Application Note

System Level Verification and Debug of DDR3/4 Memory Designs

Products:

► R&S®RTP

H. Ruckerbauer (EKH), M. Stumpf, R. Wagner | GFM340 | Version 0e | 10.2020

Please find the most up-to-date document on our homepage

http://www.rohde-schwarz.com/appnote/GFM340

Contents

1 Overview

This application note provides an introduction to the DDR memory technology and explains common challenges, related to the specific nature of DDR data, command / address and control buses and describes the typical measurements to verify and debug DDR system designs.

The paper explains the recommended test points and the connection of oscilloscope probes as well as the compensation of effects from DDR interposers via deembedding. The document describes efficient Signal Integrity verification with eye diagram measurements, advanced triggering and TDR/TDT functionality. Given the high number of signal lines and the dynamic bus termination, SSN (simultaneous switching noise) has a significant effect in DDR memory designs and Signal Integrity as well as Power Integrity is highly pattern dependent. We introduce techniques to achieve high acquisition rates and help to efficiently detect worst case scenarios, affecting the performance of the overall memory design. The document also includes a close look on Power Integrity.

Providing best practice examples in the design verification and debugging process, the document addresses all system designers and test engineers, working on DDR memory designs.

Many thanks to Mr. Hermann Ruckerbauer (EyeKnowHow) for contributing significantly to this application note. By combining his expertise in industry and technology with the Test & Measurement solution from Rohde & Schwarz, we have achieved synergies that will benefit design and test engineers.

2 DDR Memory Interface: Technology and Challenges

Memory chips are available in many different types (see overview in Fig. 2-1). Non-volatile memory stores the information even without being powered.

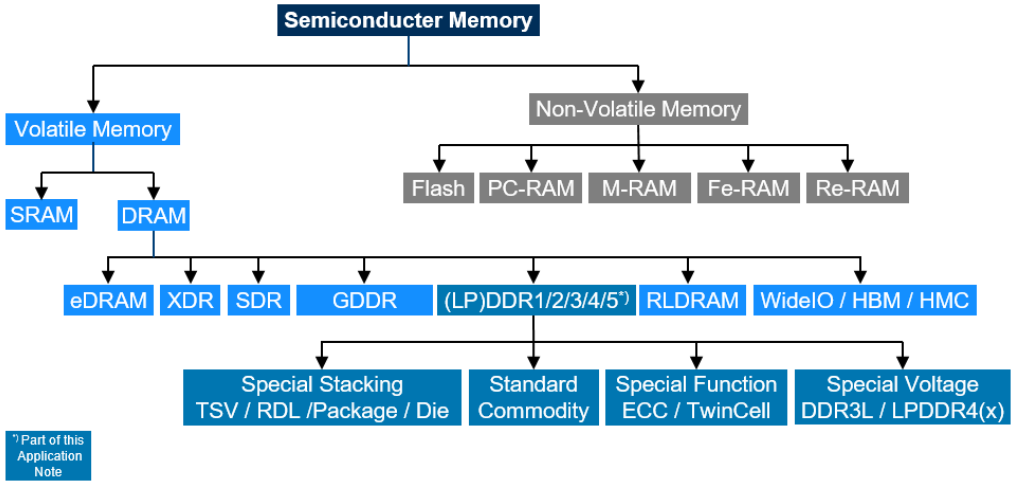

Fig. 2-1: Overview on different semiconductor memory technologies

This is used e.g. as flash memory in SSDs or USB devices. However, this kind of memory has some major disadvantages: the number of write cycles is limited, and the latency e.g. for writing is much higher than for other types of memory. Volatile memory on the other hand, loses its information when power is switched off. SRAMs (Static Random Access Memory) as shown in Fig. 2-2 (left picture) are very fast. As each bit is stored in a structure of six transistors it is very expensive due to its size. More economic are DRAMs (Dynamic Random Access Memory), that are built out of one transistor and one capacitor per bit (see Fig. 2-2 right picture).

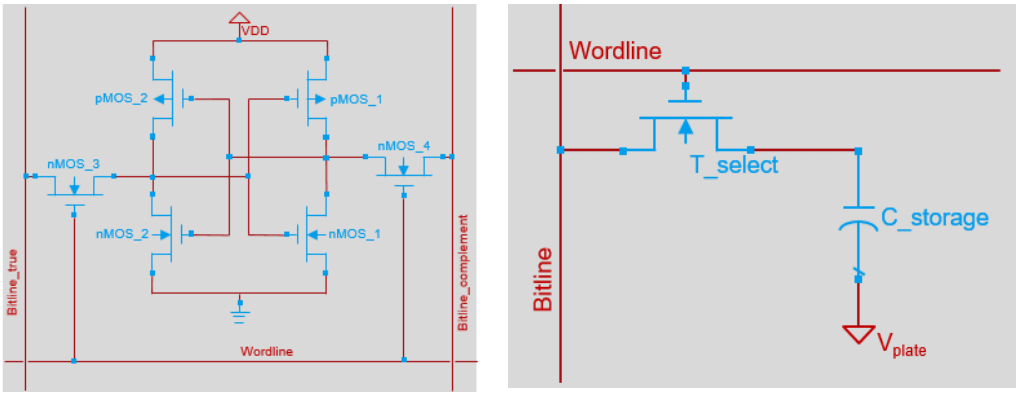

Fig. 2-2: SRAM on the left vs. DRAM circuit in the right picture

SRAMs and DRAMs also have some drawbacks. They are larger than flash memory. Additionally, the storage capacitor in DRAMs needs to be refreshed periodically. Nevertheless, the advantages of DRAMs are big enough, to make them the most important memory technology for any CPU, FPGA or MCU based architecture.

Regardless of whether you design a powerful high-end multiprocessor server or an embedded device as shown in Fig. 2-3, most likely both will have one thing in common: they will use DDR as main memory.

Today's main memory uses a Double Data Rate concept (DDR) to transfer the information on the memory bus. With each rising and falling edge of the data strobe (DQS) 1 data bit (DQ) is transferred per lane. Of course, todays servers demand lots of memory. These systems utilize terabytes of memory from DRAM devices that are mounted on Dual Inline Memory Modules (DIMM) using the latest DDR DRAM technologies. An embedded PC device on the other hand might use only one DRAM component that is directly soldered on the Printed Circuit Board (PCB) and, in many cases, DDR3 density and performance will be perfectly sufficient.

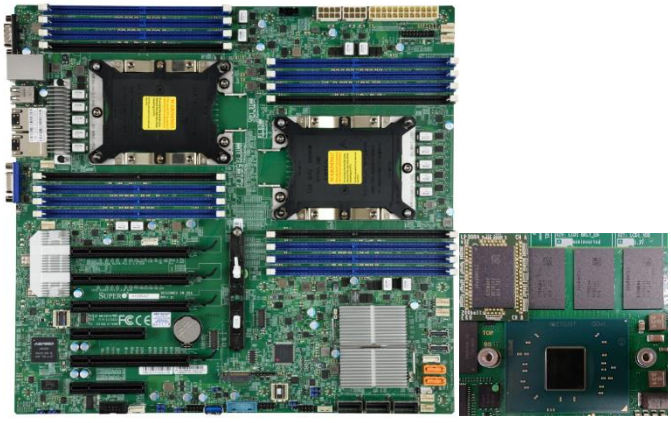

Fig. 2-3: Server system (left, scaled down) AND embedded device (right, scaled up)

DDR DRAMs come for the different generations in different types. For applications where low power consumption is important, system designers can select e.g. LPDDR4 (Low Power DDR4).

Especially small, embedded devices will gain from the pinout and reduced power of LPDDR technology. Comparing DIMM based DDR systems (see Fig. 2-4 with PCB solder down DDR design (see Fig. 2-5) shows, that both use slightly different concepts. The level of complexity is significant higher for the high processing application with DIMMs due to the larger number of DRAMs devices in one system.

Fig. 2-4: DIMM based point to multipoint data bus

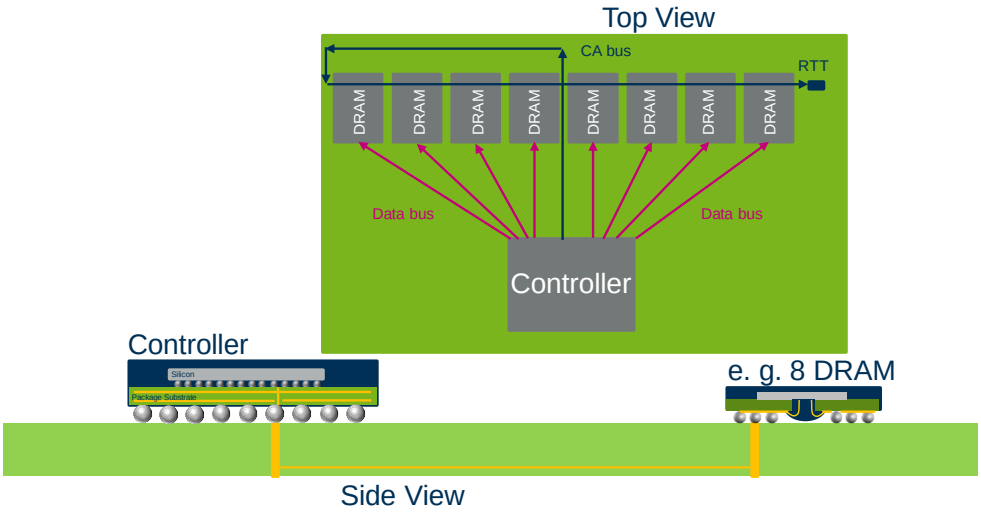

Fig. 2-5: Solder Down design with commodity DRAMs and a fly-by CA bus

However, this does not mean that the verification of a "solder down memory" based system with commodity DRAMs is simple. The memory interface has many special features that make verification difficult. Especially the fly-by Command Address (CA) bus that serves multiple DRAM devices (Fig. 2-5) can be quite challenging. The CA bus is designed by DIMM designers, and needs to be translated to the PCB for such a solder down design. Just the different stack-up and via size, driven by the PCB thickness, will change the CA bus implementation. Another example is the power delivery that is solved for DIMMs and needs to be implemented on the PCB again.

LPDDR addresses these issues. Ballout and stacking concept are optimized to create a full system with a limited amount of DRAM memory components (see Fig. 2-6).

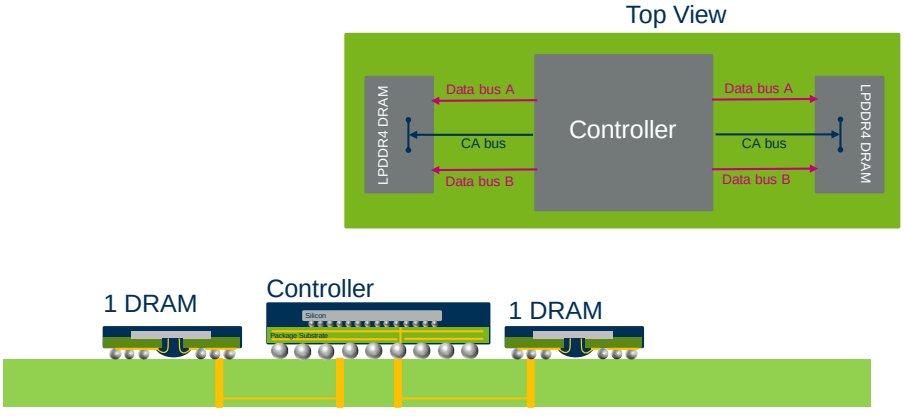

Fig. 2-6: Solder down design with LPDDR4 DRAM

2.1 DDR interface technology

DDR memory features a parallel, source synchronous interface. Source synchronous and parallel refers to the timing reference signal, provided by the source and the fact that data is transferred in parallel blocks of 8bits (1 DQS Strobe per Byte) or 4 bits (1 DQS Strobe per Nibble). A simplified overview is shown in Fig. 2-7. Here we should separate the unidirectional, single-ended CA (Command Address) bus with a differential clock as timing reference and the bidirectional, single-ended DQ (Data Query) bus with a differential DQS (Data Query Strobe) as timing reference.

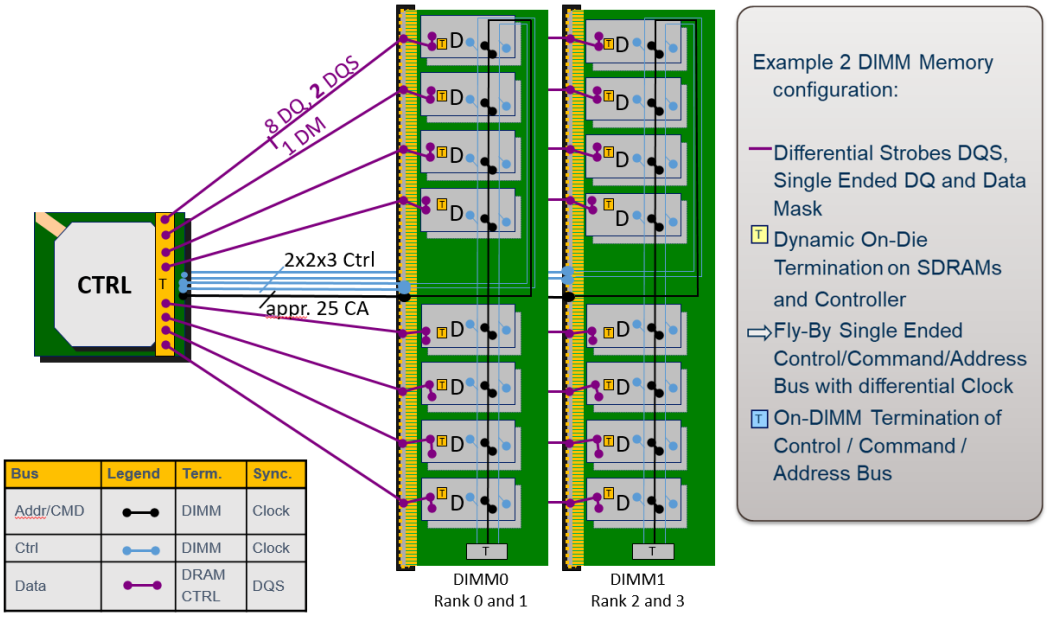

Fig. 2-7: Example: U-DIMM based (unbuffered DIMM), non-ECC memory interface with 4 ranks

Command / Address Signals:

The CA bus is single ended and unidirectional, so command and address information is only transferred from the controller to the DRAM. An overview on all signals can be found in Table 2-1.

Symbol DDR3	Symbol DDR4	Symbol LPDDR4	Type	Logical Group	DDR3	DDR4	LPDDR4
CK/Ck#	CK_t/CK_c	CK_t/CK_c_ A/B	Input	CLK	Differential clock: sampling of CCA (Control / Command/ Address) signals on rising clock edge		
CKE	CKE	CKE_A/B	Input	Ctrl	Clock enable: rank based power down control		
CS#	CS_n	CS_A/B	Input	Ctrl	Chip select: rank selection for parallel DQ lines		
ODT	ODT	ODT_CA_A/B	Input	Ctrl	One die termination : enables termination on DRAM DQ		Enables termination on DRAM CA
	ACT_n	CA[5:0]_A/B	Input	CA		Activation CMD : defines ACT command and allows usage of RAS_n/A16, CAS_n/A15, WE_n/A14 as row address.	Command address bus Independent for channel A and B Single data rate, but commands multiple cycles long
RAS# CAS# WE#	RAS_n/A16 RAS_n/A15 WE_n/A14		Input	CA	Command: select function according to truth table. In DDR4 A14-A16 are reused as additional row addresses.		
	C[0-2]		Input	CA		ChipID: selects chip inside 3D TSV stack	
BA[0-2]	BA[0-1]		Input	CA	Bank Address: defines which bank in DRAM is accessed. For MRS is defines the written register.		
	BG[0-1]		Input	CA		Bank Group: select bank group inside DRAM	
A[0-15]	A[0-17]		Input	CA	Address: multiplexed for row/column access. In DDR4 A14-A16 reused as command for column access.		
A10/AP	A10/AP		Input	CA	Auto-precharge: A10 reused for column access and defines for READ/WRITE if the DRAM executes the precharge.		
A12/BC	A12/BC_n		Input	CA	Burst chop: A12 reused for column access and defines for READ/WRITE if the burst is chopped.		
	PAR		Input	CA		parity: command address parity	

Table 2-1: CA and Ctrl bus signals and CLK for DDR3/4 and LPDDR4

In order to connect many DRAMs to a single CA bus, the fly-by routing concept was introduced with DDR3: about 25 CA signals are “flying” from the controller along all DRAMs until they are terminated on the DIMM at the end of the bus. One of the challenges of the fly-by bus is that each DRAM will see different combinations of reflections, i.e. each DRAM will see a different signal quality. A simulation example by using eye diagrams for each DRAM along a fly-by bus is shown in Fig. 2-8. If the bus is not designed well, this can easily lead to transmission failures when capturing the signals. For a first verification it might not be necessary to check the signals of all DRAMs, but to just select one or two locations. In this case, the best choice is a DRAM at the beginning of the fly-by bus. In many cases, the overlay of all reflections will cause the worst eye at the first or

second DRAM on the bus. When testing two locations, the second one should be at the end of the bus. With this approach some "extreme" locations for checking write leveling training results are covered.

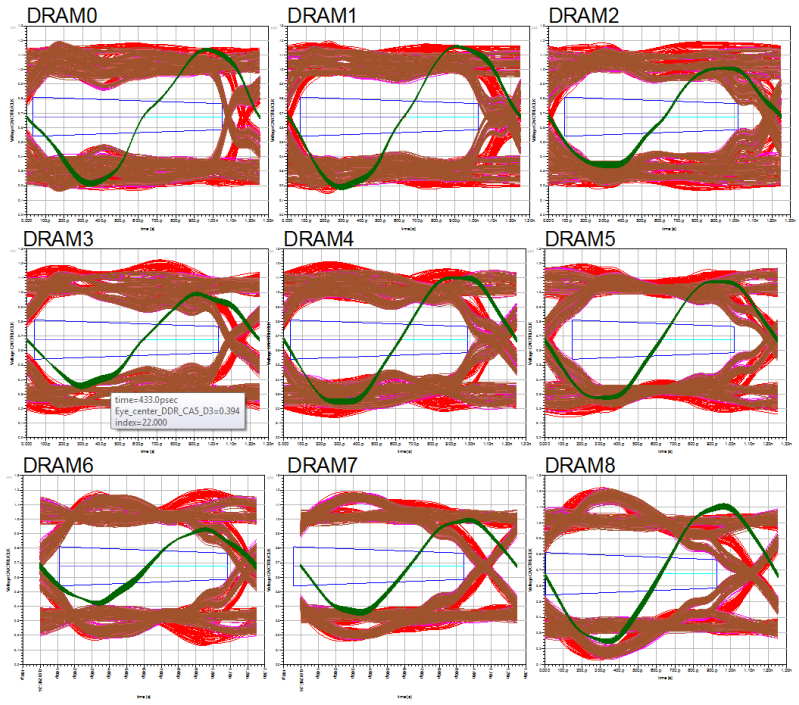

Fig. 2-8: CA Signal quality for different DRAMs along a fly-by bus

For small systems with e.g. 2 or 4 DRAMs it is also possible to implement a T-Branch topology. For this, it is more difficult to find a good termination scheme for high speed, as terminating at each DRAM would reduce the voltage level too much. So, for a T-Branch topology reflections are even more critical and will limit the maximum speed of the system.

The DRAM device captures the information on the CA signals with the rising edge of a differential clock. This interface is single data rate. DDR4-3200 (3200 MT/s) e.g. uses a 1600 MHz clock. The single data rate CA interface is captured one per clock cycle, i.e. with a data rate of 1600 MT/s.

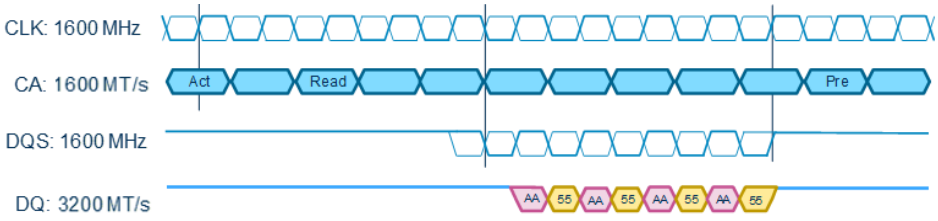

Fig. 2-9: Single data rate (CA) and double data rate (data) for DDR4 3200 MT/s (Read Burst)

Control Signals:

The Control signals (Ctrl, see Fig. 2-7: 4x3 Ctrl signals for a quad rank system) are very similar to the CA signals. However, for a single rank system, they are 1 x available, for a dual rank system, they are 2 x available and for a quad rank system, they are 4 x available. They allow to separate the DRAMs on rank 0 from the DRAMs on rank 1, rank 2 and rank 3 of the system. Each rank provides the number of data lines that the controller defines as his bus width (often 64bit). In order to connect as many DRAMs as possible the

DRAMS in different ranks are connected in parallel. To separate these parallel ranks each rank has its own set of Ctrl signals (see Fig. 2-10).

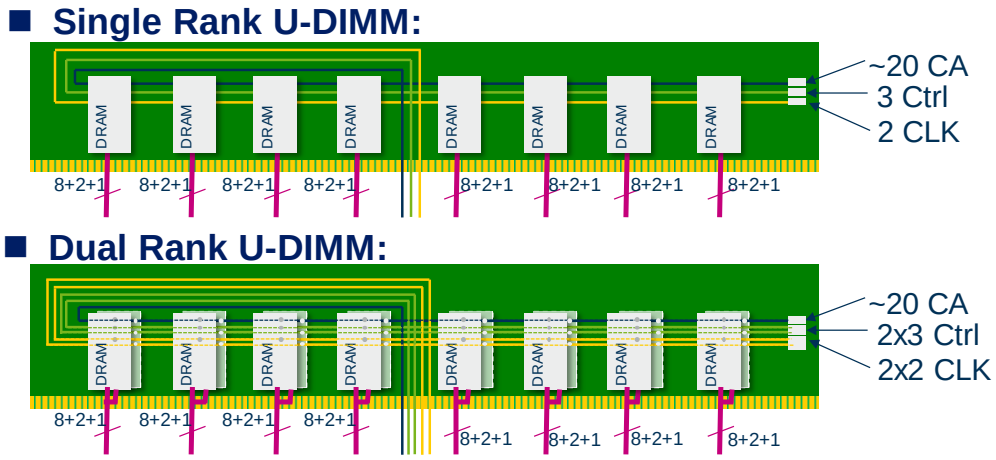

Fig. 2-10: Single rank (top only) vs. dual rank (top / bottom): Example U-DIMM (unbuffered DIMM)

It is important to notice the difference between control (Ctrl) and command / address (CA) bus. On the DQ bus (Data Query bus, 64 bit wide for non-ECC systems) it is possible to connect several DRAMs of different ranks in parallel. To allow to independently access these parallel DRAMs, each set of control signals is only routed to one rank. Therefore, they have less load than the CA signals that are routed to all DRAMs in one channel. This difference has to be taken into account during verification measurements, as the signal integrity on Ctrl and CA busses is different for multi-rank systems.

The dual rank architecture in Fig. 2-10 allows connecting the DQ signals of the DRAMs on the top and bottom of a DIMM to the same DQ lines, separating the data between rank 0 and rank 1 via the chip select signal CS.

Data Signals:

For the bidirectional DQ bus the timing reference is the differential DQS signal. This interface is working in Double Data Rate mode: the DRAM captures the signal with each rising and each falling edge of the DQS Strobe. An overview of all signals with some explanation of the function can be found in Table 2-2.

Symbol DDR3	Symbol DDR4	Symbol LPDDR4	Type	Logical Group	DDR3	DDR4	LPDDR4
DQS DQS#	DQS_t DQS_c	DQS_t_A/B DQS_c_A/B	IO	DQ	Differential data strobe: source synchronous timing information for sampling data on DQ lines. For x8/x16/x32 DRAMs: 1 DQS strobe per DQ[7:0] Byte For x4 DRAMs: 1 DQS strobe per DQ[3:0] Nibble.		
DQ	DQ	DQ_A/B	IO	DQ	Data Query: data signals. Number according to DRAM organization: x4 DRAM: 1 x DQS, 4 x DQ x8 DRAM: 1 x DQS, 8 x DQ x16 DRAM: 2 x DQS, 2x8 DQ x32 DRAM: 4 x DQS: 4x8 DQ		
DM	DM_n DBI_n TDQS_t	DMI_A/B	Input IO	DQ	Data mask: masks selected bits inside a burst	Data mask / data bit inversion / Terminate DQS: function selected by MRS. Some functions are dependent on organization.	Combined data mask or data bit inversion (selected my MR setting)
TDQS TDQS#	TDQS_c		Output IO	DQ	Terminate DQS: used if x8 and x4 DRAMs are in the same system. Different implementation for DDR3 vs. DDR4		

Table 2-2: All signals of the DQ bus for DDR3/4 and LPDDR4

The frequency of the DQS signal is the same as of the clock, e.g. 1600 MHz for DDR4-3200. As the DQ signal is captured with each rising and each falling edge of the DQS strobe, the data rate on the DQ bus is 3200 MT/s (Mega Transfers per second), as shown in Fig. 2-9.

One DQS signal is the timing reference for e.g. 8 DQ lines. However, not only the DDR capture mode is special for the DQ bus. This bus is bidirectional, so either the memory controller (Write Operation) or the DRAMs (Read Operation) are driving the data on the same signal lines. To enable bidirectional data transfer, a dynamic switchable termination is required: ODT (On-Die-Termination) serves this purpose. Dependent on the direction of the data either the DRAM (for WRITEs) or the controller (for READs) will enable the termination on their end of the bus.

Another specialty on this bus is that the data is transmitted in a burst mode, e.g. 8 bits.

As mentioned, the DQS signal is the timing reference for the DQ signals (see Fig. 2-11). There is also some relation between the central clock signal and the DQS signals, but the main important thing for the operation and validation is the timing relation of DQS and DQ. For WRITES, the DQS Strobe is center-aligned to DQ. For READS, the DQS Strobe is edge-aligned to the DQ signal.

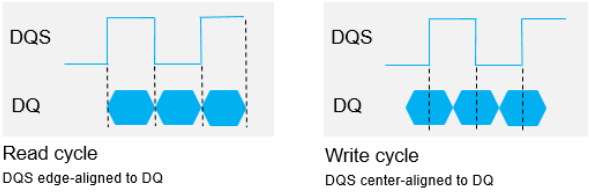

Fig. 2-11: Synchronization of DQ and DQS signals for read and write cycles.

Supply and Reference Voltages and further Signals:

There are still some more signals related to the memory interface. These include supply and reference voltages and signals for special functions (Table 2-3)

Symbol DDR3	Symbol DDR4	Symbol LPDDR4	Type	Logical Group	DDR3	DDR4	LPDDR4
VDDQ	VDDQ	VDDQ	Supply	Supply	DQ power supply: 1.1 V – 1.5 V dependent on devices (isolated from VDD)		
VSSQ	VSSQ	VSSQ	Supply	Supply	DQ GND		
VDD	VDD	VDD1/VDD2	Supply	Supply	Power supply: 1.1 V – 1.8 V dependent on devices (isolated supplies on package)		
VSS	VSS	VSS	Supply	Supply	Logic and CA GND		
	VPP		Supply	Supply	Supply voltage for internal DRAM wordline		
VrefCA	VrefCA		Misc	Misc	Reference voltage for Control / Command / Address lines (CCA)		
VrefDQ			Misc	Misc	Reference voltage for data lines (DQ)		
ZQ	ZQ	ZQ	Misc	Misc	Reference pin for ZQ calibration		
Reset#	Reset_n	Reset_n	Input	Misc	Reset: asynchronous signal for resetting the DRAM		
	Alert_n		IO	Misc		Alert: output: CRC or parity error flag Input: for connectivity TM	
	TEN		Input	Misc		Connectivity TM enable: enables TM on x16 (optional on x4/x8)	

Table 2-3: Miscellaneous signals and voltage supplies

There is still one supply missing: VTT (Termination voltage for the CA bus). Reason for this is that VTT is just a supporting system voltage to achieve good signal integrity by terminating the CA signal lines at the end of the bus. The VTT voltage from the system is not applied to any DRAM.

Creating different systems and different configurations with all of these signals can lead to quite complex system topologies and quite some challenges to verify these systems.

2.2 Typical challenges for DDR memory verification

Data Signals:

The data bus for DDR DRAM memory is challenging, due to the following attributes:

- ► Asymmetric point to multipoint connection on the data bus in systems with multiple memory ranks (e.g. 4 ranks)
- ► Parallel, source synchronous interfaces (1 DQS for e.g. 8 DQ signals)
- ► Bidirectional signaling for read/write operation on the DQ bus
- ► Dynamic termination for read/write cycles
- ► Multiple options to switch termination values: dynamic ODT and RTTnom
- ► Tri-state (high impedance) on DQ bus during non-active times

Due to the multipoint connection on the data bus of a multi-rank system, WRITE signals arrive at all connected DRAM devices (Fig. 2-12). They are assigned to the corresponding rank via the Chip Select signals CS. Dynamic on-die termination (ODT) is automatically enabled at the selected DRAM during this write cycle (RTT_WR). In this example, this is the left DRAM. Termination of the other, not selected DRAMs are enabled by a high level on the ODT signal of their Ctrl busses (RTT_Nom). In this example, only the right DRAM is terminated. During the measurement of the DRAM under test (e.g. left DRAM), the WRITE accesses to the other DRAMs will be visible as well. For proper analysis, it is necessary to filter and evaluate only the accesses to the measured rank. For multi rank systems, DRAMs from each rank need to be verified independently. In the below example, four ranks are displayed, only the last, not accessed rank is terminated with RTT_Nom. It is up to the system designer to find the best termination scheme.

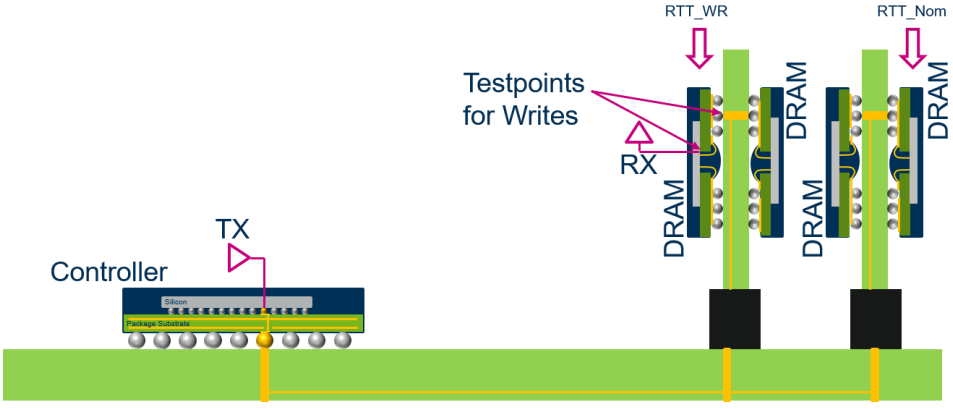

Fig. 2-12: Databus example: WRITE termination in a quad rank system (distributed on 2 DIMMs)

As the bus is asymmetric, it cannot be assumed that READ and WRTE will have the same signal quality. Both directions need to be optimized independently. WRITES should be verified at the DRAM input. For READs the best way to verify signal quality is to measure at the controller input. Fig. 2-13 shows the termination scheme for a READ cycle. The controller hereby activates the termination on its end (SoC termination).

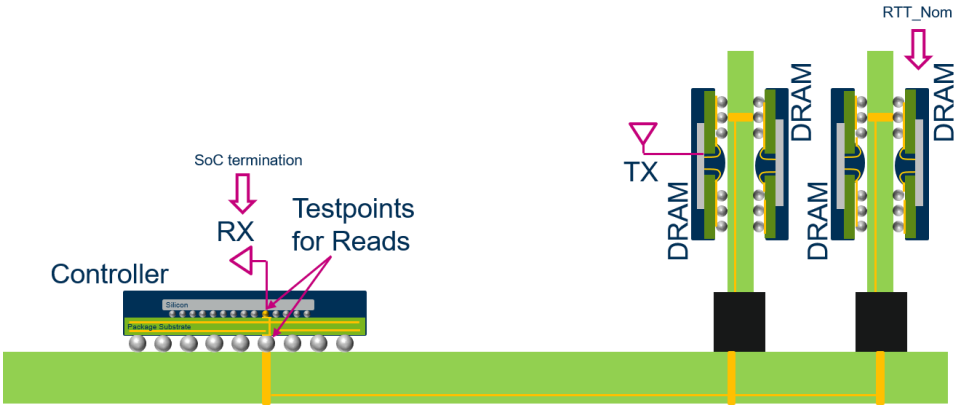

Fig. 2-13 Databus example: READ termination in a quad rank system (distributed on 2 DIMMs)

The width of the parallel bus (e.g. 64 bit total DQ width in a non-ECC system) makes it impossible to select all signals for verification or measurement. For debugging, select the DQ that shows failures. For system verification, select the worst-case routed DQ signals.

Automatic separation of READ and WRITE in a small point-to-point system is already difficult. For multi-rank systems this is even a bigger challenge due to different signal integrity characteristics and voltage levels for each of the ranks.

The system can be configured to use different termination options for different conditions in the system. For any debugging, the configuration used must be known exactly in order to understand the bus signals.

Last but not least, many data and CA signals are in tri-state mode (unterminated) most of the time. This means triggering to midlevel, or creating any data eye is not an easy task, as there are always some signals in the middle of the voltage range.

So all of these DRAM specific features will make debugging and verification of the DQ bus a very challenging task where knowledge, experience and the right test tools are required.

Control / Command / Address Signals:

The CCA bus seems to be simpler (see Fig. 2-14), as it is slower. But due to the high number of loads it can be even more challenging to handle. Besides having similar issues as the DQ bus (being source synchronous, allowing tri-state on the bus), each DRAM along the fly-by bus is confronted with a different signal quality and would need to be verified independently (see Fig. 2-8).

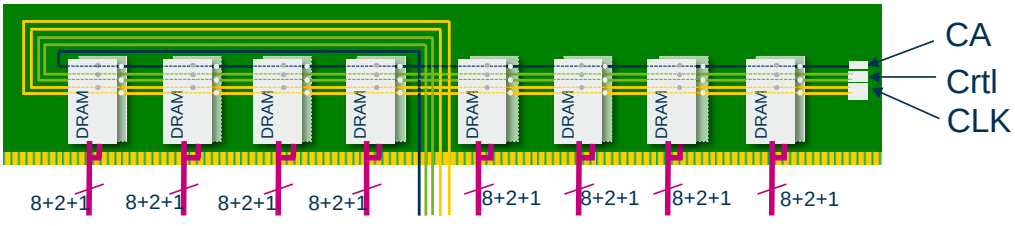

Fig. 2-14: CCA bus example with 16 loads along the fly-by bus (64 bit controller bus width, dual rank)

Further challenges:

Other specials features of the memory interfaces are:

► The JEDEC specification defines the signal quality at the ball of the DRAM device (see Fig. 2-15)

- There is no commodity specification for READ at the controller input, but each controller has its own proprietary input specification.
- There is no specification on how to execute the verification and compliance test.
- There is no defined test pattern that needs to be used during verification.
- The main memory is a required part for minimum functionality for most systems. Therefore, the memory needs to be in operational state during verification test. It is not possible to set the DRAM interface in compliance test mode for verification tests. The only option is to run a memory test (e.g. memtest86), but usually even this test is running out of the tested main memory.

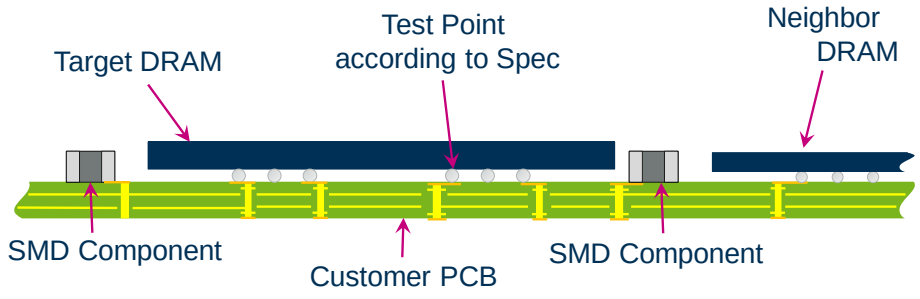

Fig. 2-15: Testpoint according to the JEDEC specification

A very thorough verification of the interfaces is also required because of the following reason: DRAMs do have a short product life cycle and sometimes the DRAM vendor releases the next shrink of the DRAM devices even before the tested system is in mass production. Because the new DRAM should be a 1:1 replacement of the old device each used device needs to be tested and released for usage in the given application. Quite often issues happen when using new DRAMs, but in most cases it is not the new DRAM that causes the problem, but the system was designed with insufficient margin and was already close to the FAIL border before.

With each new DDR DRAM generation, the verification is getting more and more difficult due to following reasons:

- Data rates and system speeds are increasing and are strongly impacting signal integrity. This also makes testing more and more difficult.
- The JEDEC DRAM specification is a specification for the DRAM device, not a system specification. It is not written for system verification.
- For system READs there is no test defined by JEDEC, as JEDEC only defines the data output from the DRAM on an ideal test board.
- Signal traces need to be embedded or de-embedded in order to see the signal quality at locations where an oscilloscope probe cannot be connected. Connecting to the required test points underneath the BGA device is not possible.

Therefore, it is already required during design and layout phase to think about the verification concept of the DRAM interface: DFT (Design For Testability) is not only nice to have, but absolutely necessary in order to get the verification done later on.

3 Performing the measurement

3.1 Connecting the oscilloscope to the test points

The test point where the measurement should take place is underneath the memory, right at the BGA interface. Therefore, the big question is: How to connect the oscilloscope probe?

In the past, vias and test points could be placed in the design to allow access to the signals with an oscilloscope probe. For DDR3 designs, this is still possible from a signal quality point of view. However, this makes it quite difficult to understand the test results because additional delays and reflections due to the location of the test point have to be taken into account. For these reasons, interposers have been invented (see Fig. 3-1). They are located between the memory and the PCB and provide solder pads to connect the scope probe, including an additional (buried) resistor to the probe and an additional via between DRAM and PCB (Fig. 3-2). The buried resistor is required to decouple the stub to the test pad from the real signal path. Without this buried resistor, reflections from the test pad would be too high and disturb the signal quality. There is a similar resistor on the active probe tip itself. It is therefore required to consider the value of the buried resistor in the interposer and reduce the resistor on the probe tip accordingly. The range of resistors used in interposers is typically 75 Ω up to 175 Ω.

To remove the effect of the interposer, the interposer needs to be deembedded from the measurement result. For small interposer size and low system speed, deembedding of the interposer only has a minor effect. For the higher speeds of DDR4 and LPDDR4 devices, deembedding of the additional interposer signal path is important to get correct measurement results. Oscilloscopes like the R&S® RTP offer deembedding functions that can compensate the transmission loss of the interposer. The Interposer supplier determines the signal path characteristics by simulation or measurement and provides it as S-parameters.

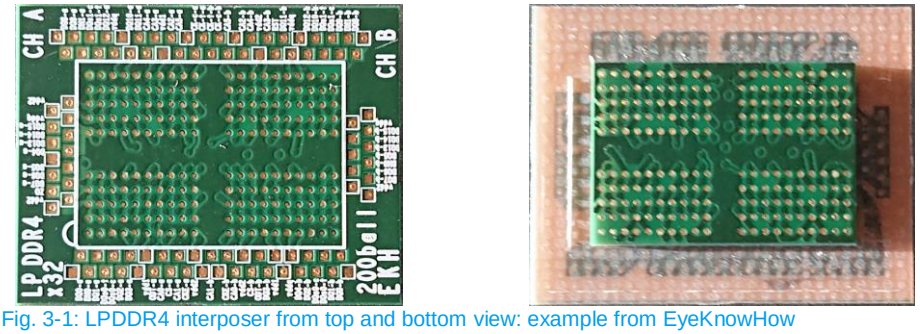

Fig. 3-1: LPDDR4 interposer from top and bottom view: example from EyeKnowHow

The use of interposers is also fueled by another industry trend: designs using HDI layout (High Density Interconnect) do not have any signals routed to the top or bottom, but everything is hidden in the inner layers, connected by blind and buried vias. Even when adding additional test points it is not possible to do this for all signals. Especially for debug purposes, it often happens that the test point is set to the wrong signal. With the right interposer you will have access to all signals.

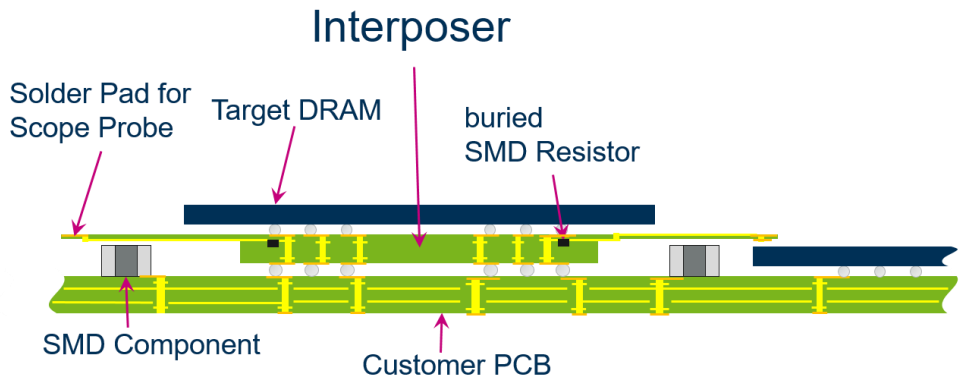

Fig. 3-2: DDR Interposer side view

In order to prepare the DUT (Device Under Test) for the measurement, getting access to all signals is just half of the job. Unfortunately, normal oscilloscopes only have up to four channels, but the DRAM does have many more signals. In order to do the measurements efficiently it is most important to prepare the measurement setup in a way to avoid any re-soldering during the measurements. Even if you try to make the setup stable, any movement should be avoided as it might damage the setup. In addition, it is just not efficient to stop the measurement for re-soldering any probes.

Even though modern differential active probes have reduced in size, it is a challenge to prepare a setup with four or more probe tips soldered to the memory interface.

By using P/N/DM/CM multimode probes like R&S® RT-ZM modular probe with compact flex tip modules R&S® RT-ZMA14 (see Fig. 3-3), one can double the number of single ended signals that can be measured with every setup by just switching the scope from p-Mode to n-Mode without a physical change of the setup. For all single ended signals, p and n should be soldered to two different signals and also be connected to GND.
Fig. 3-3 shows an example with eight multimode probes prepared and connected. At the prepared probes on the left picture, all silver wires have approximately the same length.

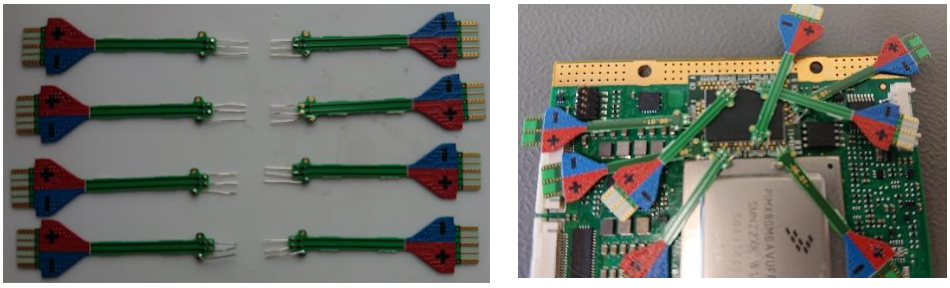

Fig. 3-3: Eight probe tips for multimode usage, prepared (left) and connected (right): Example with probe tip modules R&S® RT-ZMA14

Best practice for soldering the probes is always to prepare the probes first with the wires before soldering them to the DUT. They need to be shortened and adjusted based on the location of the probe pads on the target system. With 8 multimode probe tips it is easy to e.g. get access to differential CLK and DQS signals (besides in differential mode DM, they also can be measured in common mode CM or single ended as P and N signals) and 12 single ended signals just by changing the probe head setting. Thanks to the modular probe system, the 8 probe tips can be quickly connected to the 4 R&S® RT-ZM probe modules, connected to the oscilloscope.

This is a huge advantage over a typical four-channel test setup with one differential clock, one differential DQS, one single ended DQ and one single ended Ctrl or CA signal. And while it is getting more and more difficult in traditional setups to add further test points, this is quite simple due to the small and flexible R&S® RT-ZMA14 solder-in tips.

If you first want to solder the wire to the device under test and then connect it to the probe, make sure that the solder connections on the pad are still visible. Sometimes it is necessary to re-work this connection, what is nearly impossible if the solder joints are no longer visible.

Example probe setup for Control / Command / Address Signals:

This technique and a setup with 4 multimode probes allows verifying the eye diagrams for one CLK and up to six CCA signals in one physical setup, if the probes have been e.g. connected to CS, ODT, A3, BA2, one further command and one further address signal. These signals are particularly interesting, as they provide specific insights into the behavior of the DDR interface:

- CS is indicating an access to the given rank
- ODT can be used to understand the behavior of the termination
- A3 should switch very often
- BA2 should switch rarely
- one more command (e. g. RAS) allows to better understand the command sequence
- another address mainly for statistical reasons.

Example probe setup for Data Signals:

With the above setup, we can use four probes e.g. connecting channel 1 to DQS and channel 2, 3 and 4 to the single ended DQ signals. No re-soldering is required, just connect the probe modules to other, already pre-soldered probe tips. Using multimode probes, we can even measure one DQS and up to six DQ signals, just by switching from p- to n-mode for the single ended signals.

3.2 Deembedding the signal path from the BGA interface to the oscilloscope

For every interposer you will get S-parameters that describe its characteristics. Based on that, the oscilloscope can deembed the interposer, i.e. it can calculate how the signal would look like without this interposer. With that, the oscilloscope shows the signal at the ball of the DRAM, rather than on the test pad of the interposer. This is similar to the calculation to remove the load and frequency domain behavior of the probe. Also here, the effects of the probe are removed and the signal is shown how it would look like, if an ideal probe would be used.

With today's oscilloscopes, deembedding is typically implemented as a SW-post-processing stage after the acquisition of the waveform. From the captured signal, the wanted signal at the BGA interface is being calculated, removing the effects of interposer and probes. This approach however has 2 major drawbacks. Post-processing slows down the acquisition rate of the oscilloscope. Detecting sporadic worst case patterns on the DDR interface, driven by simultaneous switching of outputs, inputs and termination, will get more time consuming and less efficient. Besides that, signal triggering can only be performed on the non-deembedded waveform, not on the waveform of interest at the ball of the DRAM. Besides this traditional SW-based deembedding (R&S® RTP-K121), the R&S® RTP oscilloscope therefore also provides a HW-based real-time deembedding option (R&S® RTP-K122). Using the real-time architecture of the instrument, deembedding is performed in real-time on the stream of samples coming from the front-end of the instrument (see Fig. 3-4)

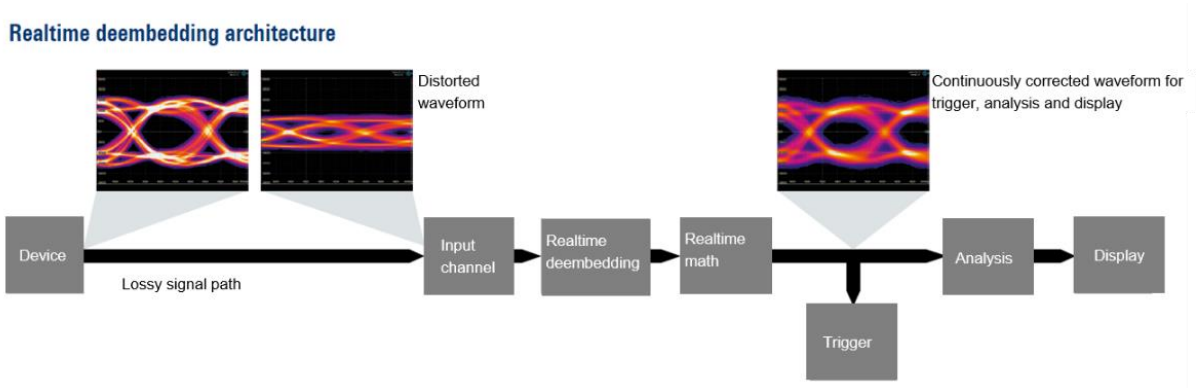

Fig. 3-4: Real-time architecture of the R&S® RTP oscilloscope

Real time deembedding in R&S® RTP (see the red rectangle in Fig. 3-5) is done in hardware. This process is very fast. You can trigger on the corrected signal, and the acquisition rate, e.g. during the eye diagram generation remains unchanged.

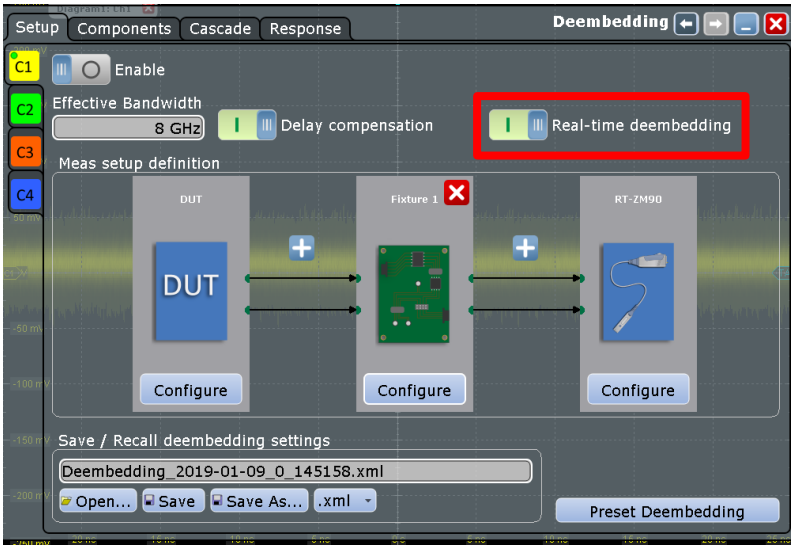

Fig. 3-5: Using real-time deembedding in R&S® RTP

3.3 Compliance test and system-level verification and debug

Beside the measurement hardware, the right software should be available for testing too. This includes two different kinds of software: The first one is used to generate the required data traffic on the DUT for testing. There are many memory test tools available (see Appendix). The right tool for the test should be selected e.g. according to target platform and required test pattern.

The other software is the DDR analysis software like R&S® RTP-K93 and R&S® RTP-K91 that runs on the oscilloscope. This software is essential as it is very difficult and time-consuming to manually perform all measurements, required to correctly verify the parameters defined in the DRAM specifications. The mentioned options support signal integrity debugging and automated compliance testing of DDR4/ LPDDR4 and DDR3/ DDR3L/ LPDDR3 interfaces respectively. A wizard guides the user via illustrated step-by-step instructions through the compliance tests. A configurable test report (see Fig. 3-6) documents the test results.

In this test report, the test engineer immediately recognizes whether the DUT has passed the test or not (pass / fail). Additionally, the options support manual signal integrity debugging by the integrated DDR eye diagram function and READ/WRITE cycle separation and decoding.

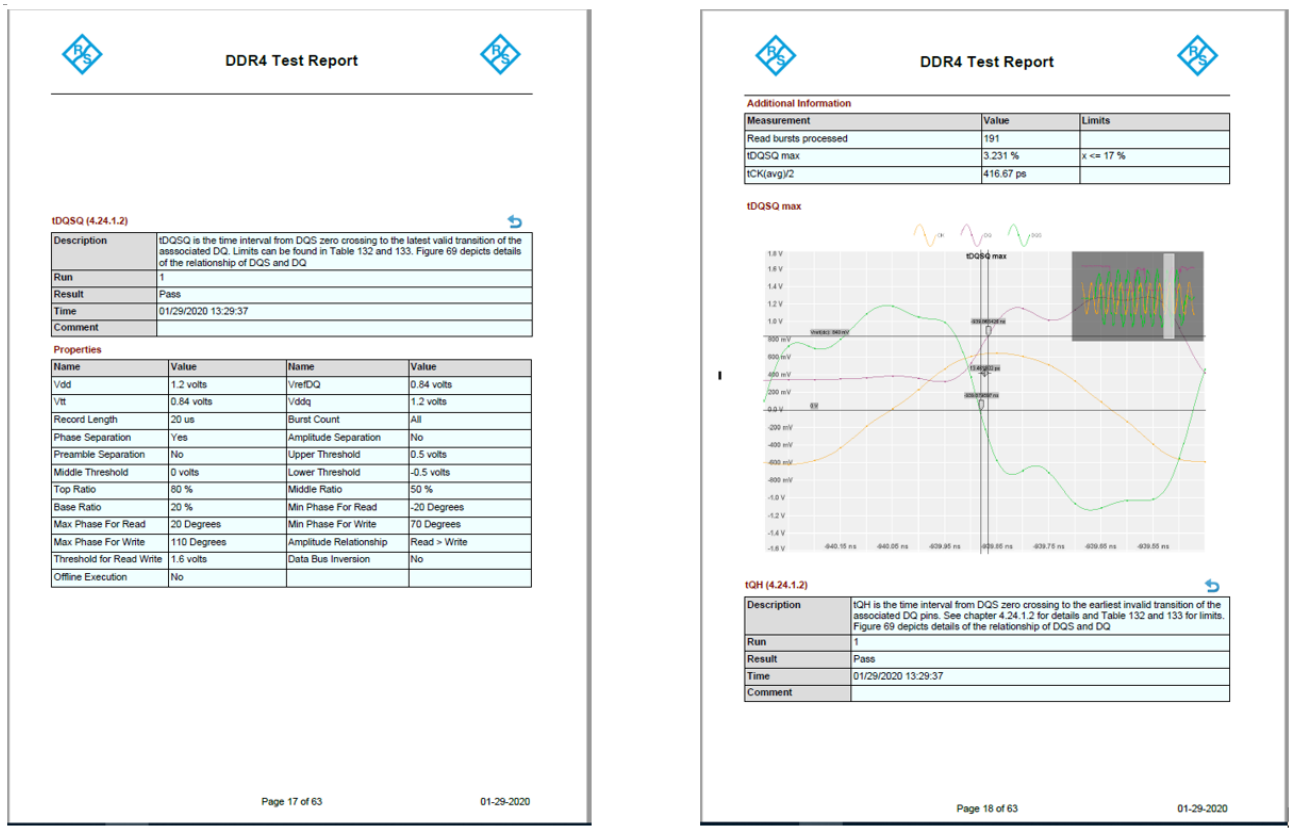

DDR4 Test Report

tDQSQ (4.24.1.2)

Description	tDQSQ is the time interval from DQS zero crossing to the latest valid transition of the associated DQ. Limits can be found in Table 132 and 133. Figure 69 depicts details of the relationship of DQS and DQ
Run	1
Result	Pass
Time	01/29/2020 13:29:37
Comment	

Properties

Name	Value	Name	Value
Vdd	1.2 volts	VrefDQ	0.84 volts
Vtt	0.84 volts	Vddq	1.2 volts
Record Length	20 us	Burst Count	All
Phase Separation	Yes	Amplitude Separation	No
Preamble Separation	No	Upper Threshold	0.5 volts
Middle Threshold	0 volts	Lower Threshold	-0.5 volts
Top Ratio	80 %	Middle Ratio	50 %
Base Ratio	20 %	Min Phase For Read	-20 Degrees
Max Phase For Read	20 Degrees	Min Phase For Write	70 Degrees
Max Phase For Write	110 Degrees	Amplitude Relationship	Read > Write
Threshold for Read Write	1.6 volts	Data Bus Inversion	No
Offline Execution	No		

Page 17 of 63 01-29-2020

DDR4 Test Report

Additional Information

Measurement	Value	Limits
Read bursts processed	191	
tDQSQ max	3.231 %	x <= 17 %
tCK(avg)/2	416.67 ps	

tDQSQ max

tQH (4.24.1.2)

Description	tQH is the time interval from DQS zero crossing to the earliest invalid transition of the associated DQ pins. See chapter 4.24.1.2 for details and Table 132 and 133 for limits. Figure 69 depicts details of the relationship of DQS and DQ
Run	1
Result	Pass
Time	01/29/2020 13:29:37
Comment	

Page 18 of 63 01-29-2020

Fig. 3-6: Example for a DDR4 compliance test report

The JEDEC DRAM specification defines both, the WRITE parameters at the memory input and the READ parameters at the memory output. Writes are best measured close to the DRAM to include the influence of the system design and to verify proper signal integrity at the memory input. The READ parameters, defined in the DRAM specification however are memory output parameters. According to the specification, they need to be measured on specific test boards with an ideal test load (see Fig. 3-7). These memory output parameters are of no real interest for system verification. They have already been tested in the compliance test of the memory device itself. Interesting on a system level is, whether the signal quality at the controller input is good enough to fulfill the controller specification, but the measurement at the controller is much more difficult. There is no common ballout for controllers and therefore usually no interposers are available. Even with an interposer, the controller package is too large to really allow judging signal quality at the ball. Therefore, it is very difficult to do a READ verification for system designs. In case the signal path can be characterized with S-Parameters, probing at the memory device and signal embedding is a possible solution.

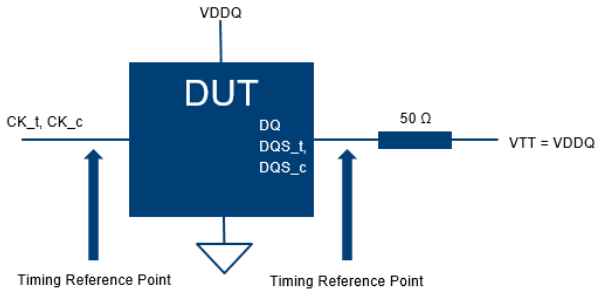

Fig. 3-7: JEDEC reference load for DDR4 DRAM output test

Additionally, it must always be taken into account that such compliance measurements are only a short snapshot with very little statistical certainty, as they are based on single microseconds-long acquisitions. And there is one thing, system verification engineers always should keep in mind: in memory design it is all about statistics and probabilities.

Just having a 20 µs trace evaluated will never capture enough accesses to get all effects of Inter Symbol Interference (ISI), crosstalk and simultaneous switching covered. Due to the nature of the DDR interface with its high number signal lines and dynamic termination, simultaneous switching noise (SSN) on the outputs and inputs (SSO and SSI) has a high impact on the Signal Integrity of the design. Being data pattern dependent, SSN produces sporadic worst case scenarios, which can only be detected with a long acquisition time, or better, a high acquisition rate of the oscilloscope. Even taking automated repetitions of the compliance app will not provide enough statistics or insight into the signal quality.

For system-level verification and debugging solid manual engineering work is required. The good news: the R&S® RTP oscilloscope provide several tools to simplify the process.

3.4 Measuring DDR Eye Diagrams for CCA and Data (WRITE/READ) Signals

In the past, the DRAM specification was purely focusing on the source synchronous edge by edge timing (setup & hold time). DRAM specifications before DDR4 don't define mask parameters for any eye diagrams. Nevertheless, an eye diagram is one of the most important ways to evaluate signal integrity on any interface. With eye diagrams, the statistical certainty of the measurement is defined by the repetition rate of the scope (dead time between two captures) and the runtime for capturing the eye diagram. Getting millions of captures is possible e.g. during overnight runtimes. By doing this, the chance increases significantly to capture critical events and catch worst case scenarios. You will also receive statistics with ISI (Inter Symbol Interference), crosstalk and SSN effects. Considering that the JEDEC DDR4 DRAM specification talks about a Bit Error Rate (BER) of 1e-16, capturing just 1e6 events is not sufficient. Eye diagrams need to be measured over a longer time to get the required statistical certainty. Based on the real-time architecture and hardware-accelerated measurement implementation in the R&S® RTP, the Options R&S® RTP-K93 and R&S® RTP-K91 include a powerful eye diagram functionality to efficiently measure and analyze DDR4 and DDR3 eye diagrams.

Control / Command / Address Signals:

Even for DDR3/DDR4 CA busses and DDR3 data busses where the specification only defines parameters like tIS/tIH and tDS/tDH (setup and hold time for CA and data signals) or tVAC (valid transition time), an eye diagram is the tool of choice for signal integrity verification and debugging. The difference of the DDR memory interface to serial interfaces such as USB or PCIe with an embedded clock is that the DDR eye diagram needs to be generated to a source synchronous reference.

Fig. 3-8: Eye diagram test of CA signals with R&S® RTP-K93

For the unidirectional control / command / address signals, this timing reference is the clock signal. The R&S® RTP-K93 / R&S® RTP-K91 DDR eye diagram function supports a comfortable setup, selecting the tested signal and the corresponding timing reference (see Fig. 3-8). Since control signals are per rank and CA signals might be connected to multiple ranks, both should be included in the verification process. Persistence mode can be used to gain higher statistical certainty over many acquisitions.

Using the multimode probe setup in Fig. 3-3, it is possible to easily probe multiple CA signals at one DRAM position. But looking to a normal DDR3/DDR4 command / address fly-by bus there will be a different signal quality at each DRAM position. Due to reflections, each DRAM gets a different signal quality and effects like attenuated clock levels and reflections change the behavior significantly. Often the first DRAMs show the worst signal quality, so DRAMs at the beginning of the bus are the location of choice if you select just one DRAM for measuring.

Data Signals:

Due to the fact, that the data bus is bidirectional it is necessary to separately create a data eye for READ and a data eye for WRITE.

To distinguish between READ and WRITE cycles, the R&S® RTP-K93 / R&S® RT-K91 options provide a DDR decode feature for DDR READ/WRITE separation (Fig. 3-9). Needed are only DQS and DQ signals. The software separates READ from WRITE accesses, mainly by looking whether the DQS edge is edge-aligned (READ) or center-aligned (Write) to the DQ signal (see Fig. 2-11):

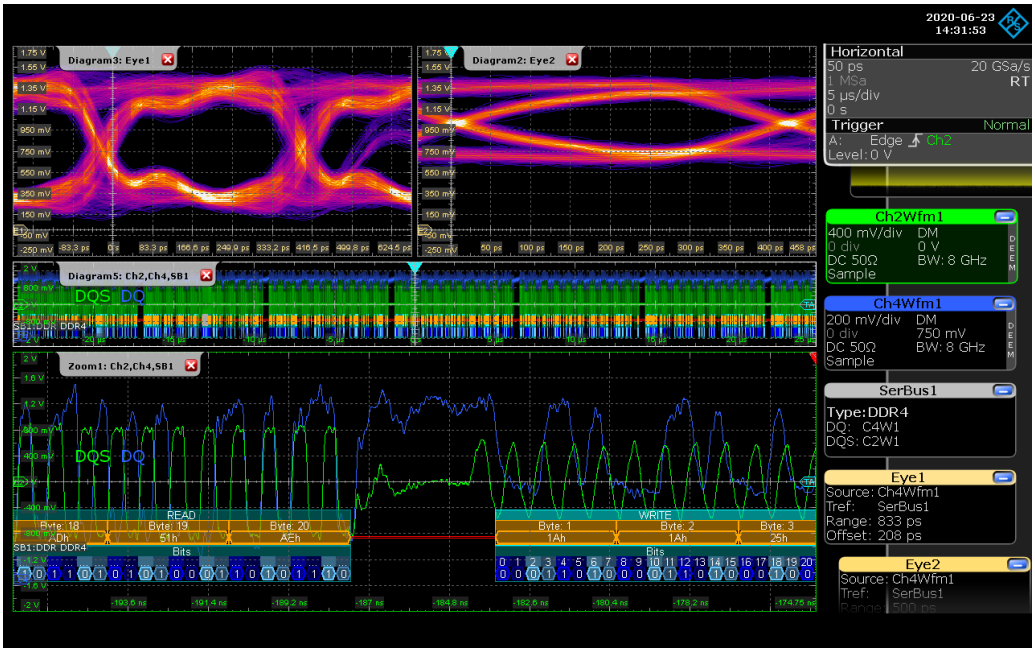

Fig. 3-9: Eye diagram test of data signals for READ and WRITE bursts with R&S® RTP-K93

Doing this for single ranks systems is somehow simple. The interposer allows measuring at the ball of the DRAM where the spec defines the measurement point for input signals. For multi-rank systems, it can be difficult to separate the WRITEs to the different ranks, but they will have significantly different Signal Integrity when measured at one DRAM. Dependent where the WRITEs are targeted to, different termination configurations will be used. In this case, most likely limiting the address range of the test software can improve the situation quite a lot and will simplify the testing.

3.5 Advanced Triggering for Measurements on CCA and Data (WRITE/READ) Signals

Eye Diagrams are an essential tool for system-level verification and debug. Nevertheless, there are many cases, where you may not want to overlay all bits from the bursts, but want to take a look at whole patterns or individual bits:

- ► For debugging, you may not want to overlay all bits from the bursts, but really want to see the burst start and burst end.
- ► Analysis of turnaround times (e.g. READ to WRITE), as they might impact the signal behavior.
- ► In DDR signaling different bits can show significant differences (e.g. odd bits at rising DQS edges and even bits at falling DQS edges). This should be visualized. One effect to cause this can be a duty cycle distortion on the incoming clock for READs. Also for WRITEs it can happen, that a duty cycle distortion on the clock will be found on the DQS and data bits as well.

The advanced triggering capabilities of R&S® RTP help to efficiently catch and display the events of interest on the DDR interface.

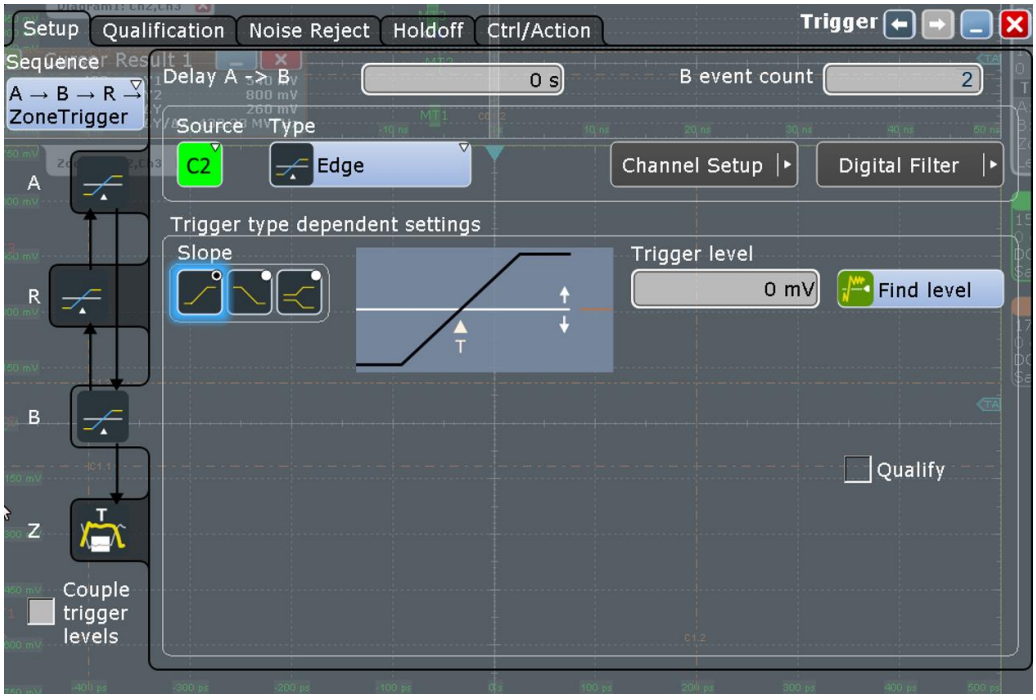

Fig. 3-10: Advanced trigger setup: two stage trigger with event count and Zone Trigger

Fig. 3-11 shows an example for such an advanced trigger setup, the trigger settings are shown in Fig. 3-10. Measured signals are CLK (channel 1, yellow), DQS (channel 2, green), DQ (channel 3, orange) and CS (channel 4, blue). An asymmetry between even and odd data bits is revealed.

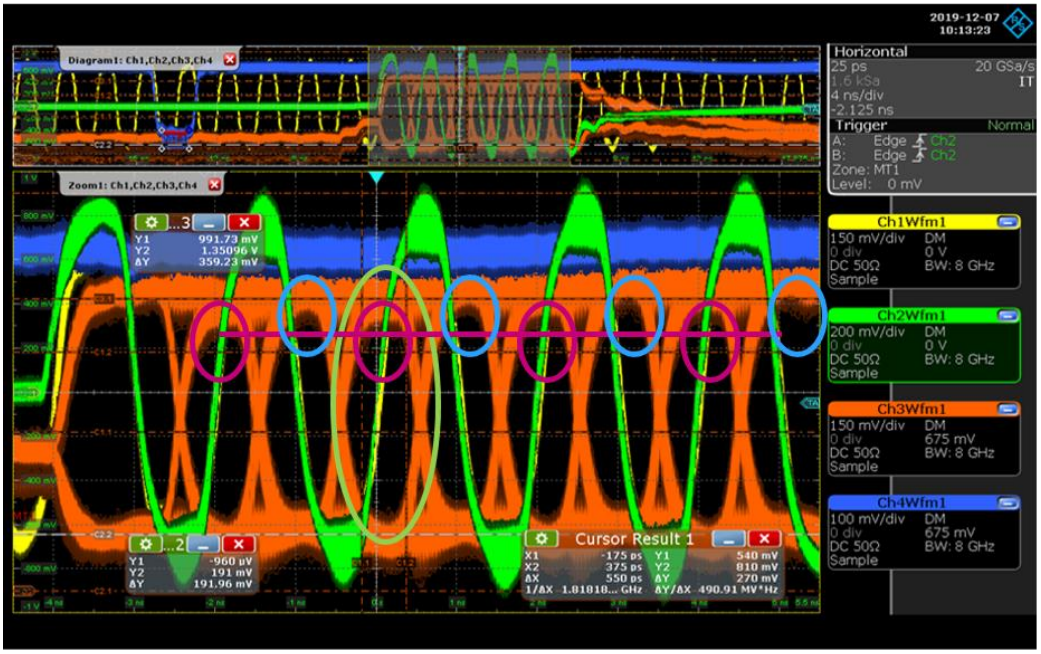

Fig. 3-11: Advanced triggering: two stage trigger with event count and Zone Trigger to capture DQ bursts with bit by bit overlay: odd bits have less margin (magenta Circles) than even bits (blue circle)

The odd bits are always a bit lower (magenta circles) than the even bits (blue circles), reducing the high side margin of the eye. An overlay of all bits in a burst would hide this behavior. The view in Fig. 3 11 allows to

individually analyze all bits inside the WRITE burst.

In this case, Zone Trigger (R&S® RTP-K19) was used in combination with the CS signal that needs to be low CWL (CAS Write Latency) before the burst starts. This is shown in the upper part of the screen with the full capture time. Additionally, the A trigger (edge) on DQS was set to -50 mV to skip the first rising edge, when the bus gets from tri-state into the WRITE cycle. With that, the trigger is armed on the second rising edge. The B trigger (edge) is set to 0 mV and an event count of 2 to trigger to the DQS zero-crossing (sampling point) of the 3rd bit of the burst (green circle indicates final real trigger location). With this, all captures are overlayed, aligned to the 3rd bit of the WRITE burst (bit of interest in this example). The eye diagram of this 3rd bit is shown source synchronous to its DQS edge, while the other bits are not.

It is very important to set the zones for the Zone trigger correctly, in order to catch all relevant events and to keep the high update rate of the scope. Last but not least, event count can be used to select the bit of interest (here the 3rd bit) in the burst.

Fig. 3-10 shows the definition of A-B trigger with zone trigger and B event count to display a special bit in the burst source synchronously to the DQS.

Fig. 3-12 shows another example, where advanced triggering can help. Measured signals here are CLK (channel 1, yellow), A8 (channel 2, green), CS (channel 3, orange) and WE (channel 4, blue). Rare issues as shown here (38 A8 mask failures in 10 million acquisitions) can only be detected in tests, where many acquisitions are collected. Memory testing is all about statistics!

The result also shows that there are many rising edges of the clock where nothing happens on the CA signals. Alternatively, the A-B trigger could be used again by first activating the trigger on the rising or falling edge of the CA signal and then starting the trigger on the next rising edge of the clock. Such techniques are also helpful in cases, where power save features are used in the system. If the controller uses clock gating, there are some dead times where no clock is driven on the bus. It is very likely that the oscilloscope's dead time just ends in such periods, so it would capture the first rising CLK edge after the power down period. At this time, the clock might need some cycles to stabilize and the clock signal will look disturbed. However starting the trigger with the active CA signal allows working with an additional trigger and ensures that the oscilloscope captures only relevant events. Last but not least, it is again possible to use the Zone Trigger to further narrow down the events of interest.

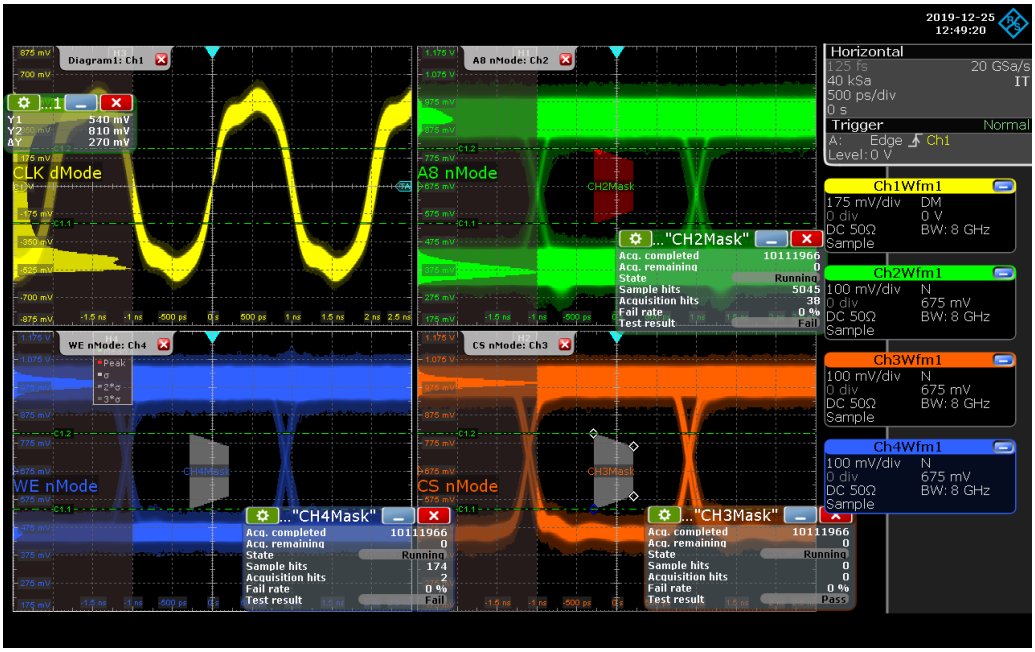

Fig. 3-12: Multiple eye diagrams and included mask tests with the fast acquisition of the R&S® RTP oscilloscope (edge trigger on CLK)

With cursors, mask testing and histograms, further analysis can be done. In this example, the vertical histogram on channel 2 shows that the critical high levels do not have the standard Gaussian level distribution. Instead, there seems to be an amount of accesses that cause non-Gaussian disturbances on the high level.

Recommended signals to start with, are CS and ODT, as these might be required for other investigations as well. Commands like RAS/CAS/WE can simplify the understanding what's going on at the DRAM bus. E.g. viewing CAS and WE as CM (common mode) signals when capturing long traces can help to understand the commands and can help to separate READ from WRITE commands. This trick does not give information on signal integrity, but can give a better indication if the commands are READs or WRITEs. For multi-rank systems it is very important to measure at least one control signal (CS, ODT, CKE) and one CA signal, as the first are per rank and the second connect multiple ranks. Due to different capacitive and inductive loading they will have different Signal Integrity.

3.6 Catching Worst Case Scenarios and Persistence Testing

Bit Error Rates (BER) are well known from serial interfaces. While e.g. PCIe interfaces take a BER of 1e-12 as limit, DRAM designs usually are considering a BER of 1e-16. While serial interfaces have quite sophisticated features to identify transmission errors (e.g. CRC) and request a new transmission to get the correct data, systems will react much more critically on DRAM errors. So, for any critical mission, error correction (e.g. Error Correction Code ECC) should be used to correct errors on the fly. ECC is much less efficient compared to CRC. DDR4 therefore introduced the WRITE CRC, while DDR5 adds the same for the READs.

In DDR systems, a robust design is critical, i.e. Signal Integrity has to be ensured even under worst case conditions. As SSN effects in DDR are highly dynamic and pattern dependent, a huge number of acquisitions is required to catch such worst case scenarios. Thanks to its real-time architecture and the hardware accelerated measurements, R&S® RTP provides a very high acquisition rate, even when deembedding is turned on. What does this mean for practical investigations?

Fig. 3-13 shows a DDR4-2400 device. The CLK frequency is 1200 MHz (clock period of 833ps). The main jitter component is Random Jitter (RJ), visible by its strong Gaussian distribution, the RJ value is 4 ps (RMS). Given the fact that RJ is unbounded, its peak-to-peak value and the related eye diagrams highly depend on the number of acquisitions.

Fig. 3-13: DDR4-2400 device, measurement of the clock jitter with RJ = 4 ps (RMS)

Fig. 3-14 shows two CS (ChipSelect) data eyes in gear-down mode of a DDR4-3200 system with a clock of 1600 MHz. In gear-down mode, the clock frequency stays unchanged, but only every second rising edge of the clock signal is used on the CCA bus. CLK is measured on channel 1 (yellow), CS on channel 4 (blue). The eye is generated by Zone triggering and infinite persistence mode. Even there is no real mask defined in the DDR4 specification, tISH and VIH/L(DC/AC) for a DRAM of this speed is used to visualize the required margins. This clock shows an RMS RJ value of 5ps (left picture, MeasGroup3). Taking the standard calculation of RJ(RMS)*16.4 to get the peak-to-peak RJ for a BER of 1e-16 this results in around 80 ps. For a system running with a UI of 312 ps (3200 MT/s) this is already 25% of the unit interval. The Total peak-to-peak jitter on CS with 10k acquisitions is about 177ps and there is no mask hit. Running the test for 2400kcycles (10 hours) results in 197 mask hits and a total peak-to-peak jitter of 242ps. In this case, the mask violations are not due to timing jitter, but due to voltage noise.

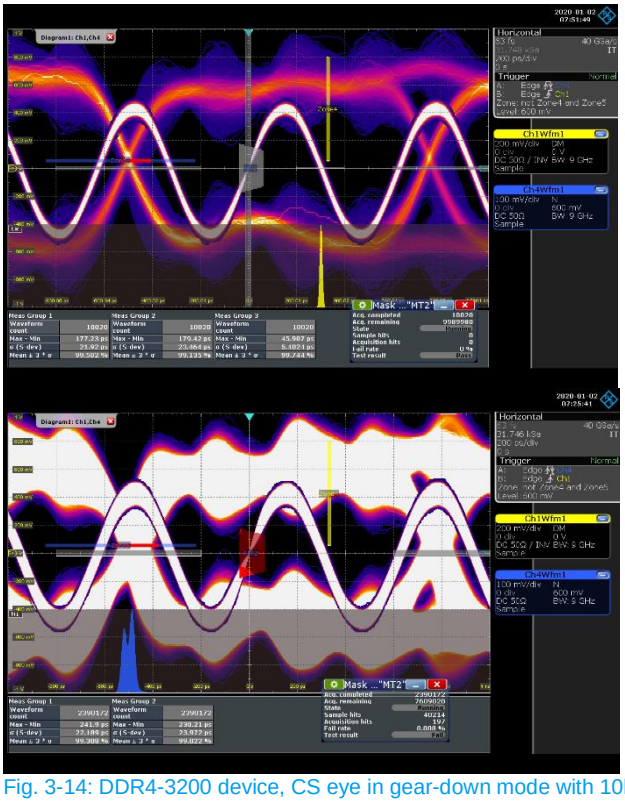

Fig. 3-14: DDR4-3200 device, CS eye in gear-down mode with 10k acquisitions (above) and 2400k acquisitions (below) and a clock RJ of 5 ps (RMS)

The mask is failing with 0.008% probability, far away from any BER of 1e-16. The mask violations however only show a potential risk, but do not necessarily mean that the DRAM captures a wrong signal. The BER of 1e-16 still might be met. In this example, the AC levels are well reached at the beginning of the eye, but are violated at the first reflection right before the sampling point (rising CLK edge). To ensure a stable operation, we highly recommend to improve the design e. g. by optimizing the termination or by applying a slight delay on the clock. Even if the BER of 1e-16 is achieved, the margin might be too low to reliably detect the CS signal.

The above examples already show that capturing a data eye over a long runtime is essential to estimate the robustness of a memory interface. Such statistics can never be generated by any compliance test. Important is to consider not just the runtime for such tests but the number of the real events of interest, captured by the oscilloscope. A high acquisition rate, real-time deembedding and efficient triggering help to catch the maximum amount of events in a given time and easily cut the time to get good statistics by a factor of 10. They help you get the most out of the valuable time you spend for these measurements.

3.7 TDR measurement

TDR / TDT (Time Domain Reflectometry / Time Domain Transmissometry) is an essential measurement in system-level verification and debug. It helps to measure insertion loss, crosstalk, reflections and discontinuities as well as impedances on PCB signal traces. TDR allows an efficient verification that the layout implementation meets the impedance requirement of the memory interface or to identify related problems. Quite often target trace impedances are just calculated, but impedance values are not verified for every manufactured PCB. In such cases, some PCBs can be out of the target impedance range (e.g. 80 Ω differential +/- 10%), causing failures in the memory operation. In these cases, it is required to measure impedances on the real failing board and best even on the failing signal itself. This requires access to the signals and a probe that is suitable for the application (see Fig. 3-15 and Fig. 3-16). It also might be necessary to remove some components (e.g. resistors or the DRAM) to connect to the PCB trace.

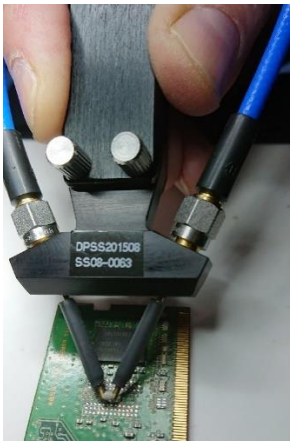

Fig. 3-15: Handheld TDR probe fitting on DRAM ballout (example Packet Micro DP-SS-201508 – 20 GHz, 0.8 mm pitch)

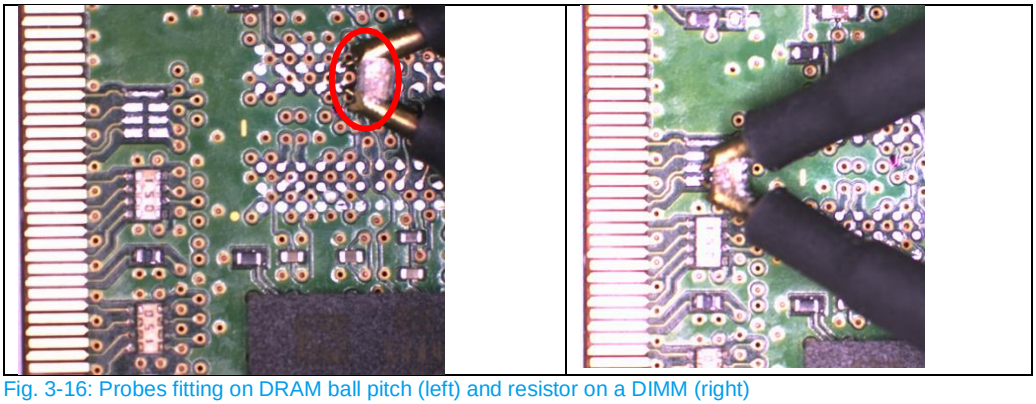

Fig. 3-16: Probes fitting on DRAM ball pitch (left) and resistor on a DIMM (right)

Such TDR check cannot determine whether the PCB stackup is correct, but impedance as the most important parameter can be verified (see Fig. 3-17)

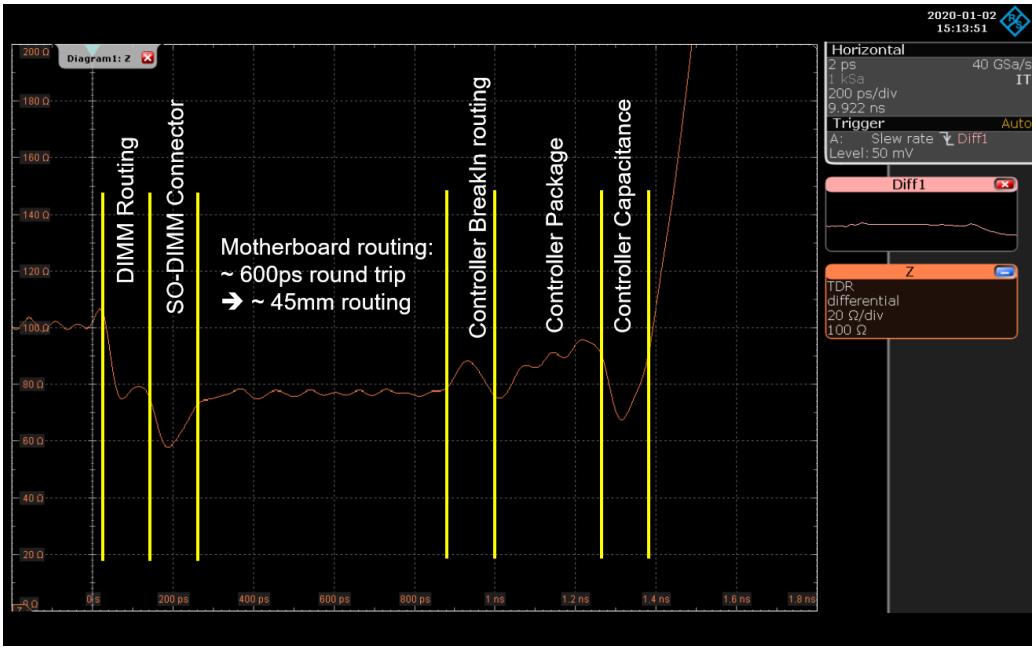

Fig. 3-17: TDR impedance profile for DQS measured from the resistor pad, using R&S® RTP oscilloscope with R&S® RTP-K130 TDR/TDT analysis option

A cross section certainly would have the advantage of showing the exact stackup, but only at one position and at the expense of being a destructive analysis. Measuring the impedance by a TDR allows to get the DUT back to work just by re-soldering the removed components on the PCB.

In addition to the basic impedance also cutouts in the GND plane can be identified. These can occur e.g. due to a layout or manufacturing error.

3.8 Power noise on VDD and Vref

Another very important parameter to be measured is power noise. What seems simple at the first glance is quite difficult if you take a closer look to the specification. VDD and Vref should be measured at the ball of the DRAM. Usually interposers are designed to measure Signal Integrity, but not Power Integrity. Even if the interposer does have connection points for power rails it might add significant noise to the measurement.

A good place for Power Integrity measurements is a capacitor that is as close to the ball as possible. A differential browser module R&S® RT-ZMA30 can be connected to one of the already available modular probes R&S® RT-ZM and used for this measurement (see Fig. 3-18, yellow circle). The probe of choice for this test however is a dedicated power rail probe like R&S® RT-ZPR. This probe includes a built-in probe meter that accurately measures the DC voltage of the power rail. Thanks to the offset compensation, only the power rail disturbances, relative to this DC voltage, are delivered to the oscilloscope and are measured with full resolution. Along with the low inherent noise of the 1:1 power rail probe and the R&S® RTP frontend, even small disturbances can be measured with high accuracy. The probe is directly soldered to the capacitor via a pig tail cable, providing a 2GHz (R&S® RT-ZPR20) or 4GHz (R&S® RT-ZPR40) measurement bandwidth (see Fig. 3-18, red circle).

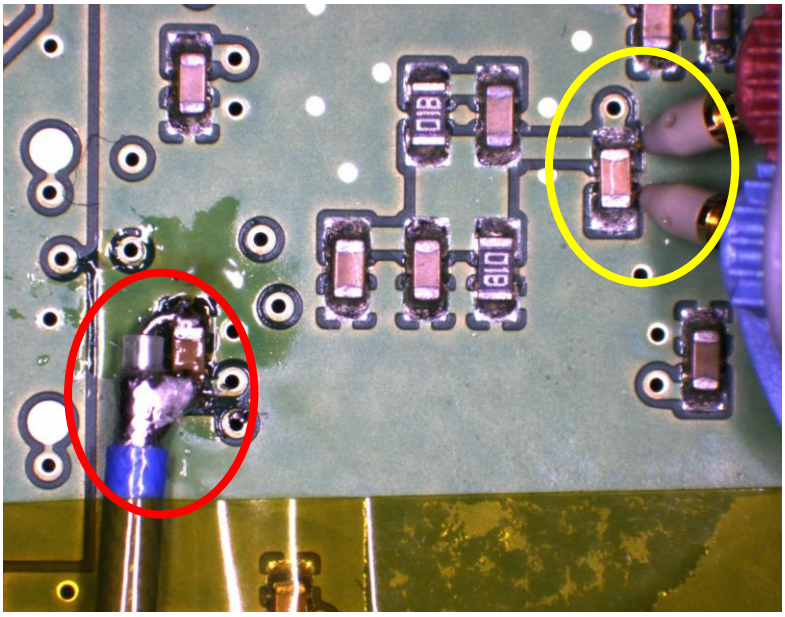

Fig. 3-18: Power supply measurement with a differential probe (yellow) and a dedicated power rail probe (red)

The next challenge is the question, what is “AC” power noise. The specification defines how to measure DC levels for the supplies or Vref, but not the AC Bandwidth for a power supply or Vref noise measurement. Often, just 20 MHz are used for power integrity measurements, but DRAM packages are small and noise up to 400 MHz and above can arrive at the DRAM silicon. It is good practice to measure power noise with sufficient bandwidth as many signal integrity problems have their root cause in poor power integrity.

See Fig. 3-19 for a comparison of power noise measured at 6 GHz bandwidth and 800MHz bandwidth.

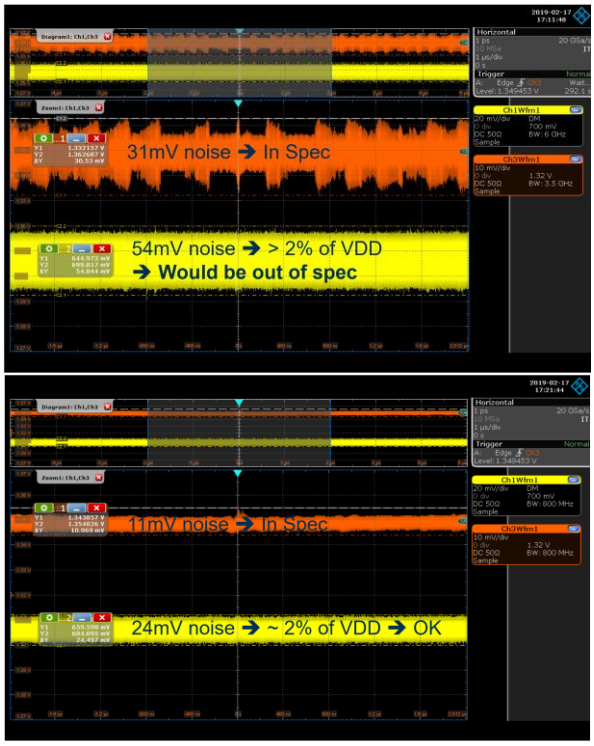

Fig. 3-19: VDD (orange) and Vref (yellow) noise at 6 GHz (above) and 800 MHz (below)

Such measurements are important to understand and identify sources of noise. Even without a formal PASS/FAIL criteria (like in 2-19, the noise in a full bandwidth measurement violates the specification but passes at 800MHz), power supply disturbances at higher frequencies can have a strong impact on the performance of the memory system. Solid engineering judgement is required and it is often helpful to analyze the power rail disturbances also in the frequency domain.

Simultaneous switching noise (SSN) is a particular challenge in DDR systems. When multiple output drivers switch simultaneously, they induce voltage transients on the power rail and the local GND, affecting the power integrity of the system. This is cumulative and increases with the number of simultaneously switching outputs (SSO). In DDR memory designs, this effect is highly dynamic and pattern dependent. Ideally, power noise measurements would be executed with worst-case traffic patterns. However, it is often difficult or even impossible to generate such worst-case patterns, especially if features like DBI (Data Bit Inversion) are used. Using an oscilloscope with a high acquisition rate like R&S® RTP and running the measurement over a longer time increases the chance to catch these worst case scenarios, the related power rail disturbances and their effect on the signal integrity of the design. Also advanced triggering, e.g. on spectrum violations, can be used when trouble shooting power integrity issues.

3.9 Clock test

Checking the clock and particularly clock jitter is another very important task for any debug or verification of the interface. Clock jitter often reveals SSN issues or switching noise of the DC/DC converters. The first step typically is to execute a clock jitter test with the compliance application (see Fig. 3-20).

Test Summary

DDR3L tCK(avg) (12.1.1)

Result	Test	Description	Run
✓	tCK(avg)	Average Clock Period tCK(avg) is calculated from a 200 moving cycle window across the recorded waveform. The measurement is described in chapter 12.1.1. and the limits can be found in chapter 13.1 Table 68 & 69.	18

DDR3L tCK(abs) (12.1.2)

Result	Test	Description	Run
✓	tCK(abs)	Absolute Clock Period tCK(abs) is measuring the minimum and maximum of Clock period. The measurement is described in chapter 12.1.2. and the limits can be found in chapter 13.1 Table 68 & 69.	18

DDR3L tCL(avg) (12.1.3)

Result	Test	Description	Run
✓	tCL(avg)	Average Low Pulse Width tCL(avg) is measuring the average duty cycle of all the negative pulse widths within a window of 200 consecutive cycles. The measurement is described in chapter 12.1.3, and the limits can be found in chapter 13.1 Table 68 & 69.	18

DDR3L tCH(avg) (12.1.3)

Result	Test	Description	Run
✓	tCH(avg)	Average High Pulse Width tCH(avg) is measuring the average duty cycle of all the positive pulse widths within a window of 200 consecutive cycles. The measurement is described in chapter 12.1.3, and the limits can be found in chapter 13.1 Table 68 & 69.	18

DDR3L tJIT(per) (12.1.4)

Result	Test	Description	Run
✗	tJIT(per)	Clock Period Jitter tJIT(per) is measuring the difference between a measured clock period and the average clock period across multiple cycles of the clock. tJIT(per) = min/max of {tCKi - tCK(avg) where i = 1 to 200}. The measurement is described in chapter 12.1.4. and the limits can be found in chapter 13.1 Table 68 & 69.	18

Fig. 3-20: Clock test in a DDR3L design with R&S® RTP-K91 compliance test option: fail in period jitter

This result already gives a first hint about the jitter performance of the system. Especially if any fail occurs, more detailed jitter investigations should be performed, e.g. using tracks, histograms and FFTs. An example is shown in Fig. 3-21.

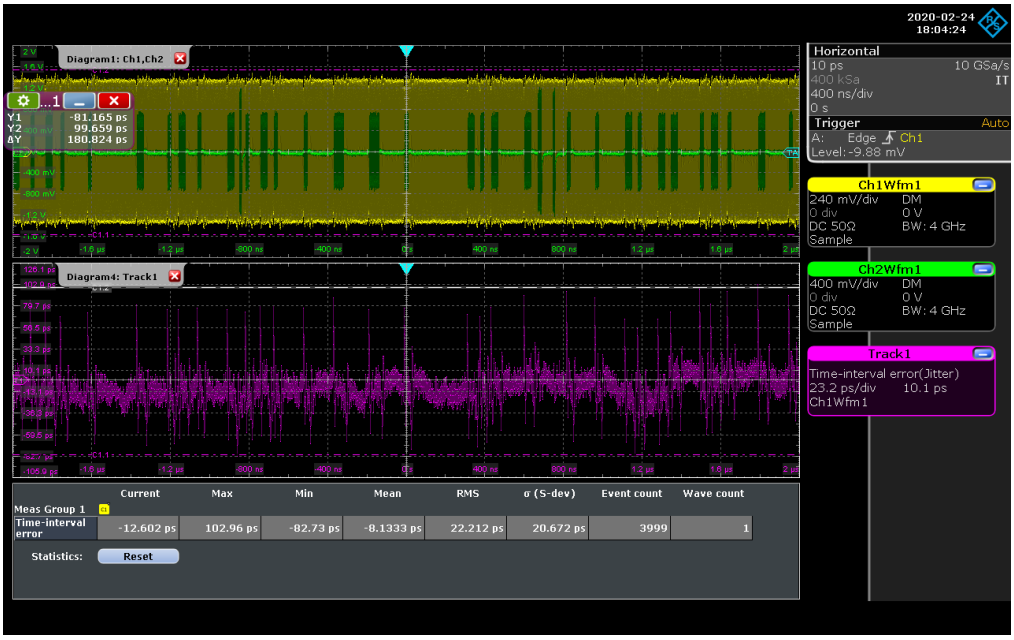

Fig. 3-21: Time Interval Error (TIE) measurement: correlation to data bursts

The yellow trace in the upper diagram represents CLK. Its TIE track for the whole acquisition length is shown in the lower half of the display (purple). All peaks to the high side of this jitter track are exactly correlated to the start of the READ bursts, which can be recognized as the smaller bursts of the DQS signal (green) in the upper half of the screen. Most of the dips down, are aligned to the end of these bursts. For the three higher green bursts (WRITES) there are no peaks in the jitter track.

In this case it is most likely that enabling the termination on the controller causes a VDD dip down on the controller die and the clock is briefly slowed down, causing the positive TIE spike. At the end of the read the termination is disabled and the VDD gets a peak up. With this, the clock runs faster for a short moment, causing the negative TIE spike. In order to verify this theory it would be possible to change termination strength on the controller and check the change in TIE behavior.

In this case again further debugging is necessary. A first step would be to check the external power rail of the memory controller.

Quite often there are also disturbances in clock jitter, that that correlate to the switching noise (switching frequency) of the DC/DC converter or other power noise events. For debugging, it is required to work with the available jitter analysis tools like track, FFT, histogram, etc. to understand deterministic vs. random jitter effects.

4 Conclusion

System level verification and debugging of memory designs needs to be done for every design and needs to be well included in the entire design process. Just performing functional test does not provide enough insight into the system to evaluate margins, e.g. for changing PCB or DRAM vendors.

Compliance test is important to evaluate the signals according to the JEDEC DRAM specification but does not deliver large statistics. Always consider, that the DRAM specification only defines how the CCA signals and the data signals (WRITEs only) need to arrive at the DRAM. For READs it is necessary to check that the signal arriving at the memory controller meets the controller inputs specifications.

In order to get statistics and catch worst case scenarios, eye diagrams with many acquisitions are required. Always consider that it is not just the runtime, but the capture rate that really defines how many events are checked. Ensure to optimize the setup for a high capture rate.

As margins are shrinking (e.g. for DDR4 3200MT/s operation) it is getting more and more important to take a closer look to new DDR DRAM designs and re-perform verification tests on existing system design.

Due to complexity (power-down mode, gear-down mode, CRC, DBI, CA Parity) of the DDR interface, it is getting more and more challenging to catch all important events on the bus. The test result examples documented in this application guide are the result of multiple optimizations of the setup for each measurement. Every engineer working on memory verification or debugging needs to understand:

- the DRAM specification,
- the basic theory about digital data transmission,
- jitter analysis,
- a bit of mathematical background on statistics
- and of course some physics to understand the background of crosstalk, ISI, power noise, insertion loss and other effects.

Solid engineering knowhow and the right tools are required for such an analysis. The R&S® RTP oscilloscope provides a large set of tools for efficient DDR system level verification and debug. Thanks to its unique real-time architecture and high-acquisition rate, measurements can be done with a high statistical certainty, even catching worst case conditions.

5 Ordering Information

Example Configuration

Designation	Type	Quantity	Order No.
High-performance oscilloscope, 8 GHz, 50 Msample memory	R&S®RTP084	1	1320.5007.08
16 GHz differential pulse source	R&S®RTP-B7	1	1333.2001.02
DDR3/DDR3L/LPDDR3 signal integrity debug and compliance test	R&S® RTP-K91	1	1337.8840.02
DDR4/LPDDR4 signal integrity debug and compliance test	R&S® RTP-K93	1	1801.3671.02
Jitter analysis	R&S® RTP-K12	1	1337.8656.02
Zone trigger	R&S® RTP-K19	1	1337.8879.02
Deembedding base option	R&S® RTP-K121	1	1326.3064.02
Realtime deembedding extension	R&S® RTP-K122	1	1326.3070.02
TDR/TDT analysis	R&S® RTP-K130	1	1326.3093.02
9.0 GHz modular probe amplifier, differential, 400 kΩ, multimode	R&S® RT-ZM90	4	1419.3205.02
Probe tip modules	R&S® RT-ZMA14	4	1338.1010.02
4.0 GHz power rail probe, 1:1, low noise, 50 kΩ, large offset range ±60 V	R&S®®RT-ZPR40	1	1800.5406.02

6 Literature

Dokuments:

Efficient eye diagram testing in DDR3/DDR4 system designs

Triggering read and write cycles of DDR3 memories

Realtime deembedding with the R&S®RTP

Optimize differential measurements on high-speed interfaces

Enhancing channel-to-channel alignment for accurate phase-coherent multichannel acquisition

Links:

DDR – double data rate memory

R&S®RTP oscilloscopes

Rohde & Schwarz Probes for oscilloscopes

Training:

EKH - EyeKnowHow "Open the Black box of Memory" DDR training

7 Appendix

Source for memory tests:

Passmark Memtest86 (www.Passmark.com) for x86 based systems

Memtester 4.30 (www.pyropus.ca) for many other platforms

Rohde & Schwarz

The Rohde & Schwarz electronics group offers innovative solutions in the following business fields: test and measurement, broadcast and media, secure communications, cybersecurity, monitoring and network testing. Founded more than 80 years ago, the independent company which is headquartered in Munich, Germany, has an extensive sales and service network with locations in more than 70 countries.

www.rohde-schwarz.com

Certified Quality Management
ISO 9001

Rohde & Schwarz training

www.training.rohde-schwarz.com

PAD-T-M: 3572.7186.00/01.00/EN

Rohde & Schwarz customer support

www.rohde-schwarz.com/support

GFM340 | Version 0e | 10.2020
Application Note | System Level Verification and Debug of DDR3/4 Memory Designs
Data without tolerance limits is not binding | Subject to change

www.rohde-schwarz.com