Testing DDR4 Memory

With Boundary-Scan/JTAG

(2nd ed.)

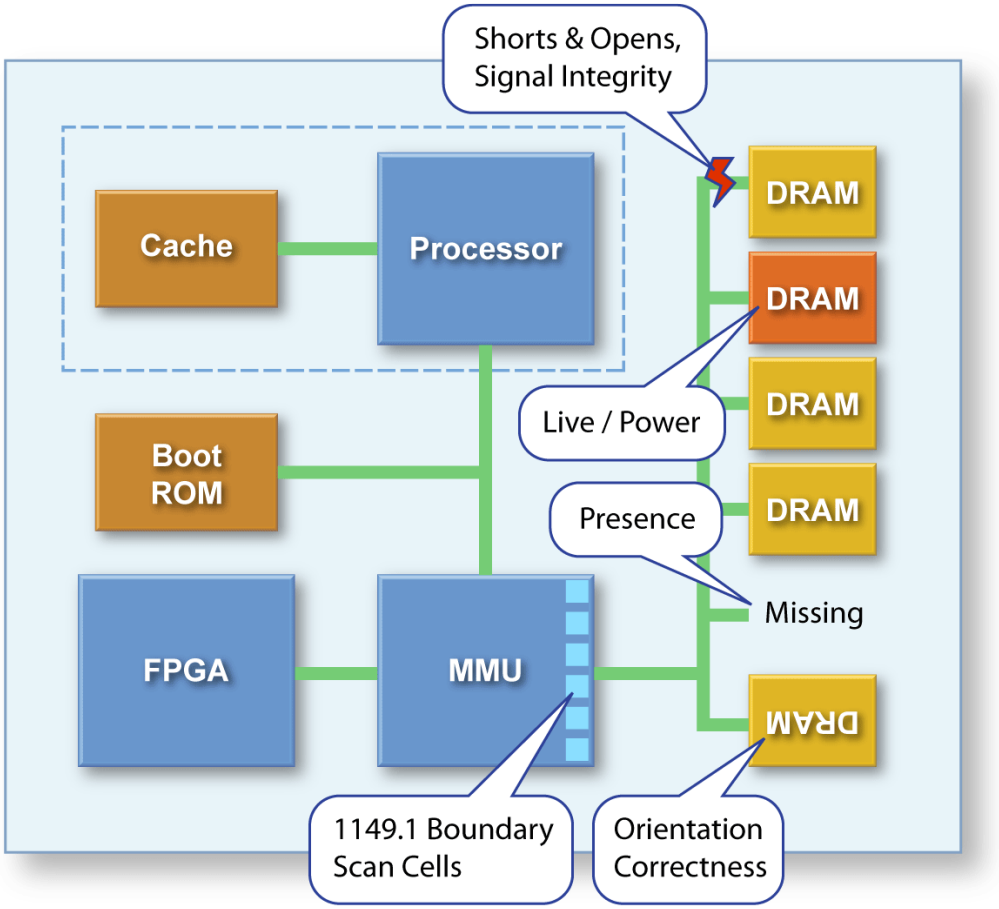

ASSET

Michael R. Johnson – Product Manager

Michael R. Johnson serves as Product Manager for ScanWorks Boundary-Scan Test (BST) for ASSET InterTech, Inc. He also serves as manager of ASSET's Application Engineering organization. As Product Manager, Michael is responsible for providing strategic direction to ASSET's ScanWorks Test Platform by listening to customers and working with cross-functional teams to deliver solutions to meet their needs.

Prior to joining ASSET in 2007 as a Sr. Applications Engineer, Michael's background included roles as a Cellular System Engineer with Nortel Networks and a Hardware Design Engineer with Alcatel USA. While at Alcatel USA, Michael's emphasis was printed circuit board design and layout for high-speed digital cross-connect and fiber-optic systems.

Michael earned a Bachelor of Science degree in Electrical Engineering from Southern University and A&M College located in Baton Rouge, Louisiana, and a Master of Business Administration with an emphasis in Strategic Leadership from Amberton University located in Garland, Texas.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

A survey by the International Electronics Manufacturing Initiative (iNEMI) asked test engineers in the electronics industry what were their biggest problems with testing circuit boards. Of the 11 possible problems listed, characterizing and testing memory soldered to circuit boards were among the top three. Memory test was at the top of the list of prevalent problems along with 'loss of access to test points' and 'the need to perform debug/diagnostics on board failures.' When asked which type of built-in self-test (BIST) instruments would solve the engineer's problem, memory BIST was rated the second most needed, virtually tied with BIST instruments for validating high-speed I/O buses. Clearly, the ability to thoroughly test, characterize and diagnose problems with soldered-down memory is one of the most pressing problems in the industry. Using Double Data Rate Fourth Generation Synchronous Dynamic Random-Access Memory (DDR4 SDRAM) memory as an example, this eBook discusses how boundary-scan test (BST) methods based on the IEEE 1149.1 standard, including the built-in Connectivity Test (CT) of DDR4 SDRAM memories and general-purpose Memory Access Verification (MAV), can be used to test and diagnose soldered-down memory devices (not to preclude use also for socketed modules, when applicable). It is assumed that a BST tool is being used to test the DDR4 SDRAM memory.

Note that circuit board test engineers participated in the iNEMI survey mentioned above. In most cases, board test engineers assume that the memory devices themselves are not causing a failure since the memory chips are tested and qualified before they are assembled on a board. As a result, a memory test failure in many cases will indicate a failure in the connectivity channel to/from the memory. For some time now, the decreasing level of test point access on circuit boards has made BST the most effective method for acquiring pin-level diagnostic on structural faults early in the production flow. BST-based CT and/or MAV can perform thorough tests on the interconnects linking memory management units (MMU), field programmable gate arrays (FPGA), and processors to memory devices. Later in the test process, functional test of some sort can be employed to ensure the quality of the assembly. The 1st edition of this eBook described DDR4 testing using CT and/or MAV actions in the production environment. For this 2nd edition, content has been added which provides an overview of DDR4 device structure, commands, and operation sequences.

Testing Memories at Every Step in the Product Life Sycle

In the broadest sense, memory testing takes place over the entire life cycle of a system, beginning with board design/development, moving into production, and culminating in post-production stages such as field service. Eventually, the cycle begins again when memory test is performed during design/development for the next generation of the board design. During each phase in the life cycle, the objectives and goals of memory test differ and the memory test process itself is typically referred to differently, according to the objectives of that phase.

During system design and new product introduction (NPI), testing memory in a timely fashion is particularly critical if the new system is to be delivered promptly to the marketplace. One of the critical points in the NPI process is the production of prototypes of the circuit board design. These prototypes must be evaluated for structural faults quickly so that board bring-up and development feedback can be completed effectively. Structural memory tests will be performed during the initial phase of board bring-up to identify the root cause(s) of faults, failures or errors in the design as well as identifying problems in the early production equipment. The intent of these tests is to correct the problems prior to the design's release to high-volume production. When memory structures on early prototypes are found to be functioning properly, functional software can be executed easily, and the functional system can be debugged sooner. The best way to quickly transition into functional test is to first identify any structural defects early in the process. This is where boundary-scan excels. Products can then flow through design faster and move into full production, shortening the product's time-to-market. Various delays and a longer-than-expected board bring-up phase could jeopardize revenues.

Once the NPI/board bring-up phase has been completed, the design is ready to transition into manufacturing. At this point, the nature of memory tests is the same, but the faults found may differ considerably. By the time design and board bring-up testing have achieved a 'known good board,' manufacturing test engineers assume that any errors in the design's memory architecture have already been corrected because the design has been released to production. Consequently, the memory tests performed during production are intended to determine whether individual circuit boards are ready to be released to users, not to identify any underlying errors in the memory architecture itself. Instead, production testing involves determining whether the manufacturing and assembly processes have added

any defects to the circuit boards. These types of memory tests with shorter test application times will be essentially go/no-go tests that can be applied quickly to keep pace with a predetermined beat rate on the production line.

When the rate of production has achieved its predetermined beat rate, boards that fail due to systemic defects must be debugged and diagnosed so that the root cause or source issue in the manufacturing process can be understood and corrected. This will improve manufacturing yields and increase the volume of product available to the market. In most cases, a test that uncovered a fault that affected manufacturing yields is retained as part of the manufacturing test suite so that this condition can be quickly identified should it occur again.

During post-production phases of the life cycle – that is, when systems have been sold and are installed in the field – memory tests are performed by repair personnel to troubleshoot malfunctioning systems and maintain user satisfaction. The two main goals during this phase are: 1) identify any reliability concerns such as memory chips or board structures that are failing sooner than expected; and 2) identify changes in the board design or component selection that may make the system better suited for deployment to the market.

Structural test techniques like boundary-scan can be quite effective during every phase in a product's life cycle for capturing faults as early and as quickly as possible.

DDR4 Structure and Basics

Before a discussion of the methods of testing DDR4 SDRAM devices can take place, an overview of the basic structure and operations of these devices is in order. As DDR3 reached its limits in terms of bandwidth and data processing speed, DDR4 has evolved as the next generation SDRAM. DDR4 delivers higher performance, higher Dual In-line Memory Module (DIMM) capacities, increased bandwidth, improved data integrity all while lowering overall power consumption as compared to DDR3.

DDR4 provides up to 50 percent increased performance and bandwidth while decreasing power consumption. This represents a significant improvement over DDR3 and a power savings up to 40 percent. In addition to optimized performance, DDR4 also provides cyclic redundancy checks (CRC) for improved data reliability, on-chip parity detection for integrity verification of 'command and address' transfers over a link, enhanced signal integrity and other robust Row Address Strobe (RAS) features.

In terms of speed, DDR4, doubles the top-end data transfer speed of DDR3 (Figure 1), and consumes less power. DDR4 devices can perform 3,200 million transfers per second (MT/s) with a 1.2 V power supply. The expected data transfer rate of DDR4 is 4,266 MT/s. DDR3's top transfer speed is 2,133 MT/s with a 1.65 V supply.

Figure 1: DDR4 peak bandwidth vs. data rate

DDR4 devices are beginning to appear in electronic devices such as smartphones, tablet or desktop computers, and designs based on system-on-a-chip (SoC) devices, such as the Zynq UltraScale+. The essential IO physical structure of a DDR4 device is displayed in Figure 2.

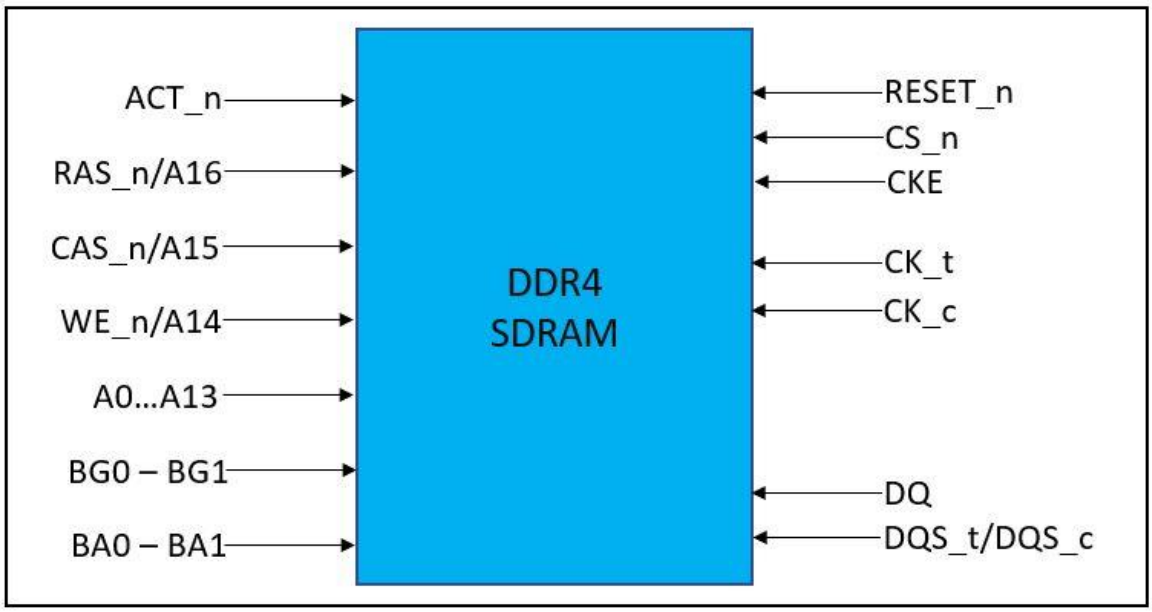

Figure 2: DDR4 top level

Each DDR4 device has a clock, reset, chip-select, address and data pins. Table 1 has more detail about each pin and its function. Please note that the table below is not a complete list of all DDR4 IOs.

Symbol	Type	Function
RESET_n	Input	DRAM is active when this signal is HIGH.
CS_n	Input	The memory looks at all the other inputs only if this signal is LOW.
CKE	Input	Clock Enable. HIGH activates internal clock signals and device input buffers and output drivers.
CK_t/CK_c	Input	Differential clock inputs. All address & control signals are sampled at the crossing of posedge of CK_t and negedge of CK_n.
DQ/DQS	Inout	Data Bus & Data Strobe. This is how data is written in and read out. The strobe is essentially a data valid flag.
RAS_n/A16 CAS_n/A15 WE_n/A14	Input	These are dual function inputs. When ACT_n & CS_n are LOW, these are interpreted as Row Address Bits. When ACT_n is HIGH, these are interpreted as command pins to indicate READ, WRITE or other commands.
ACT_n	Input	Activate command input
BG0-1 BA0-1	Input	Bank Group, Bank Address
A0-13	Input	Address inputs

Table 1: Pin function information

Bank Group, Bank, Row, and Column

The top-level picture displays the external DDR4 pin structure. The internal memory structure of DDR4 is organized in Bank Groups and Banks.

To initiate a READ from DDR4 memory, an address is provided along with additional data. The address provided is called the "logical address". The logical address is translated to a physical address before it is presented to the DDR4. The physical address is made up of the following fields:

- Bank Group
- Bank
- Row
- Column

These individual fields are then used to identify the exact location in the memory to read-from or write-to. Once the Bank Group and Bank have been identified, the Row part of the address activates a line in the memory array. This is called the "Word Line" and activating it reads data from the memory array into something called "Sense Amplifiers". The Column address then reads out a part of the word that was loaded into the Sense Amplifiers. The width of the column is called the "Bit Line".

The width of a column is standard - it is either 4 bits, 8 bits or 16 bits wide, therefore, DRAMs are classified as x4, x8 or x16 based on this column width. Another thing to note is that the width of DQ data bus is same as the column width. Therefore, DRAMs are classified based on the width of the DQ bus.

DRAM Sizing & Addressing

DRAMs come in standard sizes as specified by the Joint Electron Device Engineering Council (JEDEC). JEDEC is the standards committee creates the design and roadmap of DDR memories. The following information was obtained from section 2.7 of the DDR4 JEDEC specification (JESD79-4B). (Table 2).

2 Gb Addressing Table

Configuration		512 Mb x 4	256 Mb x8	128 Mb x16
Bank Address	# of Bank Groups	4	4	2
	BG Address	BG0~BG1	BG0~BG1	BG0
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
Row Address		A0~A14	A0~A13	A0~A13
Column Address		A0~A9	A0~A9	A0~A9
Page Size		512B	1KB	2KB

4 Gb Addressing Table

Configuration		1 Gb x 4	512 Mb x8	256 Mb x16
Bank Address	# of Bank Groups	4	4	2
	BG Address	BG0~BG1	BG0~BG1	BG0
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
Row Address		A0~A15	A0~A14	A0~A14
Column Address		A0~A9	A0~A9	A0~A9
Page Size		512B	1KB	2KB

8 Gb Addressing Table

Configuration		2 Gb x 4	1 Gb x8	512 Mb x16
Bank Address	# of Bank Groups	4	4	2
	BG Address	BG0~BG1	BG0~BG1	BG0
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
Row Address		A0~A16	A0~A15	A0~A15
Column Address		A0~A9	A0~A9	A0~A9
Page Size		512B	1KB	2KB

16 Gb Addressing Table

Configuration		4 Gb x 4	2 Gb x8	1 Gb x16
Bank Address	# of Bank Groups	4	4	2
	BG Address	BG0~BG1	BG0~BG1	BG0
	Bank Address in a BG	BA0~BA1	BA0~BA1	BA0~BA1
Row Address		A0~A17	A0~A16	A0~A16
Column Address		A0~A9	A0~A9	A0~A9
Page Size		512B	1KB	2KB

Table 2: DRAM addressing

DRAM Page Size

In the DRAM Addressing table, there is a mention of Page Size. Page size is the number of bits per row. Or it is the number of bits loaded into the Sense Amplifiers when a row is activated. Since the column address is 10 bits wide, there are 1K bit-lines per row. So, for a x4 device number of bits is 1K x 4 = 4K bits (or 512B). Similarly, for x8 device it is 1KB and for x16 it is 2KB per page.

Accessing Memory

- Read and write operations to the DDR4 are burst oriented. The operations begin at a selected location (as specified by the user provided address) and continues for a burst length of eight or a 'chopped' burst of four.
- Read and write operations are a 2-step process. The operations begin with the ACTIVATE Command (ACT_n & CS_n are made LOW for a clock cycle), which is then followed by a RD or WR command.
- The address bits registered coincident with the ACTIVATE Command are used to select the Bank Group, Bank and Row to be activated (BG0-BG1 in x4/8 and BG0 in x16 selects the Bank Group; BA0-BA1 select the Bank; A0-A17 select the row). This step is also called Row Address Strobe (RAS).
- The address bits registered coincident with the Read or Write command are used to select the starting column location for the burst operation. This step is also referred to as Column Address Strobe (CAS).
- Each Bank has only one set of Sense Amplifiers. Before a read/write to a different row in the same Bank can be performed, the current open row must be de-activated using a PRECHARGE command. PRECHARGE causes the data in the Sense Amplifiers to be written back into the row.
- Instead of issuing an explicit PRECHARGE command to deactivate a row, the Read with Auto-Precharge (RDA) and Write with Auto-Precharge (WRA) commands can be used. These commands tell the DRAM to automatically deactivate/precharge the row once the read or write operation is complete. Since column address uses only address bits A0-A9, A10 which is an unused bit during CAS is overloaded to indicate Auto-Precharge.

Command Truth Table

To issue commands to the DRAM such as READ, WRITE, ACTIVATE, PRECHARGE, the ACT_n, RAS_n, CAS_n & WE_n inputs are interpreted as commands based on the truth table (Table 3).

Function	Shortcode	CS_n	ACT_n	RAS_n/A16	CAS_n/A15	WE_n/A4	A10/AP
Refresh	REF	L	H	L	L	H	H or L
Single Bank Precharge	PRE	L	H	L	H	L	L
Bank Activate	ACT	L	L	Row Address			
Write	WR	L	H	H	L	L	L
Write with Auto-Precharge	WRA	L	H	H	L	L	H
Read	RD	L	H	H	L	H	L
Read with Auto-Precharge	RDA	L	H	H	L	H	H

Table 3: Partial DDR4 truth table

The table above is only a subset of commands you can issue to the DRAM. The entire DDR4 command truth table is specified in section 4.1 of the JEDEC specification JESD79-4B.

Read Operation

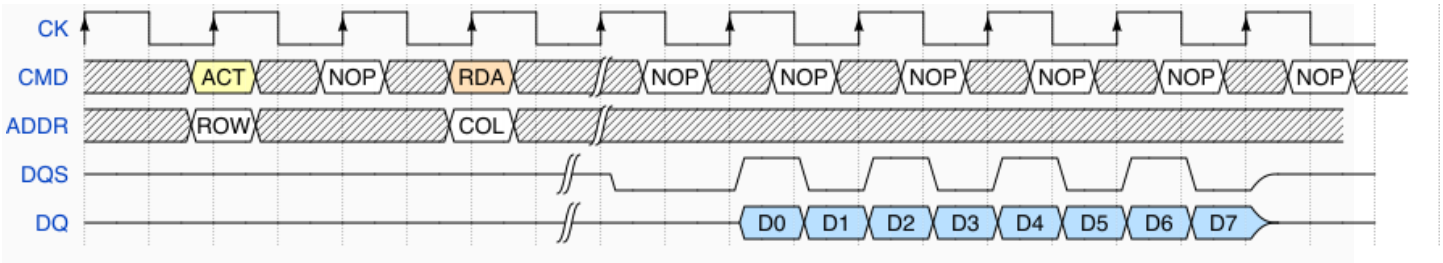

Figure 3: Read operation

Figure 3 shows the timing diagram of a READ operation with burst length of 8 (BL8).

- The first step is an ACT command. The value on the address bus indicates the row address.
- The second step in an RDA command. The value on the address bus is the column address.
- The RDA command tells the DRAM to automatically PRECHARGEs the Bank after the read is complete

Write Operation

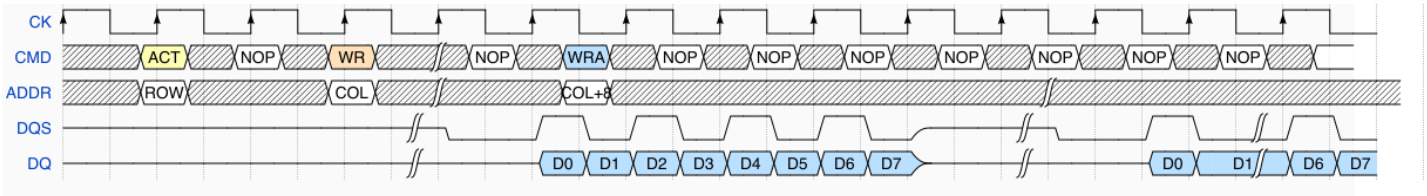

Figure 4: Write operation

Figure 4 shows the timing diagram of a WRITE operation.

- The first step activates a row.
- Then 2 WRITE commands are issued. The first one to address COL and second one to COL+8.
- The second write operation does not need an ACT before it because the row we intend to write to is already active in the Sense Amplifiers.
- Also note that the first command is a plain WR, so this leaves the row active. The second command is a WRA which de-activates the row after the write completes.

A16, A15 & A14 are not the only address bits with dual function. The auto precharge command is issued via A10 and select BurstChop4 (BC4) or BurstLength8 (BL8) mode is selected via A12, if enabled in the mode register.

Operation Summary

As a summary, before continuing to CT and MAV DDR4 testing methods:

- The DRAM is organized as Bank Groups, Bank, Row & Columns
- The address issued by the user is called the Logical Address and it is converted to a Physical Address by the DRAM controller, before it presented to the memory
- DDR4 DRAMs are classified as x4, x8 or x16 based on the width of the DQ data bus
- DDR4 DRAMs can depth cascaded, or width cascaded, to achieve the required size
- Read and write operations are a 2-step process. The 1st step activates a row, while the 2nd step reads or write to the memory.

Connectivity Test (CT) of DDR4 memories

One method for conducting structural testing for shorts and opens on interconnects between boundary-scan devices and soldered down DDR4 memory is with a built-in test mode known as Connectivity Test (CT). When the CT mode is activated, the DDR4 memory device pins are divided into a set of inputs and a set of outputs. A user can check for shorts and opens on interconnects with a boundary-scan tool by sending a known test stimulus to the input pins and checking for a pre-determined stimulus on the output pins. The expected output stimulus is defined by the Joint Electron Device Engineering Council (JEDEC®) and the standard number is JESD-79-4 (currently in Revision B).

The CT mode is enabled through the assertion of an external pin, labeled as Test Enable (TEN). The TEN pin activates the built-in test mode and asynchronous logic gates between the input pins and the output pins. When CT mode is enabled, Boolean exclusive OR (XOR) functions are invoked at pins of the memory device, assigning some pins as CT inputs and others as CT outputs. In this way, this test function of the DDR4 device is fully static and completely independent of the content and control of the memory array.

Of course, the TEN pin must be 1) present on the DDR4 memory device and 2) it must be accessible to a boundary-scan cell or by some other external means for assertion. Therefore, in order to make use of the CT method, PCB design-for-test must be considered to make the TEN pin accessible.

The organization of the internal logic is different depending on the size of the DDR4 memory device (e.g., x4, x8, or x16). To present an example of how the CT operates, we will analyze a simple case using a DDR4 memory device, size x16. In this case:

Output Data 0 (DQ0) is a function of three inputs (A1, A6, PAR):

$$DQ0 = MT0 \qquad \textit{where} \qquad MT0 = XOR\ (A1, A6, PAR)$$

This means that when there is an odd count of 1s across the three inputs of the XOR gate, the output will go to 1; otherwise, when there is an even count of 1s, the output will go to 0. See Figure 5 and Table 4.

Figure 5: 3-input XOR gate

Exclusive-OR Gate - Truth Table			
A (A1)	B (A6)	C (PAR)	Y (DQ0)
0	0	0	0
1	0	0	1
0	1	0	1
1	1	0	0
0	0	1	1
1	0	1	0
0	1	1	0
1	1	1	1

Table 4: 3-input XOR truth table

The boundary-scan function within the memory controller, which is typically a chip like an FPGA, ASIC, or microprocessor, can be used to stimulate the CT inputs and monitor the CT outputs. Then, using the XOR truth table for evaluating the outputs as a function of the inputs, the interconnects between the memory and the boundary-scan device can be easily tested by the JTAG tool and diagnosed down to net- and pin-level.

The CT mode is invaluable in prototype bring-up when systems do not boot, or you have other issues that are difficult to diagnose with traditional functional methods. CT mode is also a very fast and accurate test to use during volume manufacturing, since the application of the boundary-scan vectors is very fast, and diagnosis is very accurate and immediate. Figure 6 presents a schematic drawing of a memory interconnect being tested through DDR4 CT mode.

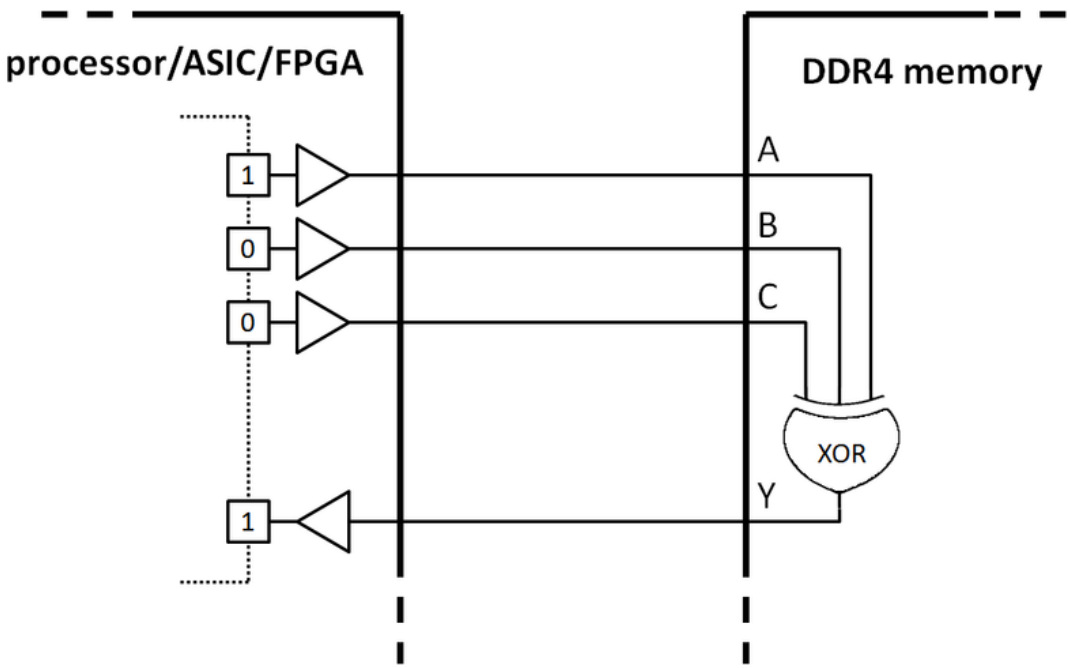

Figure 6: Schematic drawing of a memory interconnect test through DDR4 CT mode

While the CT mode provides highly effective and efficient testing of shorts and opens, it may not always be applicable, such as due to inaccessibility to TEN pin, as described above. Further, the TEN pin was not mandated to be operable for all DDR4 device types, depending on capacity and organization. As such, a generalized means of boundary-scan test for DDR4 of any capacity and for all organizations (x4, x8 and x16) and regardless of whether the TEN pin is operable or accessible, may be required.

What is a Memory Access Verification (MAV) Test?

Like conventional memory tests, a Memory Access Verification (MAV) test is based on reading from and writing to memory. By contrast, as boundary-scan tests, MAV are directed at structural coverage and yet are still based on pseudo-functional algorithms. In fact, these same types of tests are used by many memory device suppliers to test their memory chips and/or logic chips that contain embedded memory. It must be noted that board test requirements for memory testing are less rigorous than chip test requirements. As a result, board memory test algorithms are generally simpler and require fewer operations.

For example, a simple board-level interconnect test may be accomplished by writing and reading the physical equivalent of “5s” and then “As” on the data bus, simultaneous with applying these same values to select locations representing 5 and A on the physical address lines (true board routes may not be in “logical” order but may be physically scrambled). This ensures that all routes are stressed with a checkerboard pattern consisting of 1s surrounded by 0s and 0s surrounded by 1s. Other more advanced memory interconnect tests may use an algorithm like Wagner, which will achieve high fault coverage while reducing test execution time and the number of test vectors applied. (NOTE: Simply stated, the Wagner algorithm assigns a unique ID with minimal length (number of bits) to each net to be tested and uses each such IDs to express the states of each net across an equal number of test steps such that it can be assured that any given net, with respect to any other, will be guaranteed to differ on at least one step where said net has state logic one and at least one step where said net has state logic zero.)

Boundary-Scan Testing of DDR4 Memories

Although memory chips typically do not conform to the boundary-scan standard (IEEE 1149.1 JTAG) – meaning they do not have their own IEEE 1149.1 Test Access Port (TAP) or dedicated boundary-scan registers on chip – they can be tested from the boundary-scan facilities of a connected device, such as an MMU, microprocessor or an FPGA, as long as power and a clock are present on the interconnects and access to the board's IEEE 1149.1 TAP is provided. The boundary-scan registers on a MMU can be appropriated and directed to test the shorts and opens on the interconnect routes to and from memories. In some cases, BST will be the only alternative to test these interconnects. For example, bed-of-nails fixtures that are essential to in-circuit test (ICT) systems frequently do not have access to the memories because there are no test pads or just a few on the board. Moreover, functional tests for the board are typically not yet available when board bring-up is being performed on prototypes.

When the boundary-scan registers on a device that is connected to the memory device are used to test memory interconnects, either CT (when the TEN pin is operable and accessible) or MAV (more generally) are used. The CT, being directed expressly at coverage of connectivity faults, is highly efficient for this purpose. MAV, on the other hand, as described above, employs algorithmic writes and reads that are executed on memory locations in particular sequences.

Naturally, the more complex the test algorithms, the longer the testing will take because a boundary-scan 'scan operation' (ScanDR) must be conducted for each read or write. In many cases, when multiple Read, Write, Stall, Output_Enable and other control signals must be operated at different times during a test, then multiple ScanDRs must be executed to operate those signals in the proper sequence for each read or write action. Fortunately, applying Wagner patterns significantly reduces the number of patterns BST requires to test the memory, reducing to a matter of seconds the test execution time on interconnects to memories of any size.

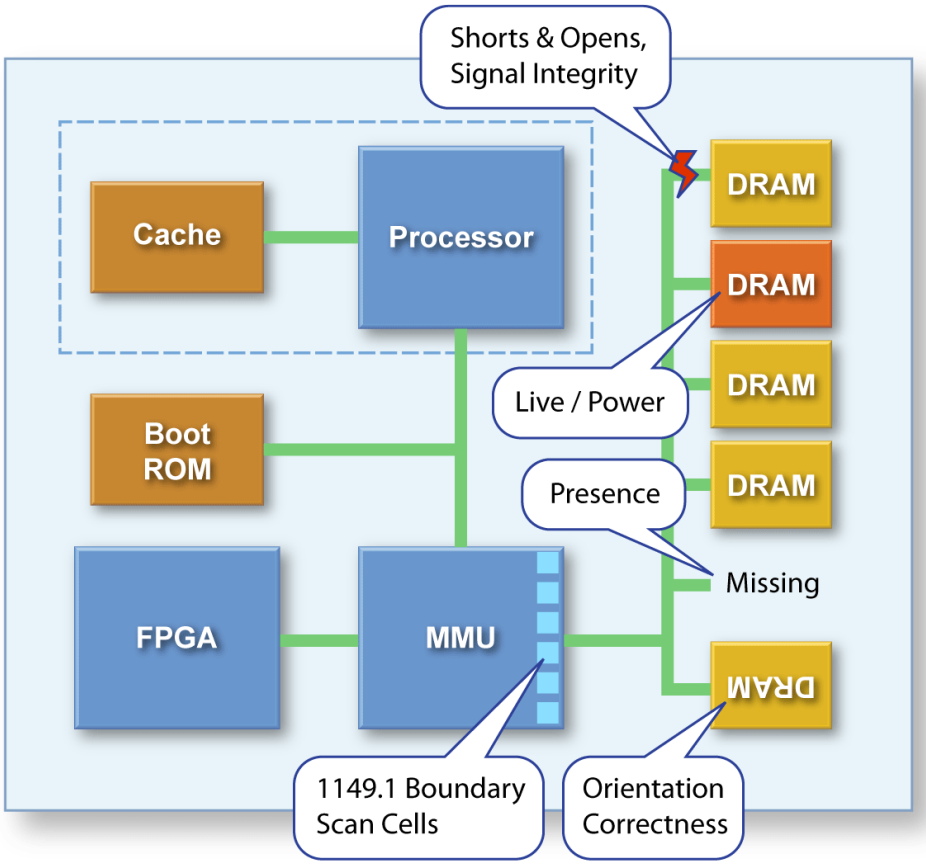

Figure 7: Example board with memory architecture

Figure 7 shows that a MMU with a boundary-scan Test Access Port (TAP) can test connected memory devices for specific faults, such as those defined by the PCOL/SO board test fault spectrum. (PCOLA/SOQ is a fault spectrum originally defined by Agilent Technologies and now adopted and supported by iNEMI. For a description, see ASSET Test Data Out blog post, "Ever heard of PCOLA/SOQ/FAM?" 7 illustrates a case where shorts or opens have corrupted the interconnect to the first DRAM device, the next DRAM does not have power, another DRAM is missing and the last memory device is the wrong type of memory and its orientation is incorrect (upside down). All of these fault types are included in the PCOLA/SOQ fault spectrum definition. (Note that the "A" and "Q" are inspection items and not covered by boundary-scan electrical tests.)

The rate and efficiency at which boundary-scan tests can be applied to memory buses are limited by the clock speed of the boundary-scan resources in the chips on the board and the length of a particular

device's boundary-scan register. For example, 5 MHz may be the maximum boundary-scan clock speed of the slowest device on the JTAG scan chain. This limits the speed of the entire chain to 5 MHz. In Figure 7, the capabilities of the MMU would dictate these factors. Moreover, multiple scans typically comprise a memory test. When each scan can only be applied at a slow speed, the time needed to apply an entire memory test made up of multiple scans is affected. Fortunately, smart test patterns, like the Wagner algorithm, effectively address these issues so that reasonable test times are achieved by applying boundary-scan tests to memory interconnects. Further still, when CT is available and accessible, the number of scans required is just a small fraction of those required of other methods.

Testing DDR4 with a Boundary-Scan Tool

Every BST tool is different. The functionality and requirements described below may or may not apply to a particular BST tool. In fact, certain less powerful tools may not be capable of some of the functionality explained here.

For a BST tool to test DDR4 memory, it must be provided certain information, including the following:

- Structure and composition of the boundary-scan chain
- The order of the boundary-scan devices on the scan chain. This can be entered manually by the engineer or some tools are able to automatically extract this information from the board design's netlist.
- The Boundary-Scan Description Language (BSDL) file for each device on the chain. BSDL files are usually provided by the chip supplier and can be downloaded from the supplier's web site.
- How the DDR4 memory is connected to the boundary-scan device from which boundary-scan tests will be launched. Some tools can extract this information from the board's netlist, which often can be imported by the boundary-scan tool from a computer-aided design (CAD) tool.
- Model(s) of the DDR4 device(s) to be tested with boundary-scan. Such models are required to associate physical pins to logical signals and to describe the essential algorithms for initialization, stimulus and response, as applicable the chosen test type, whether CT or MAV.

Figure 8 shows the general architecture of a circuit board with a bank of DDR4 memory and the buses that connect it to a boundary-scan device.

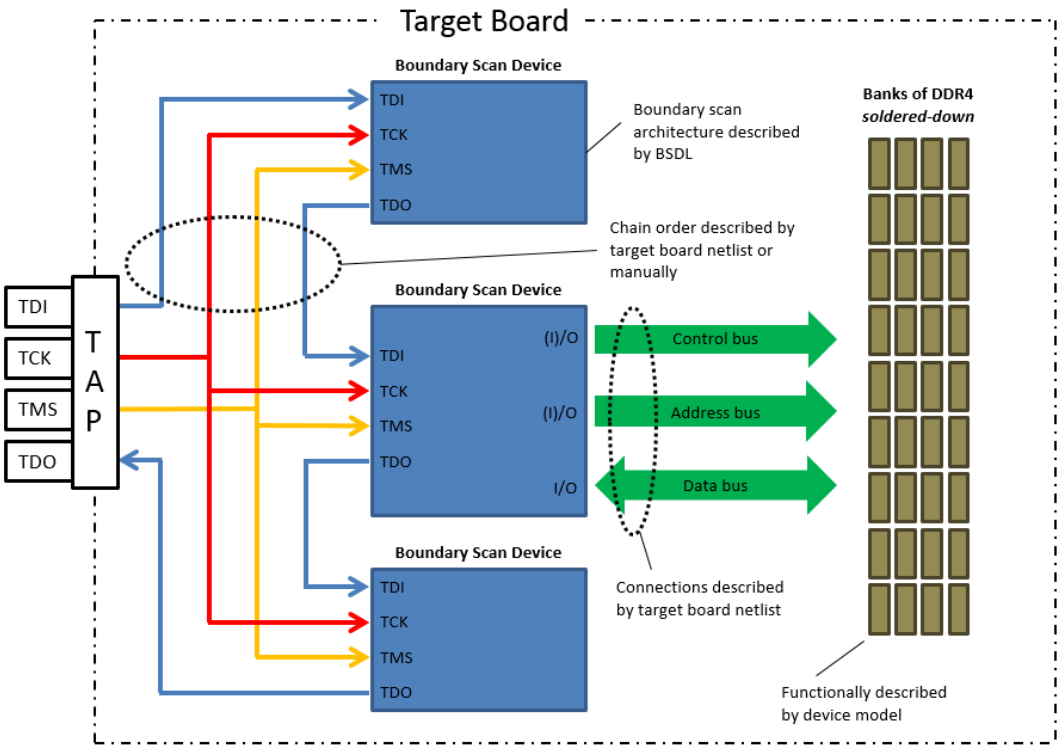

Figure 8: General architecture of a circuit board with memory

On many board designs the connection between a boundary-scan device and the DDR4 memory array will also contain non-boundary-scan devices such as series resistors, buffers and other types of devices. When this is the case, information on these so-called 'cluster' devices must be provided to the boundary-scan test tool. Model-based boundary-scan test tools capture the needed information on these cluster devices in device models. The more accomplished suppliers of boundary-scan test tools provide these models to users in a cluster device library. Advanced tools can then automatically generate memory test sequences and patterns using boundary-scan actions.

Summary and Conclusions

Modern circuit boards have complex memory architectures that are harder to test because of their high speeds, the increasing frequency of data transfers over memory buses, the escalating complexity of communications protocols across these buses, the disappearance of test points on circuit boards and the fact that test probes may cause unwanted signal integrity issues when they are placed on an interconnect. Memory test and validation coverage from intrusive probe-based methods such as oscilloscopes, flying probe testers and ICT systems is rapidly eroding. Frankly, these legacy test methods are quite challenged by today's aggressive test goals.

Non-intrusive memory test methodologies such as BST-based CT and/or MAV can effectively replace most of these legacy intrusive methods. In addition, CT and/or MAV can function as an effective bridge into the non-intrusive world of test, validation, and debug technologies. Today's highly automated boundary-scan tools can be applied through cost-effective, compact, and standalone or integrated testers that provide comprehensive fault coverage, including production-related memory diagnostics that isolate faults to pin or net levels in a matter of seconds.

In addition to the memory test coverage from BST, other non-intrusive memory test methodologies can complement this coverage, including processor-controlled test (PCT), processor-based functional test of DDR (PFTDDR), and FPGA-controlled test (FCT). Separately or together these methodologies have reached the point where they can exceed the memory test coverage provided by legacy intrusive methods. These non-intrusive technologies are available as automated software tools that can be applied through cost-effective, compact, and standalone testers. When these non-intrusive memory test technologies are deployed together, the coverage derived from each complements the others and, taken together, they can comprise a memory test toolkit capable of meeting the requirements of any test strategy. Because these non-intrusive test technologies are software-based, they are extremely flexible and, as a result, can be targeted to the specific characteristics and restrictions of a particular board design. Combining the strengths of all these non-intrusive test methods creates a powerful and versatile memory test platform that solves the test problems created by today's high-speed memory and memory architectures.

References

1. *DDR SDRAM* – Retrieved from: https://en.wikipedia.org/wiki/DDR4_SDRAM

2. *DDR4 SDRAM.* Retrieved from: https://www.systemverilog.io/ddr4-basics#physical-structure

3. *What is DDR4 RAM?*. Retrieved from: https://www.lenovo.com/us/en/faqs/pc-life-faqs/what-is-ddr4-ram/

4. *DDR4 Memory Standard.* Retrieved from: https://www.kingston.com/unitedstates/us/memory/ddr4-overview

5. *DesignWare Technical Bulletin – DDR4 Bank Groups in Embedded Applications.* Retrieved from: https://www.synopsys.com/designware-ip/technical-bulletin/ddr4-bank-groups.html