridor

4.4 On-Package Slot Inductor Isolation

Given the restrictive package form factor and package layers, some of the power rails (of the same voltage level) had to be shorted together. This could cause large noise coupling, especially when one aggressor power is noise dominant in comparison to the victim power rails. To mitigate this, a slot inductor was implemented on the package to isolate the noisy power rail from the quiet rail (Figure 13). This was an effective technique and significantly reduced the victim power noise without the need for discrete filter components.



Figure 13: Use of on-package slot inductor for noise isolation

5.0 Power Integrity Results

Once the compromises had been made and a balanced strategy achieved, the package routing commenced and the engineering teams were able to analyze the performance of the automotive-compliant package relative to the same silicon in a consumer-grade mobile package. The high level fundamentals were common between the mobile and the automotive SoCs, namely the use of the same silicon die, a multi-layer FCBGA package substrate with package decoupling, and fully integrated voltage regulators (FIVR) and embedded air core inductors (ACI). As discussed in the previous sections, the automotive SoC differed from the mobile by the addition of an IHS, the use of DSC instead of LSC, a thicker core, BLUF between the SoC and the PCB, and 2 additional package layers. The ACIs in the automotive SoC were slightly modified from mobile due to the thicker core, but they were still 1-2T racetrack and snake inductors designed to meet specific performance parameters. Figure 14 shows cross sections of this automotive PDN and a mobile PDN for reference.



Figure 14: Cross section of (a) the automotive PDN (b) the mobile PDN

FIVR, as the name implies, is a switching voltage regulator that is fully integrated into the SoC silicon and package substrate. The controller and switching FETs are integrated into the silicon side along with the bulk output decoupling capacitance while the inductive component of the regulator is implemented in the package substrate by means of routing fat traces configured in 'coils', or ACIs, to provide the necessary inductance. There are myriad of benefits to this scheme, the main ones being lower platform cost, granular power gating and control, and better transient performance [6]. A higher voltage input is fed into the balls of the SoC where it is then down regulated by the FIVR components and distributed to the necessary circuitry. Figure 15 shows a simplified schematic of a single FIVR domain as described in [7] and its partition at the die, package, and motherboard.



Figure 15: Simplified FIVR implementation

As discussed, the automotive design drivers degraded PI performance the most and necessitated PDN improvements to meet the silicon design targets. The PI analysis can be broken down into three sections; high current FIVR with input power, Low current FIVR with input power, and non-FIVR. The analysis of each will be presented along with the mobile SoC to provide a baseline for comparison.

5.1 Air Core Inductors (ACI)

All the ACIs were re-designed according to the automotive package stackup and form factor. Figure 16 shows an example of the IA Core ACI in an automotive vs. mobile space. As can be seen, the footprints are different, but all the automotive ACIs had similar ripple and compensator characteristics as the mobile design.



Figure 16: IA Core ACI in (a) mobile and (b) automotive

5.2 High Current FIVR with Input Power

The high current FIVR output rails do not have package caps, but instead require only that the ACI be designed to meet the specified performance parameters, such as the resistance and inductance at the switching frequency, Q factor, unity gain bandwidth, phase margin, and gain margin. The FIVR input rail, however, is highly sensitive to noise and consumes the majority of the on-package decoupling capacitors. The power supply rejection ratio is similar to that of discrete VR components, so the output noise of the FIVR is very much a function of the noise on the input side. The automotive FIVR input power network suffered heavily as a result of the compromises made. Even with all the optimizations, the performance of the automotive SoC was quite a bit worse than the mobile. Figure 17 shows the

PDN impedance of the input power on automotive vs mobile as a function of frequency. The 1st droop frequency was quite a bit lower and featured a peak impedance that was 50% higher than mobile, both of these being indicative of a more inductive path from the die to the first stage of decoupling. In other words, this degraded impedance was a direct result of using low numbers of DSC that could not be placed as close to the die as desired. Figure 18 shows a relative transient performance with ~35% more droop and 120% more overshoot in the automotive SoC than the mobile.



Figure 17: PDN impedance profile for high current FIVR rail



5.3 Lower Current FIVR with Input Power

In contrast to the high current FIVR output rails, the lower current FIVR output rails typically do have package caps. The lower current FIVR rails typically power I/O blocks, which are typically distributed over a long, thin strip of silicon, and require localized decoupling to mitigate noise at all power inputs of the I/O. As with the high current FIVR, the input rail is also still sensitive to noise. The high current FIVR can create noise transients that are so high on the input rail that even after being filtered through the VRM, they would still cause spec violations. For this reason, the lower current FIVR domains are typically powered by their own isolated input power rail. This separate input power rail still has a requirement for package decoupling, but the requirement is less than that on the high current FIVR output domain. As was seen with the higher current case, the automotive SoC performed much worse than the mobile SoC, but the design was still able to meet the transient noise specification.



Figure 19: PDN impedance profile for low current FIVR rail

5.4 Non-FIVR Power Rails

In addition to the FIVR rails, the automotive and mobile SoCs also featured several more traditional non-FIVR power rails. These power rails are sourced from typical VRM components on the motherboard (though in the case of automotive, these were automotive graded components). Some of these rails also feature on-package decoupling, but these rails are in general less critical and have higher noise tolerances. The non-FIVR rails on the automotive lent themselves well to the adjustment to DSC and the design team was in most cases able to meet the performance envelope set by the mobile component, or were able to meet the allowable noise thresholds. The main reason this was the case with the non-FIVR rails was because similar numbers of DSC on automotive were applied as were LSC on mobile. For the FIVR input and output rails, this was not the case as the constraints discussed earlier precluded the ability to match the number of package caps 1:1 between automotive and mobile, not to mention farther placement from the SoC die. Figure 20(a) show an example of a non-FIVR LPDDR4 core power rail with comparable impedance between the automotive and mobile packages. Figure 20(b) depicts another non-FIVR example of a PCH I/O power rail, where the automotive package results in higher peak-to-peak noise than the mobile package, but still within the allowable level.



Figure 20: (a) PDN impedance for non-FIVR LPDDR4 rail (b) Transient simulation for non-FIVR PCH I/O rail

6.0 Signal Integrity Optimizations

Section 5.0 demonstrated that the power integrity solution was compromised but mostly able to meet the performance targets set forth by the SoC design team, except for the high current FIVR input power rail. Unfortunately, the same could not be said for the signal integrity. As discussed in section 3, the modifications to the PDN necessitated the implementation of various changes that, when summed together, had a very significant effect on the signaling and electrical performance of some of the I/O interfaces. Several of these interfaces were broken from an electrical margin perspective and it then

became the focus of the signal integrity team to pull out all the stops and maximize channel margins wherever possible. The major pillars of this strategy are outlined below.

6.1 Split Reference Mitigation

Section 3.2.2 discussed the dual referencing scheme that was implemented for LPDDR4 and the split on the power side reference that was a result of a non-optimized silicon bump placement. The SI team worked with the package layout team to minimize the power plane cut and to bring it as close as possible to the die bumps. The reason for this was to bring the discontinuity of the split as close to the die as possible to allow for the on die decoupling capacitance to provide the return path for current coming back on the POWER1 plane. Even after making these allowances, the effect of the split is still easily seen in the channel response. Figure 21 compares the mobile SoC, featuring single referenced stripline traces, with the automotive SoC, with all the split plane mitigation strategies employed. Even with all the mitigation in place, the NEXT and FEXT for a representative DQ bit were significantly worse in the automotive SoC than they were in mobile.



Figure 21: Effect of the split dual referencing scheme on crosstalk (a) NEXT (b) FEXT

6.2 Via Transition Optimization

The signal integrity detractors invoked by the PDN modifications were largely caused by an increase in crosstalk, and it has been seen in the past that vertical crosstalk through long vias can be a significant issue that is often overlooked [8]. Some simple studies on single ended and differential via crosstalk were carried out that demonstrated the effect of vertical crosstalk in function of spacing and with or without optimally placed return vias. Figure 22(a) shows a single ended via system with via diameter D and spacing S. Figure 23 shows how crosstalk can be impacted by as much as 14dB in the case of single ended signals being tightly packed with no adjacent return path. Even with an increase of signal to signal spacing, adding a return path via has significant impact on reducing crosstalk. Figure 22(b) shows a similar study was carried out for differential vias with pair spacing DS and pair to pair spacing S. It was found that the differential nature of the buses meant that crosstalk was less as shown in Figure 24. The team still provided nearby returns for the differential pair vias in the package, but it had less of an impact than on the single ended signals.



Figure 22: Via models (a) single ended w/ and w/o GND via, (b) differential w/o GND via



Figure 23: Impact of adding GND via on single ended crosstalk (a) NEXT (b) FEXT



Figure 24: Impact of differential victim-aggressor spacing on crosstalk (a) differential (b) mixed mode

To this end, the SIE team methodically investigated all vertical connections on the high risk busses. Post-layout extraction were utilized to identify worst case signals and focus was placed on normalizing the performance across all pins. This performance normalization was achieved through placement of return vias nearby to all signal transition points. Often times, this required shuffling signals slightly and dealing with power plane voiding. The additional power plane voids degraded the power integrity slightly, but the buyback it provided in terms of total system margin made the changes worthwhile. The SoC ballmap was also altered to provide good isolation and crosstalk performance for these sensitive signals. For example, the final LPDDR4 bus featured a 1:1 S:G ratio for the DQ bits and also employed a full column of ground balls to isolate between adjacent bytelanes as seen in Figure 25.



Figure 25: BGA isolation between bytelanes on LPDDR4 interface

6.3 Signal Separation and ROV

The final pillar of the SI optimization strategy was to employ good design practices and make the most effective use of available space. The ground flood that typically fills in the gaps between signals on routing layers was reduced slightly and the signal traces were spaced out as much as the design rules allowed while still meeting the SoC footprint requirements set by the product team. Even spacing opportunities of a few tens of microns were captured. The team also prioritized the highly critical LPDDR4 bus and set a firm rule that there was to be zero ROV allowed for this bus. This was challenging, as there were well over a dozen DSC voids to contend with.

6.4 Effect of Extended Temperature on Low Speed I/O

One surprising challenge that came forth was that some of the lower speed busses (like eMMC/SPI) were also impacted negatively by the automotive use case conditions. The challenge here came about due to the mechanism by which buffer impedance and slew is calibrated coupled with the extended temperature requirements from Table 1. During normal operation, the system is powered on and the low speed I/O busses are tuned and compensated for the process, voltage, and temperature conditions at bootup. The problem is that if the part is powered up at -40°C and the part then later drifts up to the Tjmax of 110°C, the initial compensation no longer applies and the buffer is far weaker than desired. Under the standard mobile use case, this 85°C Δ T poses no issue, but the 150° Δ T experienced in automotive causes issues with channel timing. The preferred solution would have been to implement a dynamically updating compensation scheme, but such a scheme was not feasible in the timeframe required. Therefore, platform components had to be selected carefully that had low Tco and setup/hold requirements, trace spacing had to be increased, and the total allowable channel routing was reduced in order to maintain an operational bus.

7.0 Signal Integrity Results

The process to converge on a solution was an interactive one, and there were multiple back/forth cycles with the package and PCB layout teams to flesh out the most viable alternatives. As part of this process, to assess the impact on channel margin and performance, 3D models were created and BER contour channel simulations were run. The following subsections will discuss the key findings of the final iterations and outputs of this analysis effort.

7.1 LPDDR4

One of the most challenging buses on any platform is the DRAM interface. Large channel width coupled with high switching frequencies and multiple loads all create an environment that's very

challenging to maintain open eyes. 3D models of the automotive and mobile packages were extracted and a typical LPDDR4 channel was attached to the BGA balls of the 3D SoC extractions. Figure 26 compares a full channel with the mobile SoC versus the automotive SoC for DQ on the LPDDR4 interface, both with and without the SI optimizations discussed in section 5. Prior to optimization, increased crosstalk and severely degraded eye margins can be seen in the automotive case that have appreciable violations of the LPDDR4 mask requirements. The eye in this non-optimized automotive case has been closed by 34ps and 59mV and violates the eye mask. A small part of the degradation was caused by increased crosstalk due to the thick core, but most of the degradation came as a result of the PI requirements that were driven by the automotive use condition requirements. The dual referencing, split reference plane, vertical crosstalk, ROV, and constrained routing all contributed to this eye closure.



Figure 26: Frequency and time domain analysis of single reference and dual reference design (a) NEXT (b) FEXT (c) 10e-16 BER eye contour

Full channel simulation was also performed to ensure that there were no bits that were stragglers in terms of performance. Even just a single underperforming bit could cause overall speed degradation in the LPDDR4 bus as a whole. Figure 27 demonstrates how the eye width (EW) and eye height (EH) were improved as the optimizations detailed in section 5 were implemented in the package and the PCB.



Figure 27: Eye margin across full channel width (a) eye height (b) eye width

7.2 High Speed I/O (HSIO)

HSIO buses typically operate at much faster switching speeds than parallel busses like LPDDR4, but the differential nature of these interconnects and the feature set they implement (namely equalization) make them less problematic and more impervious (though not immune) to ROV and split reference routing. Table 3 shows the major differences between the mobile and automotive package design that posed potential signal integrity problem areas. Automotive requirements forced the use of dual reference on the package, there were unavoidable ROV, and the interfaces had to contend with larger via diameter. To minimize the impact on SI, a close coupling differential pair design was used. This approach minimized the effect of impedance change due to dual reference, ROV, and larger via, and also had the effect of minimizing crosstalk noise coupling. The only drawbacks was that without paying attention to the phase relation within the pair, common mode current could be generated in the differential pair that is then coupled to other pair, producing more jitter.

	Mobile	Automotive
PCIE/USB_TX	SL (3f)	SL (3f) dual reference
PCIE/USB_RX	MSL (surface)	SL (5f) over cap void
Core	X	3.5X
via diameter	small	large

Table 3: Mobile vs. automotive SoC package differences

7.2.1 PCI-Express Gen3

As with the LPDDR4 analysis, 3D modeling and channel simulations were performed to ensure that the optimizations made were sufficient to meet the requirements and operating speeds specified by the product team. Figure 28 shows a comparison between the mobile and automotive designs for a PCIe gen3 interface with a standard calibration channel (22" trace length) with and without RX DFE equalization. TX preset P1 with 0dB preshoot and -3.5dB de-emphasis was used. This test was selected to study the effect of the SoC package alone on overall performance. Further study was done using short and long calibration channels with different TX presets and optimized RX DFE. Eye margins are seen to be healthy and do not preclude full speed operation of the bus. Figure 29 shows a 6 lane PCIE

comparison of eye height and eye width between mobile and automotive designs after all optimization. The figure shows a comparable results with mobile design having higher margin over automotive.



Figure 28: PCIe gen3 eye diagrams for mobile (a) without RX DFE (b) with RX DFE; automotive (c) without RX DFE, (d) with RX DFE



Figure 29: PCIe gen3 eye metrics (a) eye height w/o RX DFE, (b) eye width w/o RX DFE, (c) eye height w RX DFE, (d) eye width w/ RX DFE

7.2.2 USB3.1

USB3.1 was another bus that suffered from degradation from the automotive use case conditions. Figure 30 shows a comparison in eye height between the automotive and mobile SoCs for 6 different USB differential pair. This analysis featured 9" PCB routing, a standard Type A connector, and a 1m cable assembly. While the automotive case had up to 30mV less margin, the eye openings were easily sufficient to pass the spec. Eye width was only degraded a small handful of picoseconds.



Figure 30: USB3.1 eye metrics (a) eye height (b) eye width

8.0 Summary

Designing high performance compute platforms for use in harsh conditions such as automotive presents unique and often surprising challenges to the SI and PI teams involved in product development. Strict automotive use conditions drive sub-optimal design decisions that force SI and PI teams to work collaboratively and to balance/optimize performance from a holistic, full system perspective. Even with a tight coupling between signal and power integrity, it is very difficult to design a system that reaches full performance parity with a platform that operates under less exacting conditions. This paper presented an SoC package development case study where the automotive use conditions created issues that were mitigated as best as they could without increasing size/cost beyond the point where the product became infeasible. The automotive performance did not come for free, however, and the product paid the price in terms of design complexity, engineering effort, and a slightly more expensive package substrate. Performance degradation was seen from a purely electrical sense, but was managed through careful and considerate balancing of the tradeoffs between signal and power integrity. For the end product, the degradation realized did not have any appreciable effect in reducing operating speeds or reducing feature set.

References

- Y. F. Shen, "Power integrity challenges of re-designing a mobile SoC with fully integrated voltage regulator to IoT applications," 2018 IEEE 22nd Workshop on Signal and Power Integrity (SPI), 2018.
- [2] Automotive Electronics Council Component Technical Committee, "AEC-Q200 Rev-D," June 2010.
- [3] Automotive Electronics Council Component Technical Committee, "AEC-Q100 Rev-H," September 2014.
- [4] S. Tandon, "Package reliability challenges in autonomous systems," *Fraunhofer EMFT & Tech Search Workshop*, pp 19, November 2017.
- [5] A. Waizman and C. Y. Chung, "Resonant free power network design using extended adaptive voltage positioning (EAVP) methodology," *IEEE Transactions on Advanced Packaging*, vol. 24, pp. 236-244, August 2001.
- [6] E. A. Burton, et al., "FIVR Fully Integrated Voltage Regulators on 4th Generation Intel® Core[™] SoCs," *IEEE Applied Power Electronics Conference and Exposition (APEC) 2014 Twenty-Ninth Annual IEEE*, pp. 432-439, March 2014.
- [7] K. Bharath and S. Venkatraman. "Power Deliver Design and Analysis of 14nm Multicore Server CPUs with Integrated Voltage Regulators," *Proc. of the 2016 IEEE Electronic Components and Technology Conference (ECTC).*
- [8] M. White, et al., "Cost Effective Crosstalk Management & Decision Making," *DesignCon 2010*, pp. 502-520, February 2010.