

DesignCon 2009

Analyzing Signal and Power Integrity Limitations for Mobile Memory Systems in 3D Packaging Environments

Ralf Schmitt, Rambus Inc.

[Email: rschmitt@rambus.com]

Joong-Ho Kim, Rambus Inc.

June Feng, Rambus Inc.

Dan Oh, Rambus Inc.

Chuck Yuan, Rambus Inc.

Abstract

Mobile memory systems, optimized for low power dissipation, use advanced 3D packaging strategies like 'Package-on-Package (PoP)' and stacked memory devices to achieve high system density. As the bandwidth requirements of mobile systems increases, data rate scaling is limited by SI and PI effects. In this paper, we analyze a typical mobile memory system to identify the dominating Si and PI effects limiting data rate scaling in these systems. We present a model of the 3D system, analyze the impact of SI and PI effects on system performance, and discuss design improvements to increase the achievable data rate.

Authors Biography

Ralf Schmitt is Senior Engineering Manager at Rambus Inc. He received his Ph.D. in Electrical Engineering from the Technical University of Berlin, Germany. Since 2002, he is with Rambus Inc, Los Altos, California, where he is a senior engineering manager responsible for power integrity on chip, package, and system level. His professional interests include on-chip signal integrity, power integrity, timing analysis, clock distribution, and high-speed digital circuit design.

Joong-Ho Kim received his M.S. and Ph. D. degrees in electrical and computer engineering from Georgia Institute of Technology, Atlanta, Georgia. He is currently a signal integrity engineer at Rambus Inc. He is responsible for product design and analysis in signal/power integrity (SI/PI) area for high performance memory interface products such as XDR, DDR, GDDR, and low power memory systems. Previously, he was working on SI/PI analysis for CMOS microprocessors and in-house tool developments at Intel Corporation. Dr. Kim has thirty five publications in refereed journals and conferences, eight issued patents. His current research interests are systematic approach considering both signal and power integrity simultaneously, the characterization of high-speed interconnects, S-parameter based simulation using VNA measurements or full-wave solvers, and the macro-modeling for circuit simulations.

June Feng received her MS from University of California at Davis, and BS from Beijing University in China. From 1998 to 2000, She was with Amkor Technology, Chandler, AZ. She was responsible for BGA package substrate modeling and design and PCB characterization. In 2000, she joined Rambus Inc and is currently a senior member of technical staff. She is in charge of performing detailed analysis, modeling, design and characterization in a variety of areas including high-speed, low cost PCB layout and device packaging Her interests include high-speed interconnects modeling, channel VT budget simulation, power delivery network modeling and high-frequency measurements.

Kyung Suk (Dan) Oh is Senior Engineering Manager at Rambus Inc. He received the B.S., M.S., and Ph.D. in electrical engineering from the University of Illinois, at Urbana-Champaign in 1990, 1992, and 1995, respectively. His doctoral research was in the area of computational electromagnetics applied to transmission line modeling and simulation. Since 2000, he is with Rambus Inc., Los Altos, CA. His group is responsible for providing signal integrity analysis for various products including serial, parallel, and memory interfaces. Additional responsibilities of his group include an advance CAD tool development for high-speed link simulation. His current interests include advance signal and power integrity modeling and simulation techniques, optimization of channel designs for various standard or proprietary I/O links, and application of signaling techniques to high speed digital links.

Xingchao (Chuck) Yuan received his M.S. and Ph.D. degrees in Electrical Engineering from Syracuse University, in 1983 and 1987, respectively, Syracuse, New York. Since 1998, he is with Rambus Inc, Los Altos, California, where is a director of signal integrity engineering responsible for designing, modeling, and implementing Rambus multi-gigahertz signaling technologies using conventional interconnect technologies. He has over seventy publications in technical journals and conferences. His professional interest includes computational electromagnetic methods and high speed signaling technologies for multi-gigabit applications.

I. Introduction

Mobile devices like cell phones and portable game consoles require memory systems that meet two key requirements: low power dissipation and high packaging density. In order to meet these requirements, mobile memories are implementing changes compared to traditional memories that impact signal and power integrity.

Mobile memories achieve low power dissipation partially due to changes in the memory core itself, e.g. reducing the power necessary for refresh operations, and partially due to changes to the interface circuits. In particular, static power dissipation in the interface is reduced or removed by increasing the termination impedance of the memory bus. High packaging density is achieved by extensive use of 3D packaging. For example, memory dies are stacked in the same package to increase memory capacity for a given footprint. The consequence is an increased bus loading and larger crosstalk and increased supply impedance due to longer bond wires. Another 3D packaging option popular in mobile systems are 'Package-on-Package (PoP)' structures, where the memory device package is stacked directly on top of the controller package. The interface between the control and the memory devices in a PoP structure is in general pin-limited, and the power supply system in PoP systems has to be routed through a stack of two packages, both increasing the power supply network of the memory device. Thus, 3D die staking and packaging options introduce additional SI and PI challenges for mobile systems that can limit the achievable data rate of these systems.

Current mobile memories like LPDDR are targeting data rates of 266MHz and below, and until recently, such data rates were sufficient for mobile devices like cell phones. Lately, however, as mobile devices are increasingly used to generate and access multimedia data, the memory bandwidth and capacity requirements have increased dramatically, raising the need to increase the data rate of mobile memory significantly. Consequently, the next generation of mobile memories will target data rates of 667MHz and above.

In this paper we present an analysis of the SI and PI challenges that limit the data rate scaling of mobile memory systems derived from current specifications. We review the current specification for mobile memory solutions and identify design characteristics that limit the scaling of these memory interfaces to higher data rates.

Next we analyze the performance of the mobile memory interface in a channel and supply environment typical for high-performance mobile systems, using stacked memory devices. The impact of ISI, crosstalk, and SSO noise on the data eye opening is analyzed and the dominating effects limiting the achievable data rate are identified. Changes to the specification of the mobile memory system are investigated that improve the data rate scaling but do not have a significant impact on the power dissipation.

Finally, the maximum achievable data rate of the mobile system is discussed depending on the packaging environments used in the mobile system.

Our analysis results show that SSO noise and crosstalk are the two dominating factors limiting the achievable data rate in current mobile memory systems. 800Mbps is feasible with some challenges. But 1066Mbps appears to be very challenging. Scaling the data rate of single-ended un-terminated (or lightly terminated) interfaces beyond 1066MHz

will likely require to trade off power dissipation against performance, removing the power saving features from the mobile memory design, which is not attractive. A low power memory system based on differential signaling is required to satisfy the memory bandwidth requirement of future multimedia mobile devices.

II. Characteristics of Current Mobile Memory Systems

In mobile systems maximizing battery time is a major design target. Every component of a mobile system is designed for low power dissipation to meet the tight system power target and achieve an acceptable battery life. Therefore, special memories are developed for mobile applications that show significantly lower power dissipation than standard memory devices.

The primary target for power reduction in mobile memories is the stand-by power of the memory devices. Many components of a mobile system like a smart phone are in stand-by mode for most of the time, and reducing the power dissipation in this mode is crucial for battery time. The second target for power reduction in mobile memories is the power dissipation during active memory access. Existing mobile memory specifications address both of these power reduction targets.

II.1. Characteristics of Mobile Memory LPDDR Device

In 2007, JEDEC released a specification for a “Low Power Double Data Rate (LPDDR) SDRAM” device [1] that is largely based on the existing specification for ‘standard’ the DDR-I and DDR-2 memory devices, but implements changes compared to these specifications that reduce the power dissipation as well as the pin count. Several of these changes address the memory core itself and have no direct impact on the signal integrity of the interface circuits. Examples of these changes are:

- A Power-Down and a Deep-Power-Down mode that deactivate many internal circuits to reduce stand-by current.
- Partial Array Self Refresh (PASR) that restricts the self-refresh to a part of the memory array. This avoids the power-consuming refresh of memory array regions that are currently not used.
- Temperature Compensated Self Refresh (TCSR). Since the leakage of memory cells is a strong function of temperature, then self-refresh rate can be adjusted based on the device temperature.

The LPDDR devices described in this specification target data rates up to 266 Mbps, much lower than e.g. DDR-2 devices. Furthermore, mobile memories are primarily used at very short channels with only few device loadings, i.e. few memory devices, on the channel. As a result, the requirements for signal integrity on these channels are by far not as strict as for ‘standard’ memories with higher data rate. Therefore, the interface circuits can be simplified as well, reducing the power dissipation of the interface circuits for the price of a margin loss that is tolerable at these data rates and channel topologies. These changes are:

- **No bus termination, in particular no On-Die Termination (ODT)**

Terminating the channel driven by a push-pull driver, as used in DDR memories, will cause static current draw in at least one of the logical state. Omitting the termination of the bus creates reflections at the end of the channel that increase the voltage swing at the receiver, but causes Inter-Signal-Interference (ISI) due to the reflected wave.

- **No external reference voltage V_{ref} for the receivers**

A reference voltage for the receivers is typically generated using resistive voltage divider. Voltage dividers, however, dissipate static power. Omitting a reference voltage for the receivers reduces the static power dissipation of the memory system, but leads to a single-ended receiver design, i.e. a receiver similar to a simple inverter. Such a receiver has typically smaller input voltage sensitivity than differential receiver designs. But this was tolerated due to the large input swing on the un-terminated bus and the low target data rate.

- **High ratio of power supply to signal pads, especially on CA bus**

For the data bus, the ratio between VDDQ/VSSQ supply pads and DQ signal pads is approximately 2, which is comparable to standard memory devices. On the Control and Address bus, however, only a single pair of VDD/VSS pads supports up to 24 input receivers. This very large ratio between power supply pads and signal pads is only possible due to the lack of on-die termination. Thus, no termination current has to be supported at the receivers.

Figure 1 shows a simple schematic of the DQ and CA bus for an LPDDR memory system. It consists of a full-rail voltage-mode driver connected to an un-terminated bus with a single-ended receiver.

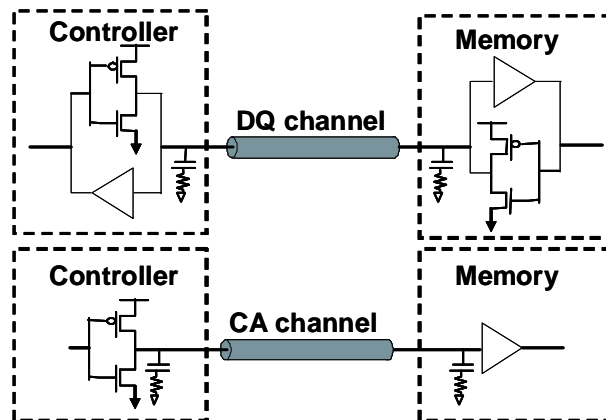


Figure 1: Schematic of LPDDR memory channel

11.2. Characteristics of Mobile Memory Channels

Mobile systems are usually designed in a small form factor, e.g. a handheld device. This requires a very compact system design with components placed as close to each other as possible. In order to accommodate such a compact placement, mobile systems often use PCBs generated in a build-up process with numerous signal layers that supports

stripline routing and via structures with fine pitch. As a result, the channel length in mobile systems is usually very short, often less than 3”.

In order to achieve a compact design mobile systems also make extensive use of chip stacks, where several dies are stacked on top of each other. There are several solutions available to generate such device stacks using memory dies. Some of these solutions use unpackaged bare dies that are integrated directly into the system design, other solutions use memory die stacks in a standard package. In the scope of this analysis we will focus on memory die stacks in standard packages. These stacked memory devices can be integrated into the system as a single device.

Figure 2 shows two different memory device stacks frequently used in mobile systems.




Figure 2: Example of Memory Devices with Stacked Dies

In a homogeneous stack several memory dies of the same kind are stacked to generate a memory with larger capacity. The corresponding signal buses of the different dies, with the exception of Chip Select signals, are connected to the same bonding pad. This generates two loadings on the channel separated by the bond wires as short stubs.

In a heterogeneous stack memory dies of different kind are stacked in the same package. A typical combination of memory dies is a flash memory with large capacity but poor WRITE performance, and a DRAM with smaller capacity but much faster access time. Typically, each of the memory device kinds uses its own set of buses and pins in the package. For the channel analysis, each of the two dies can be analyzed separately.

Commonly, hybrid solutions are used that combine e.g. a single Flash die with two or more DRAM dies. These memory stacks provide large flexibility to the system design. Dependent on the type and number of dies in the memory stack the system memory configuration can be changed with minimum impact on the system design. So, it is possible to offer systems with different memory capacity by changing only a single component.

11.3. Package-on-Package (PoP) Configuration

Package-on-Package (PoP) configurations are a convenient way to combine the packing density typical for vertical device stacks with the use of standardized memory devices that can be provided by multiple vendors. Figure 3 shows a PoP combination of a controller chip and a stacked memory device. On the left, both devices are shown separately, each with their own package. On the right, the combination is shown after placing the memory device on top of the controller package.

In a PoP configuration, the memory device, which by itself can be a stack of multiple memory dies, is placed on top of the controller package. A typical configuration uses a hybrid memory stack of Flash and DRAM memory that provide all memory needed by the controller. The controller package provides pads for the BGA balls of the memory device on the top side of the controller package. In order to provide space for the controller chip itself, the BGA balls in the center of the memory component have to be de-populated.




Figure 3: Controller and Memory Device in PoP Configuration

There are standards for PoP memory components defining the size and the pad-out for different bus width and different memory die stacks. Standard PoP memory components provide a ring of BGA balls two rows deep at the periphery of the device package. As a result, memory components in PoP packages are usually pin limited, especially for hybrid memory stacks with separate buses for e.g. Flash and DRAM dies. Figure 4 shows an example of a hybrid memory stack containing Flash and 32-bit LPDDR dies in a 128-pin PoP package. The LPDDR buses are located on the left (DQ bus) and right (CA bus), while the top and bottom side are largely reserved for the Flash interface.

Compared to a side-by-side placement of controller and memory component on a PCB, PoP configurations typically have a shorter channel length, consisting only of the routing inside the two packages. The power supply system of a PoP memory component, however, has typically a highly inductance than the supply system of a stand-alone memory component, since the power supply for the memory component in a PoP environment has to be routed through the controller package. Additionally, the interface between controller and memory component in a PoP environment is often pin limited, as can be seen in Figure 4. Therefore, the number of power supply pins is often smaller than for stand-alone packages, leading to a smaller number of highly inductive supply traces in the (typically) 2-layer PoP memory package.

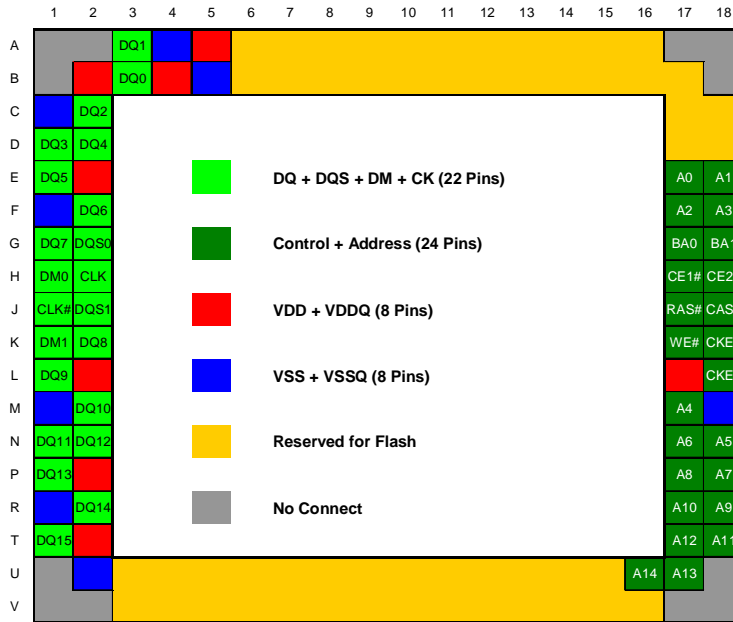


Figure 4: PoP Package Pin-Out Example

III. Mobile Memory System Model

In this section we will analyze the SI and PI effects limiting data rate scaling in current mobile memory systems. For this, we analyze a typical mobile memory system based on the LPDDR specification listed before, but incorporating some improvements that are likely necessary to achieve higher data rates while keeping the power dissipation low.

The mobile memory system under consideration consists of a stacked memory component with either one or two memory dies, placed on the PCB aside of the controller chip. We intentionally chose a PCB placement instead of a PoP placement to cover a wider range of typical system implementations. The PoP placement will be covered as a special case of this test system where the PCB channel length is negligible and the memory component package has an increased package inductance.

In contrast to the LPDDR specification, the output driver supply voltage of the memory system was reduced from $VDDQ=1.8V$ to $VDDQ=1.2V$. This reduction is already offered in some commercially available LPDDR device implementations and will likely be widely used in the future to further reduce the power dissipation of these devices.

In our memory system model we also assume a differential receiver with an external V_{ref} voltage, again in contrast to the LPDDR specification. Single-ended receivers have a much lower sensitivity and require a large V_{in} especially over different PVT corners. This would limit the data rate dramatically especially if the output driver voltage is reduced to $VDDQ=1.2V$. Therefore, we decided to simulate a differential receiver, but the results are valid for single-ended receiver as well, assuming a larger V_{in} requirement.

III.1. Memory Component Model

The mobile memory system we analyze uses a memory component with either one or two identical mobile memory dies, stacked in the same package as shown on the left side of Figure 2. For the case of a two-die stack, corresponding pads of both dies are connected to the same bonding pad inside the package with the exception of the individual chip-select signals of both dies. In our analysis we will simulate both a single memory die as well as a two-die stack configuration to investigate the impact of device loading on the channel performance.

We assume a 4-layer package for the memory component, providing supply power planes and microstrip routing for the signals. This is a high-end solution for a memory package, providing a best-case for the signal and power integrity of the memory package. In many cases, memory devices as the one we are using in our analysis will be assembled in two-layer CSP packages, resulting in higher crosstalk and higher package supply inductance. But once we have identified the SI and PI performance limitations of the mobile memory system we can estimate the impact of the 2-layer package on system performance.

The microstrip signal traces in the memory package were modeled using transmission line models. The bond wires of the memory device(s) were extracted using a 3D field solver. Figure 5 shows the bond wires of the two-die memory component. Due to the close vertical proximity significant crosstalk has to be expected between the bond wires of both devices. Therefore, the bond wires of both devices were extracted together to capture this crosstalk.




Figure 5: Wirebond section of two-die memory package (DQ bus)

III.2. Controller Model

Mobile memory controllers are typically assembled in high-performance flip-chip packages. Stripline routing is assumed for the signal routing inside the controller package, with the stripline impedance closely matching the output driver impedance. Power supply planes and a sufficient number of vias in the package provide a low-inductance power supply system for the controller.

III.3. PCB Model

Mobile systems often use PCBs in build-up technology with 8 or more layers, providing ample of signal layers for stripline routing and aggressive design rules. In our simulation model PCB signals are routed as striplines with a characteristic impedance of 50Ω . The typical length of the PCB routing in our model is 1.5”.

III.4. Driver Impedance

In our simulations we are assuming a nominal driver impedance of 40Ω . This is approximately equivalent to the half-strength mode in the LPDDR specification, which is intended for lighter loads and point-to-point environment. The mobile memory system we are analyzing provides a point-to-point environment even in the two-die memory stack configuration, since both loads are at the end of the channel with very short stubs between them.

III.5. Channel Model of Mobile Memory Test System

Figure 6 shows the channel model of the mobile test system for the two-die memory stack configuration. It follows the same modeling strategy as shown in [2]. For the single-die configurations one of the memory dies together with its bond wires are removed.

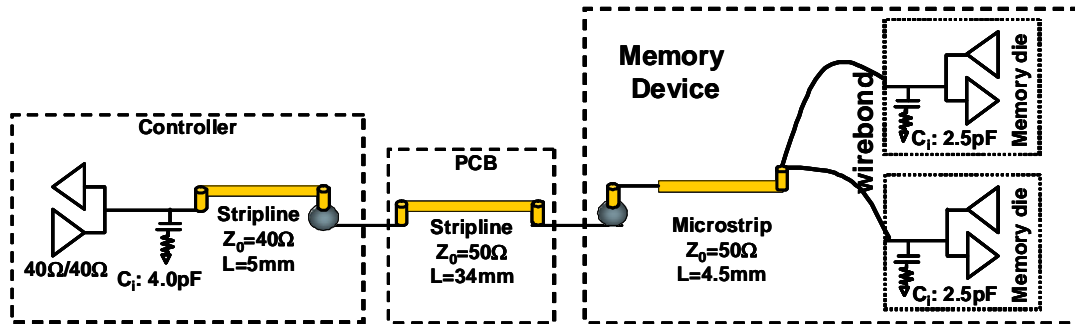


Figure 6: Channel Model of Mobile Memory Test System

III.6. SI/PI Co-Simulation Model of Mobile Memory System

In order to analyze the impact of signal as well as power integrity issues on the channel performance, a SI/PI co-simulation deck was derived from the component models listed above. The modeling methodology follows the same strategy as described in [3] and [4]. A schematic of the SI/PI modeling methodology is shown in Figure 7.

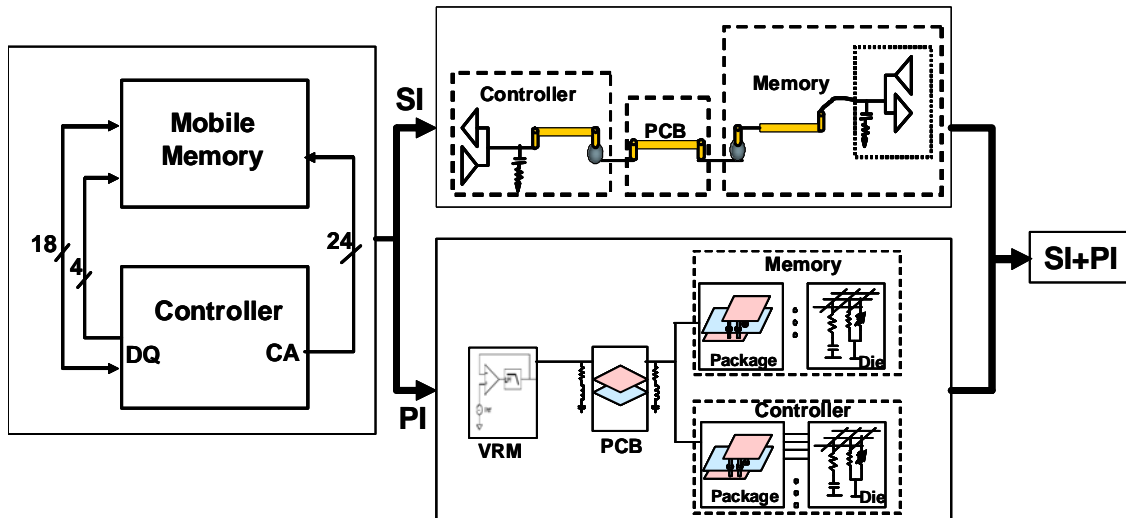


Figure 7: Schematic of SI/PI Model Generation Strategy

First, a compact SI models is developed for the memory systems. This model can be used stand-alone to analyze SI effects like crosstalk and ISI. In parallel, a detailed power supply model is generated modeling each component of the power delivery network. From this detailed supply model a reduced-order supply model is derived with greatly reduced complexity. This compact power supply model is combined with the channel model to generate a SI/PI co-simulation deck. With this co-simulation deck the impact of supply noise in general and SSO in particular on the channel can be simulated.

For the Mobile Memory Test System, the generation of a SI/PI co-simulation deck was further divided into a DQ model and a CA model. Figure 8 shows a block diagram of the resulting model.

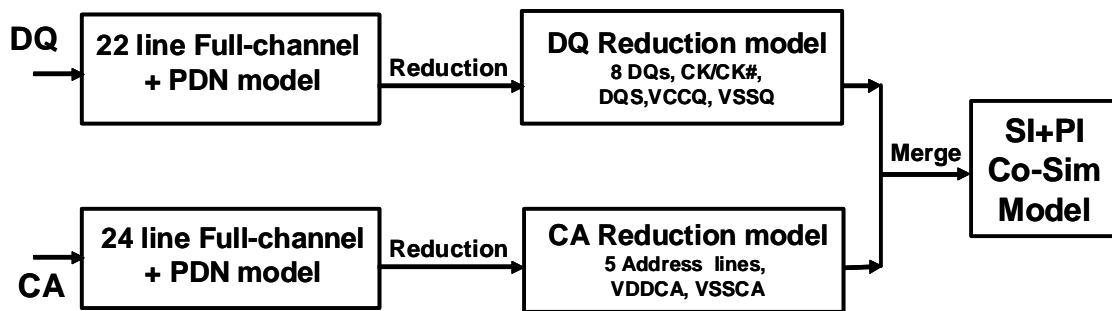


Figure 8: Block Diagram of Mobile Memory Channel Co-Simulation Deck

Since the CA and the DQ bus in the Mobile Memory device are entirely separated, even using pads on opposite sides of the die, the signal and power supply models of both channels can be separated as well. The final model, however, still contains models for both buses together so that the transfer of supply noise from one side to the other side can be modeled.

IV. Mobile Memory SI Analysis Results

IV.1. Channel ISI and Crosstalk

In a first set of simulations, the impact of Inter Signal Interference (ISI) and crosstalk on the eye opening of the data channel was investigated. For the ISI simulation, only the victim signal is transmitting while all other signals are kept quiet. For crosstalk simulations, all signals are transmitting data, however, the data pattern transmitted by the aggressor lines is different from the data transmitted on the victim line. All simulations use an ideal power supply for the transmitter and the receiver of the interface system. In order to achieve the best possible performance in this system, a nearly optimal PCB routing was assumed with minimum crosstalk. PCB signals were routed as striplines with a large spacing of $3w$ (three times the width of the signal line itself) between adjacent lines.

The eye opening at the receiver was simulated for Read and for Write accesses using a data rate of 667MHz. Simulations were done for two different memory configurations: in the first configuration, a standard memory device with a single memory die was used. The second configuration used a memory device where two memory dies are stacked on top of each other, each of the signals connected in parallel to the same package pad. Figure 9 shows the eye diagram results.

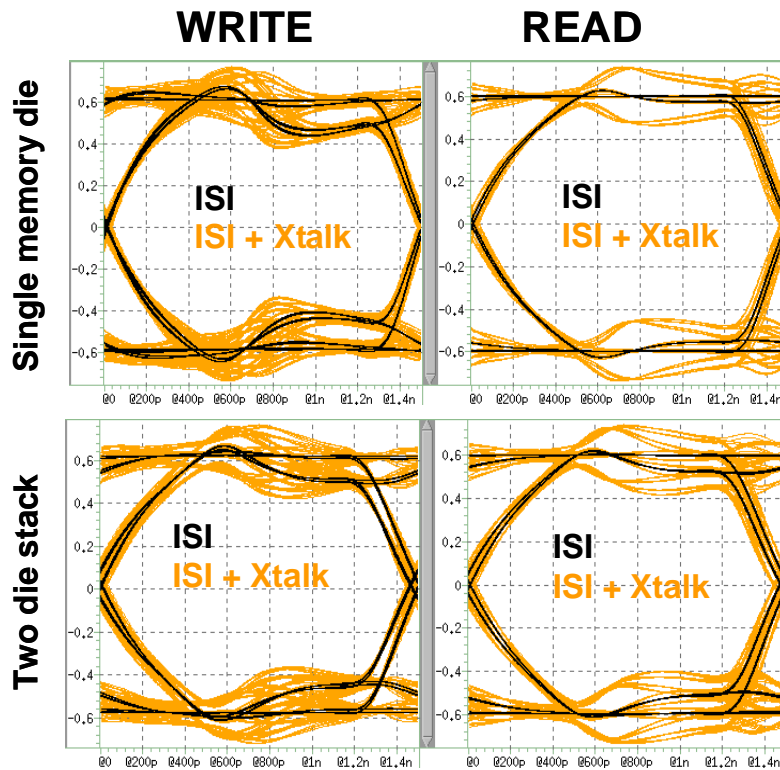


Figure 9: Eye diagrams at 667MHz for ISI and Crosstalk

Analyzing only the effect of ISI, both Read and Write case show a wide open data eye. The reflection at the un-terminated end of the channel cause some margin loss, visible in the second half of the data eye especially for the Write case, but this margin loss is tolerable.

Both memory configurations, the case of a single memory die and the case of a stack of two memory dies in parallel, achieve comparable eye openings, although the capacitive loading for the configuration with the memory stack is higher. This is caused by the missing termination of the channel. The capacitive loading reduces the impedance at the end of the channel, bringing it closer to the characteristic impedance of the channel.

Considering both ISI and crosstalk, the size of the data eye is decreasing, but it still appears acceptable. The eye height for Read and Write case, for a single die load and for a two die memory stack, are all comparable. Breaking down the crosstalk contributions of different channel components shows that the wirebond section of the package dominates the crosstalk, contributing approximately 50% to the total crosstalk, while PCB routing and the routing in the memory package are each contributing approximately 25% to the total crosstalk. This is the result of the very optimistic routing assumptions in the packages and the PCB. In the controller package, short stripline routing is used, and in the PCB stripines with large signal spacing ($3w$) were used. Finally, the memory uses a 4-layer package with stripline routing with large ($2w$) spacing. These are optimistic assumptions that might not be achievable in a real system, so the results in Figure 9 represent the best-case for margin loss due to crosstalk. In an actual system memory devices often use two-layer packages, especially in PoP environments that will show significantly increased crosstalk between the package traces. Also, providing a spacing of $3w$ for the PCB routing might not be possible in many cases. In order to analyze the margin sensitivity to this routing spacing we simulated the eye opening in Write mode for different PCB wire spacing. Figure 10 shows the result for a data rate of 667MHz.

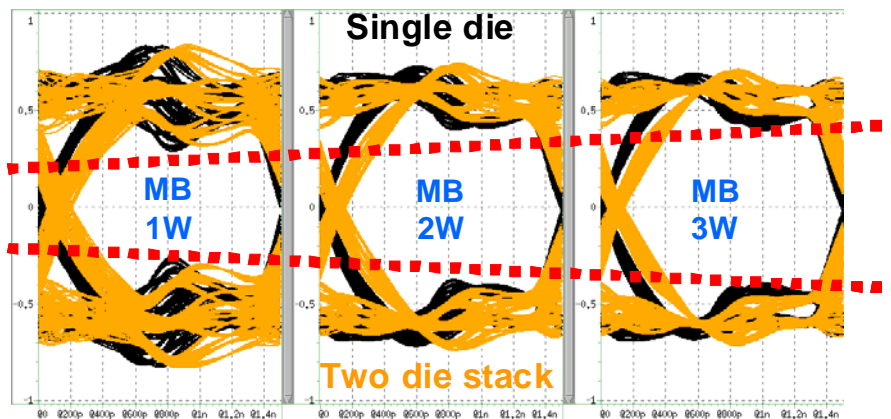


Figure 10: Write Eye Diagrams for Different PCB Routing Spacing

The results shown in Figure 10 show a large sensitivity to the spacing used in the PCB routing. For a narrow the voltage margin is reduced substantially and jitter is increased due to increased crosstalk. The same can be expected for increased crosstalk in the memory package.

IV.2. Power Supply Noise Analysis

For most systems the power supply noise inside the memory device will be higher than at the controller PHY, since controller flip-chip packages typically allow power supply designs with significantly lower inductance. For power supply analysis we therefore focus our attention on the memory device. Figure 11 shows the schematic of the power supply system for the example of a memory component with two stacked memory dies.

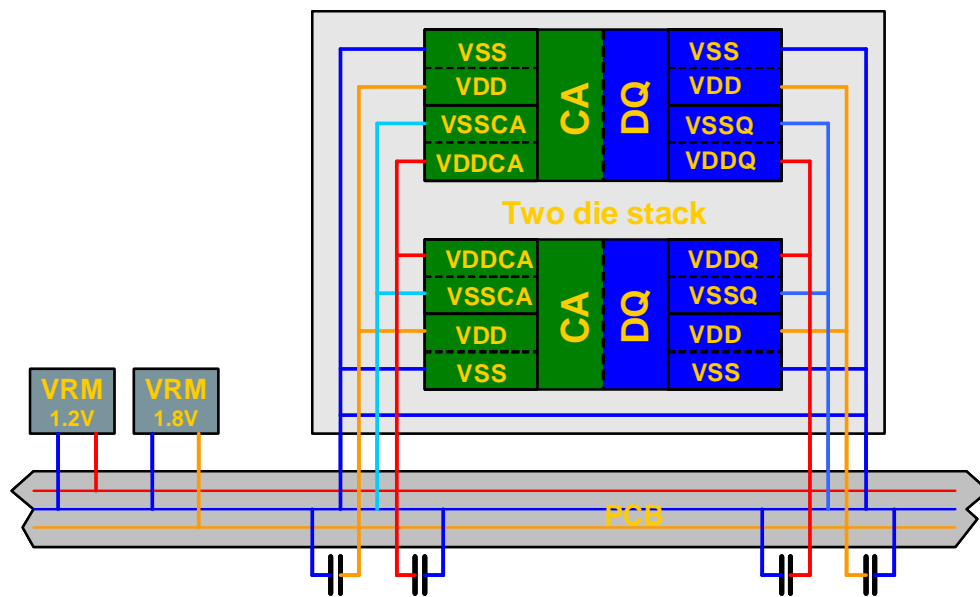


Figure 11: Power Supply Schematic of Two Die Memory Component

The power supply pairs for the DQ channels and the CA channel are separated to two sides of the memory sides, named VDDQ/VSSQ and VDDCA/VSSCA respectively. The ground rails, although named differently in the schematic, can be shared throughout the package and even silicon routing. Decoupling capacitors will be typically placed on the back side of the PCB to avoid interference with signal escape routing.

For each of the two output supply pairs we are interested in self-induced noise, i.e. the noise generated by the current draw in this supply pair itself, as well as transfer noise, i.e. the amount of noise coupled from one supply pair to the other [5]. Figure 12 shows the self-impedance and the transfer-impedance of the interface supply system.

During Write access, both the DQ and the CA signals are receiving data. Due to the lack of termination on all signal lines, current draw on VDDQ and VDDCA will be minimal in this case, and so we expect only little supply noise.




Figure 12: Self- and Transfer Impedances of Memory Supply Systems

During Read access, however, the DQ drivers inside the memory are transmitting, and VDDQ has to supply the output drivers, while the CA bus is still receiving. Consequently, the VDDQ power supply system requires lower supply impedance than the VDDCA supply system, which is reflected in Figure 12. A critical concern during Read is how much of the noise generated by the DQ output drivers is coupled to the CA receivers. Figure 11 shows, however, that the corresponding transfer impedance is very low. This is a result of the separation between VDDQ and VDDCA all the way through the package.

It is interesting to notice that the peak supply impedances for a stack of two dies are in general lower than the impedance for a single die device. In a two die stack, the on-chip decoupling capacitors of both dies are shorted together inside the package, providing a larger amount of capacitance to the die that is active at any given moment. Therefore, it can be expected that the supply noise for devices with several stacked dies is lower than the supply for a single die device.

Figure 13 shows the simulated power supply noise inside the memory package during Read and Write access. As expected, there is little power supply noise during Write access, since no output drivers are active inside the memory. During Read access, VDDQ shows significant supply noise due to SSO. A part of this noise is coupled into VDDCA, but the noise amplitude on VDDCA is still significantly lower than on VDDQ. As expected, supply noise for two die device stack is smaller than for a single die due to the shared on-chip decoupling capacitors in the two die configuration.

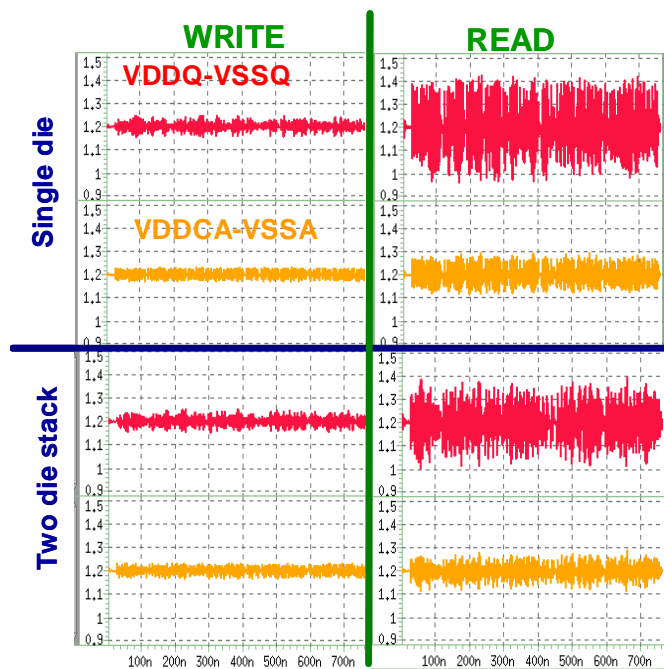


Figure 13: Supply Noise Waveforms inside the Memory Device

IV.3. SSO Analysis

Figure 14 shows the eye diagram including the effect of SSO supply noise.

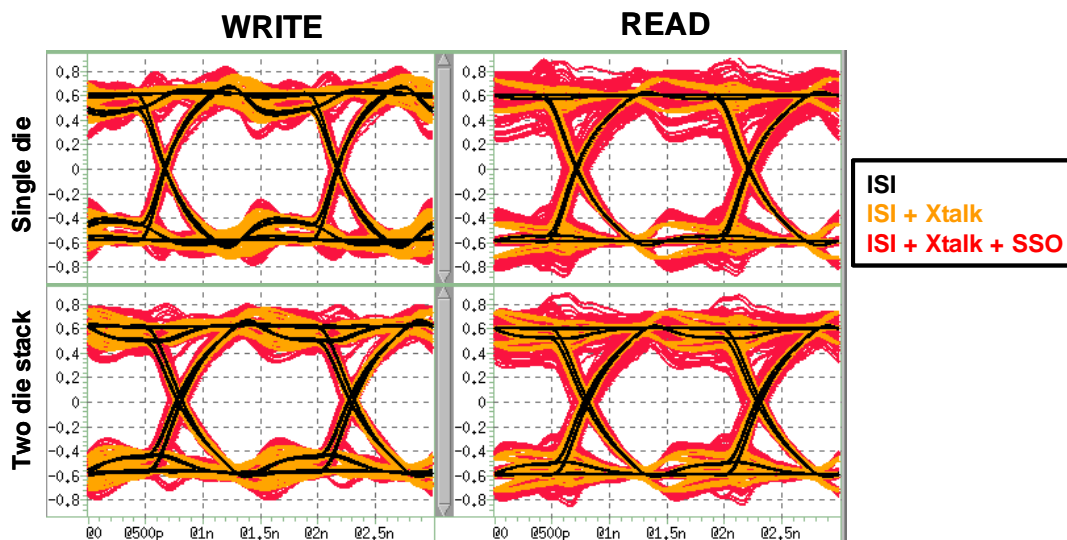


Figure 14: Eye Diagrams including SSO Noise Impact

Figure 14 shows that during Read access, when the output drivers in the memory device are transmitting, there is a significant voltage and timing margin loss due to SSO

noise. During Write access, the additional loss due to SSO is significantly smaller, reflecting the smaller supply noise expected at the controller.

IV.4. Margin Loss Summary

In order to compare the relative eye height for different simulation cases, the minimum eye height in a window of $\pm 1/4$ of the bit time around the eye center is measured. Additionally, the jitter at the reference voltage level is measured to compare the timing jitter of different simulation results. Figure 15 shows a schematic of the measurements.

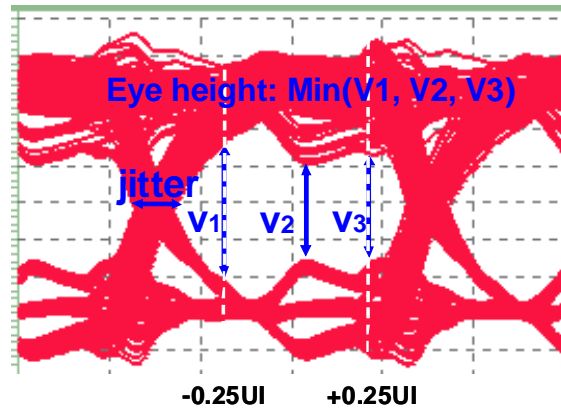


Figure 15: Measurements to determine Eye Height and Jitter

Figure 16 shows the margin parameters considering ISI, crosstalk, and SSO impacts for a data rate of 667MHz.

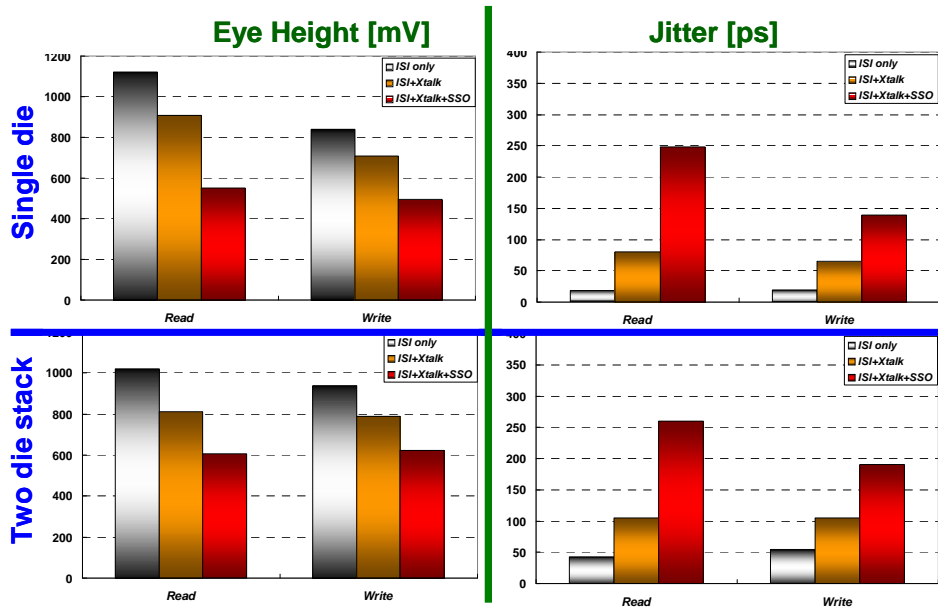


Figure 16: Eye Margin Breakdown at 667MHz Data Rate

Figure 16 shows that SSO has a large impact on signal jitter, especially for Read access. In this mode, SSO contributes for more than half of the total eye jitter. During Write access the jitter is smaller, due to the lower supply noise expected in the controller, but SSO is still the dominant jitter source. SSO also has a large impact on eye height, as large as or even larger than crosstalk.

Figure 17 shows the margin breakdown for a data rate of 800MHz




Figure 17: Eye Margin Breakdown at 800MHz Data Rate

Figure 17 shows that at higher data rate SSO increasingly becomes the dominant contributor to margin loss, both for jitter as well as eye height. It also shows that at higher data rates the configuration with two memory dies achieve lower margin than a configuration with a single die, although at 800MHz this difference is still small. For a configuration with two dies the eye height is reduced to 400mV and the signal jitter consumes more than 25% of the bit time.

V. Summary and Conclusions

In this paper we have analyzed the performance limits of current mobile memory systems which are based on push-pull voltage mode drivers on un-terminated channels. We reviewed typical system implementations for mobile system and discussed 3D packaging options like die stacking and PoP often used in mobile systems. Next, we generated a SI/PI co-simulation model of a typical mobile system and analyzed the impact of ISI, crosstalk and SSO on the data eye margin.

Based on the analysis results, SSO noise and crosstalk are the two dominating factors limiting the achievable data rate in current mobile memory systems. The increased channel loading caused by two stacked memory chips instead of a single chip, on the other side, has only minor impact on the achievable performance.

Our analysis suggests that a data rate of 667MHz should be achievable in realistic mobile system environments, including a PoP environment. However, it likely requires a differential receiver instead of the single-ended receiver proposed in the current mobile memory specification. A typical DDR system, which uses a differential receiver, requires an eye opening of approximately 400mV to reliably detect a data signal. The analyzed mobile memory achieves this eye opening at 667MHz. A single-ended receiver, however, would likely require a larger eye opening, especially when considering parameter variations.

A data rate of 800MHz appears possible, however it will require that the crosstalk in PCB and package is minimized and the power supply impedance of the memory package is sufficiently small, similar to the assumptions in our analysis. For PoP environments this will be difficult to achieve, since the supply impedance is typically too high. Taking component variations typical for high volume production into consideration, a yield loss is likely at this data rate even for a high-performance mobile system environment.

Achieving a data rate of 1066MHz appears challenging and will likely require improvements in the design of the memory component. Possible improvements are

- **A noticeable increase in the ration of supply pads to signal pads**

Increasing the number of supply pads will improve the power supply impedance, reducing supply noise during SSO events, and can be used to improve the shielding between signal wires. This reduces the crosstalk in the bond wires and package routing.

- **‘Light’ termination of the signal channel**

Currently, the signal channels are un-terminated to avoid static current flow in each of the logical states. Reflections could be reduced by partially terminating the signal lines. Unfortunately, this will cause static current dissipation in the output drivers, which is not appreciated. But it might be possible to keep the termination impedance high enough, limiting the static current dissipation, and still achieve the necessary improvement in signal quality.

Scaling the data rate of single-ended un-terminated (or lightly terminated) interfaces beyond 1066MHz will likely require to trade off power dissipation against performance, removing the power saving features from the mobile memory design, which is not attractive.

Differential signaling provides a much more promising solution for a data rate scaling beyond 1066MHz. Differential signaling shows very little sensitivity to crosstalk or power supply noise, the two dominant SI/PI effects limiting the data rate of current single-ended solutions. It is therefore possible to operate differential interfaces at multi-gigabit data rates even in low-cost package environment with large crosstalk and power supply noise [6]. It was also shown in [7] that differential links, using voltage mode signaling, can be implemented with very low power dissipation, achieving power efficiency as low as 2.25mW/Gb/s per link. Such transceiver architecture is a promising basis for a mobile memory system that can achieves multi-gigabit data rates at low power in high-density 3D mobile system environments, that would provide the memory bandwidth required for future multimedia mobile devices.

VI. References

- [1] JEDEC Standard JESD209, "Low Power Double Data Rate (LPDDR) SDRAM Specification", JEDEC Solid State Technology Association, August 2007.
- [2] J.-H. Kim, W. Kim, D. Oh, R. Schmitt, J. Feng, C. Yuan, L. Luo, and J. Wilson, "Performance impact of simultaneous switching output noise on graphic memory systems," *In Proc. 16th Topical Meeting on Elect. Perform. Electron. Packag.*, pp. 197-200, Oct. 2007.
- [3] R. Schmitt, J.-H. Kim, W. Kim, D. Oh, J. Feng, C. Yuan, L. Luo, and J. Wilson, "Analyzing the Impact of Simultaneous Switching Noise on System Margin in Gigabit Single-Ended Memory Systems", *DesignCon 2008*, Feb. 3-6, 2008, Santa Clara.
- [4] D. Oh, W. Kim, J.-H. Kim, J. Wilson, R. Schmitt, C. Yuan, E. Lozano, L. Luo, J. Kizer, J. Eble, and F. Ware, "Study of Signal and Power Integrity Challenges in High-Speed Memory I/O Designs Using Single-Ended Signaling Schemes", *DesignCon 2008*, Feb. 3-6, 2008, Santa Clara
- [5] R. Schmitt, J.-H. Kim, C. Yuan, J. Feng, W. Kim, and D. Oh, "Power integrity analysis of DDR2 memory systems during simultaneous switching events", *DesignCon 2006*, Feb. 2006, Santa Clara.
- [6] J.-H. Kim, R. Schmitt, D. Oh, W. Beyenne, M. Li, and A. Vaidyanath, "Feasibility of Multi-Gigabit Memory Interface in LQFP Packages", *DesignCon 2009*, Feb. 2-5, 2009, Santa Clara
- [7] J. Poulton, R. Palmer, A. M. Fuller, T. Greer, J. Eyles, W. J. Dally, and M. Horowitz, "A 14-mW 6.25Gb/s Transceiver in 90-nm CMOS", *IEEE Journal of Solid-State Circuits*, vol. 42, no. 12, pp. 2745-2757, December 2007