

Integrating and Operating HBM2E Memory

Achieve Maximum Performance From the Highest-Bandwidth DRAM in Micron’s Ultra-Bandwidth Solutions

Introducing HBM2E

High-bandwidth memory (HBM) is the fastest DRAM on the planet, designed for applications that demand the maximum bandwidth between memory and processing. By tightly integrating through-silicon-via (TSV) stacked memory die with the host application-specific integrated circuit (ASIC), whether CPU, GPU, TPU or other, in the same chip package, Micron delivers the best bandwidth possible. Plus, our HBM2E device provides top energy efficiency and high capacity in a very small footprint (Figure 1). Since collaboration among the engineering ecosystem is critical to integrate and operate HBM2E efficiently, Micron shares our insights and best practices in this tech brief.

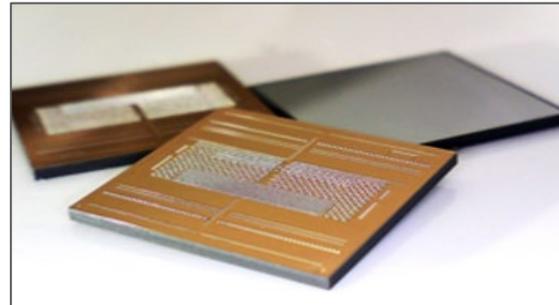


Figure 1: Micron HBM2E DRAM

HBM2E, an offering in Micron’s Ultra-Bandwidth Solutions, is well suited to meet the extremely high-memory bandwidth needs of high-performance computing (HPC), artificial intelligence (AI) and other compute-intensive applications. HBM2E is often an alternative to the well-established GDDR6 and GDDR6X SGRAM, which also provide very high memory bandwidth. Table 1 compares some of the key features of both DRAM technologies.

Table 1: HBM2E DRAM vs. GDDR6(X) SGRAM

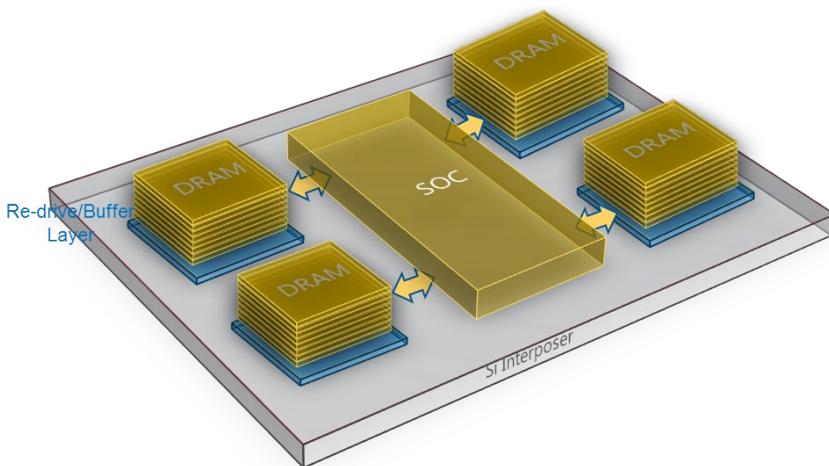
Parameter	HBM2E DRAM	GDDR6(X) SGRAM
Memory capacity per device	8GB, 16GB	1GB, 2GB
Channel count per device	8	2
Memory capacity per channel	1GB, 2GB	0.5GB, 1GB
Devices per system	4 to 6	up to 12 (up to 24 in clamshell configuration)
Total memory capacity	up to 96GB	up to 48GB
Maximum per-pin data rate	3.2 Gb/s	16 Gb/s (GDDR6) 24 Gb/s (GDDR6X)
I/O width	1024	32
Memory bandwidth per device	410 GB/s	64 GB/s
Total memory bandwidth	up to 2,450 GB/s	up to 768 GB/s (GDDR6) up to 1,152 GB/s (GDDR6X)
Package	KGSD (known good stacked die)	BGA-180

HBM2E DRAM Stacked Architecture

HBM2E DRAM is a 3D integration of four or eight DRAM memory layers in a single stacked device. (Twelve DRAM layers are planned for the future.) An additional logic or base layer at the bottom of the stack interfaces to the host ASIC and adds other functions to the HBM2E DRAM.

HBM2E is shipped as a known good stacked die (KGSD). Customers must enlist their outsourced semiconductor assembly partners (OSATs) to integrate the HBM2E with their host ASIC (system on a chip, or SoC) as a system-in-package (SiP), as illustrated in Figure 2. A typical system comprises four to six HBM2E DRAMs.

The main goal of such tight integration is greater power efficiency of the whole memory subsystem (SoC and DRAM) than a traditional memory DRAM and printed circuit board (PCB) assembly. The signal routing constraints



of state-of-the-art PCBs limit the number of data lines per SoC to about 384 lines or pins (DQs). With the number of signal lines being so constrained, increasing the per-pin data rate as shown in Table 1 for GDDR6 and GDDR6X is the only way to achieve a higher bandwidth. The drawback of such very high data rates, however, is that significant power must be dedicated to the data transmission between SoC and DRAM. It requires techniques like on-die signal termination (ODT), high-speed input buffers and clocking, delay-locked loop or phase-locked loop (DLL or PLL), clock-data recovery (CDR), and constant interface training and tracking.

Figure 1: Example of a System-in-Package (SiP) With HBM2E DRAMs

The benefit of this very high system integration is that all memory traffic now fully resides within the SiP. Using advanced packaging technologies like silicon interposers enables the data bus to run much wider than is possible with discrete DRAMs (e.g., 1,024 I/O lines per HBM2E device) and at a much lower per-pin data rate (e.g., 3.2 Gb/s/pin). At such comparably low data rates, none of the sophisticated and power-hungry techniques described above are required.

Beside the much better power efficiency, HBM2E also delivers higher capacity per SiP in a very small footprint. Figure 3 shows a cross section of a typical SiP. The HBM2E DRAM consists of four or eight memory core dies. The four stack or eight stack of memory dies are the same, just the height of the top die is different, so both configurations have the same cube height. In fact, the cube height matches the height of the host ASIC and enables the use of a planar cooling device for the SiP.

The memory dies within the HBM2E communicate vertically with the base die by through-silicon-vias (TSVs). A total of about 5,000 TSVs per HBM2E DRAM layer connect signals as well as power and ground supplies.

The HBM2E DRAM is soldered to a silicon interposer that routes all interface signals in very tiny traces to the host ASIC. The power and ground lines are routed via the silicon interposer and package substrate to the SiP's package balls (Figure 3).

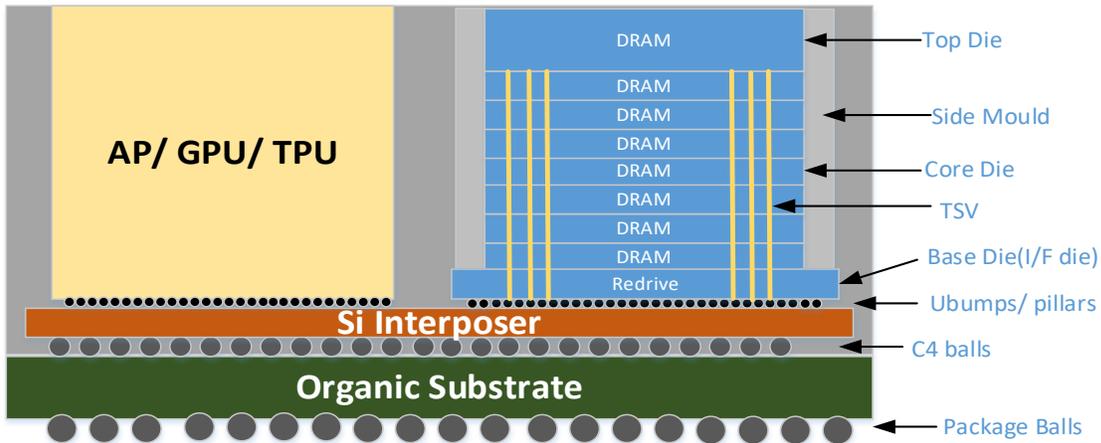


Figure 3: Cross Section of a System-in-Package (SiP) With HBM2E

HBM2E Independent Channel Architecture

Channel Architecture

The HBM2E DRAM is internally organized as eight independent channels A to H (Figure 4) for both four-high and eight-high DRAM configurations. Each channel is equipped with its own clock, command/address and data interface, and can be operated fully independent of the other channels. The 204 signals per channel are summarized in Table 2. Note that the channel count and thus the memory bandwidth are the same for four-high and eight-high configurations; the eight-high configuration just doubles the memory capacity of each channel.

The channel concept encourages efficient use of the total memory bandwidth. Accesses to the DRAM in many cases are only 32- or 64-bytes wide; these accesses can be served much more efficiently when the internal memory organization is segmented, allowing several (e.g., eight) of such fine-grain and independent memory accesses to be executed in parallel.

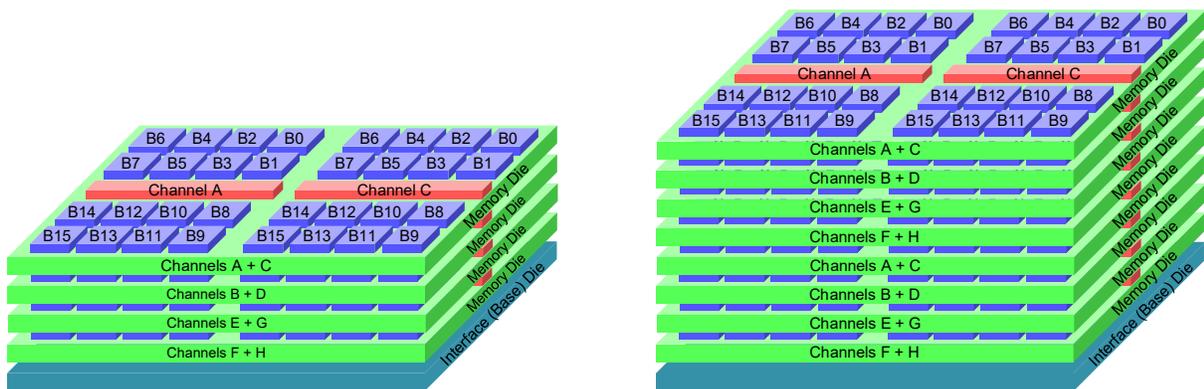


Figure 4: Example of HBM2E Channel Organization

Table 2: Per-Channel Signal List

Symbol	Type	Description
CK_t, CK_c	Input	Command clock
CKE	Input	Clock enable
C[8:0]	Input	Column command and address
R[6:0]	Input	Row command and address
DQ[127:0]	I/O	Data input/output: 128-bit data bus
DBI[15:0]	I/O	Data bus inversion
DM[15:0]	I/O	Data mask or error-correction code (ECC) data
PAR[3:0]	I/O	Data parity
DERR[3:0]	Output	Data parity error
AERR	Output	Address parity error
WDQS[3:0]_t, WDQS[3:0]_c	Input	Write data strobe
RDQS[3:0]_t, RDQS[3:0]_c	Output	Read data strobe

Pseudo Channel Mode

Each of the eight channels of the HBM2E DRAM is further subdivided into two so-called pseudo channels (PCs) that operate semi-independently. Both PCs share a common command clock (CK) and row (R) and column (C) command/address inputs (Figure 5). Effectively, this results in virtually 16 channels per device. The command encoding includes a bit that determines which of the two PCs (PC0 or PC1) executes the command. Each WRITE or READ operation lasts for two clock cycles and transfers 256 bits of data as a burst of four (BL4) on 64 DQs (DQ[63:0] for PC0 and DQ[127:64] for PC1). The four-unit intervals of the BL4 burst are denoted as D0 to D3. The corresponding READ and WRITE commands are one clock-cycle long. Issuing READ or WRITE commands alternating between PC0 and PC1 on the column commands bus, as illustrated in Figure 6, delivers 100% data bus utilization on both PCs. Row activation and precharge operations can be initiated in parallel on the row command bus without impacting reads or writes.

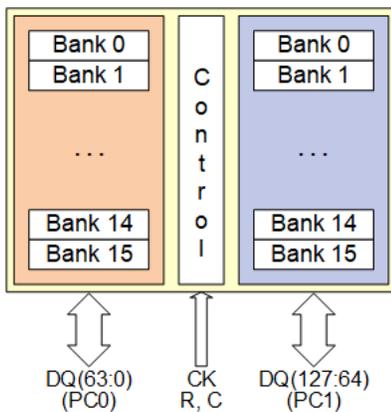


Figure 5: Pseudo Channel Mode Architecture (Single Channel)

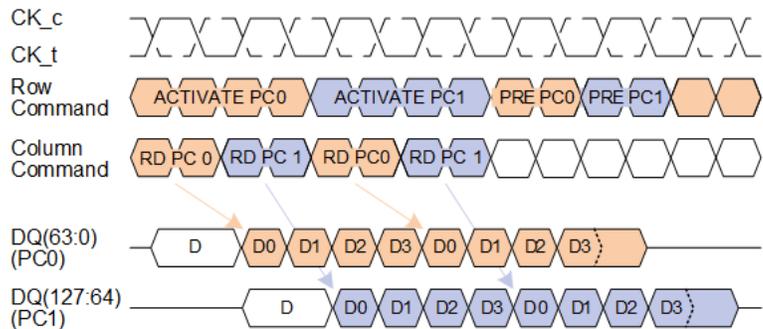


Figure 6: Pseudo Channel Mode Operation (Single Channel)

Addressing

The per-channel addressing parameters are depicted in Table 3. As illustrated in Figure 4, the eight-high configuration doubles the capacity of a four-high configuration by adding a second memory die to each channel. This is reflected in the addressing by the additional “stack ID” (SID) bit, which acts as fifth bank address bit and selects between the lower and upper memory die and thus between the lower and upper 16 banks.

Table 3: Per-Channel Addressing

Parameter	8GB (4-High)	16GB (8-High)
Density per channel	8Gb	16Gb
Density per pseudo channel (PC)	4Gb	8Gb
Prefetch size per PC (bits)	256	
Bank count	16	2 x 16
Bank address	BA[3:0]	SID, BA[3:0]
Row address	RA[14:0]	
Column address	CA[5:1]	
Page size per PC	1KB	

Note: Density, prefetch size and page size do not include the ECC bits.

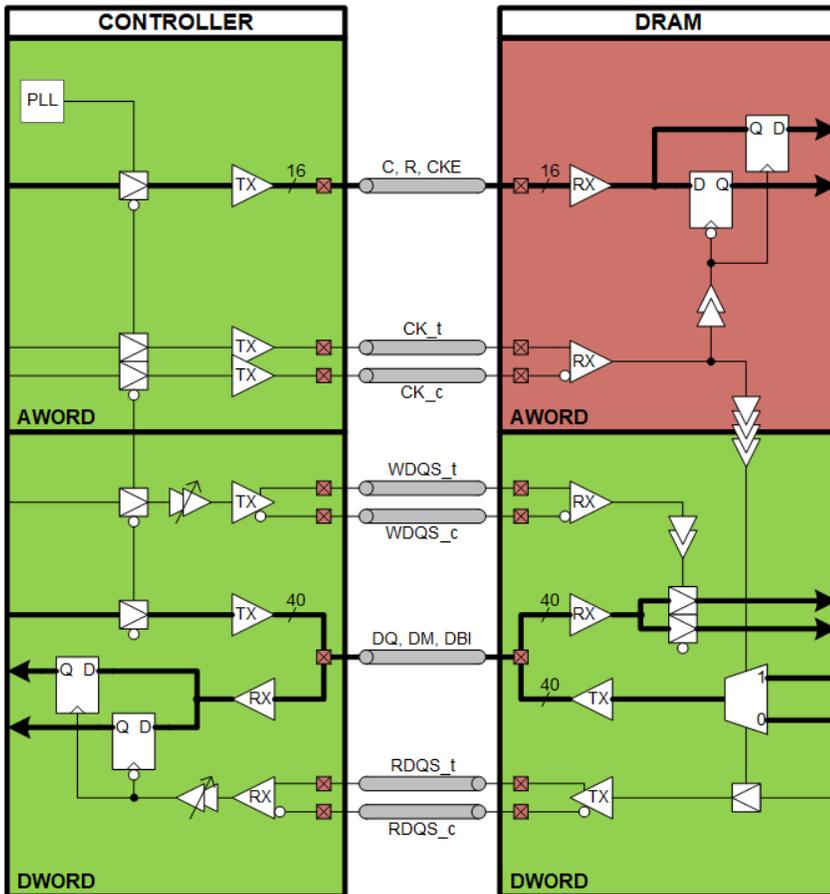


Figure 7: HBM2E DRAM Clocking Scheme

HBM2E DRAM Operation

Clocking

Figure 7 shows the example clocking system of an HBM2E-DRAM-based system, and Table 4 adds typical frequencies and data rates:

- A differential CK clock (CK_t, CK_c) latches the command and address inputs (R, C, CKE).
- Write data are received center aligned with the differential write data strobe (WDQS_t, WDQS_c); one WDQS strobe pair is associated with four data bytes.
- Read data are transmitted edge aligned with the differential read data strobe (RDQS_t, RDQS_c); one RDQS strobe pair is associated with four data bytes.

This clocking scheme is very similar to DDR DRAMs, except that unidirectional WDQS and RDQS strobes are provided compared to bidirectional data strobes

(DQS) in DDR DRAMs. The unidirectional strobes provide a faster bus turnaround between reads and writes.

To reduce operating power, no PLL or DLL, like in other DRAMs, is associated with the memory interface.

The HBM2E DRAM supports a wide contiguous operating frequency range, from 50MHz up to the maximum rated clock frequency, which allows the host a wide range in which to adapt the clock frequency and thus the power of a channel to the current workload.

Table 1: Example Data Rates

Signal	Value	Type
Differential CK clock	1.6GHz	
CKE	1.6 Gb/s	SDR (CK)
Command/address inputs (R, C)	3.2 Gb/s	DDR (CK)
Differential write and read data strobes (WDQS; RDQS)	1.6 GHz	
Data bus (DQ; DM, DBI)	3.2 Gb/s	DDR (WDQS or RDQS)

Note: SDR = single data rate: signals are referenced to the rising clock edge.
 DDR = double data rate: signals are referenced to both edges of the clock or data strobe.

Signaling

The HBM2E interfaces use regular pull/push drivers with programmable driver strength for signaling. The significantly large signal attenuation of the channel does not require termination at the receive end. This termination abandonment is one reason for the superior energy efficiency of HBM2E DRAMs.

Write Data Mask

The data mask (DM) pins of the HBM2E DRAM's data interface can be configured to either serve as write DM with byte-wide granularity or as additional data bits for system ECC (see below). When configured for DM support, write data received on the DQ inputs are ignored (masked) in case DM is sampled high; write data are written to the memory array in case DM is sampled low. Write data mask is common to most DRAM devices.

Error-Correcting Code

When error-correcting code (ECC) mode is enabled, the DM pins are treated as additional data I/Os, effectively adding a ninth bit to each data byte. The total data bus width then is 144 (128 + 16) per channel or 1,152 (1,024 + 128) in total. The array provides additional memory cells for this ninth data bit, and the effective total memory capacity is 9GB (for four-high DRAM) and 18GB (for eight-high DRAM).

These ECC bits are intended for system parity bits. Internally, they are treated as regular data bits, which offers a very efficient way to transfer these system parity bits along with the corresponding data in a single WRITE or READ command. It is important to highlight that the HBM2E DRAM does not perform any error checking on these bits and that these ECC bits should not be mixed up with on-die ECC functions that other DRAM may support.

Data Bus Inversion

The HBM2E DRAM supports a byte granular data bus inversion coding (DBIac) during write and read operations (Figure 8). The function can be enabled or disabled independently for writes and reads. The purpose of the DBI function is to reduce the data toggle rate on the external data lines between the host ASIC and the HBM2E, thus reducing the overall power as well as the power-supply-induced jitter (PSIJ).

The DBIac function is symmetric for the host ASIC and DRAM and works as follows:

- Transmitter (e.g., HBM2E DRAM on reads) counts the number of DQs transitioning from the previous state. The transmitter inverts read data and sets DBI high when the number of transitioning data bits within a byte is greater than four or when the number of transitioning data bits within a byte equals four and DBI was high. Otherwise, the transmitter does not invert the read data and sets DBI low.
- Receiver (e.g., HBM2E DRAM on writes) inverts data received on DQ in case DBI is sampled high or leaves the data uninverted in case DBI is sampled low.

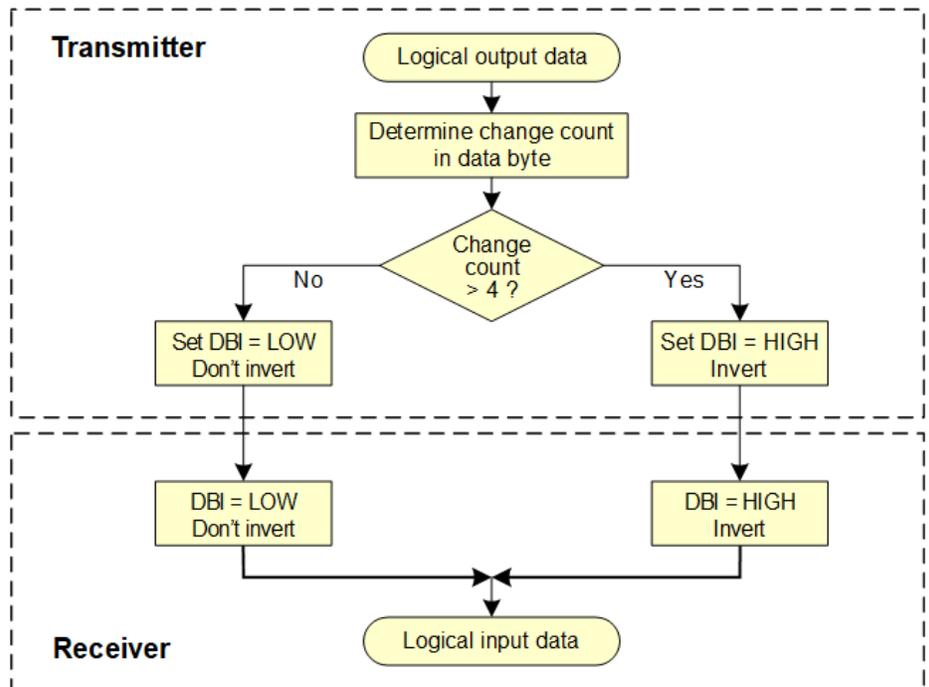


Figure 8: Data Bus Inversion Coding Algorithm

Basic Write and Read Operation

A basic write and read operation is illustrated in Figure 9:

1. The ACTIVATE command opens a row in a bank for a subsequent write or read access, and the bank and row addresses are provided with the ACTIVATE command. Note: ACTIVATE is a two-cycle command, and the related row-to-column delay for writes (t_{RCDWR}) is referenced to the second cycle of the command when all required address bits are received.
2. Once the t_{RCDWR} timing has elapsed, a WRITE command is issued, and the bank and column addresses are provided with the command.
3. The write data are received as a burst of 4 (D0 to D3) after the write latency (WL). Write data are latched on both edges of the write data strobe (WDQS), which must be provided center-aligned with the write data. An additional preamble WDQS pulse prior to the write data ensures that the WDQS signal waveform has settled for the actual latching of the write data.
4. After the write-to-read timing (t_{WTR}), a READ command is issued; the bank and column addresses are provided with the command.
5. The read data are transmitted as a burst of 4 (D0 to D3) after the read latency (RL); read data strobe pulses (RDQS) are transmitted edge-aligned with the read data; and an additional preamble RDQS pulse prior to the read data ensures that the RDQS signal waveform has settled for the actual latching of the read data by the host ASIC.
6. The PRECHARGE command closes the open row, and the bank address is provided with the command. The command is issued once both the activate-to-precharge delay (t_{RAS}) and the read-to-precharge (t_{RTP}) timing have been met. The bank is available for a next row activation after the row precharge time (t_{RP}).

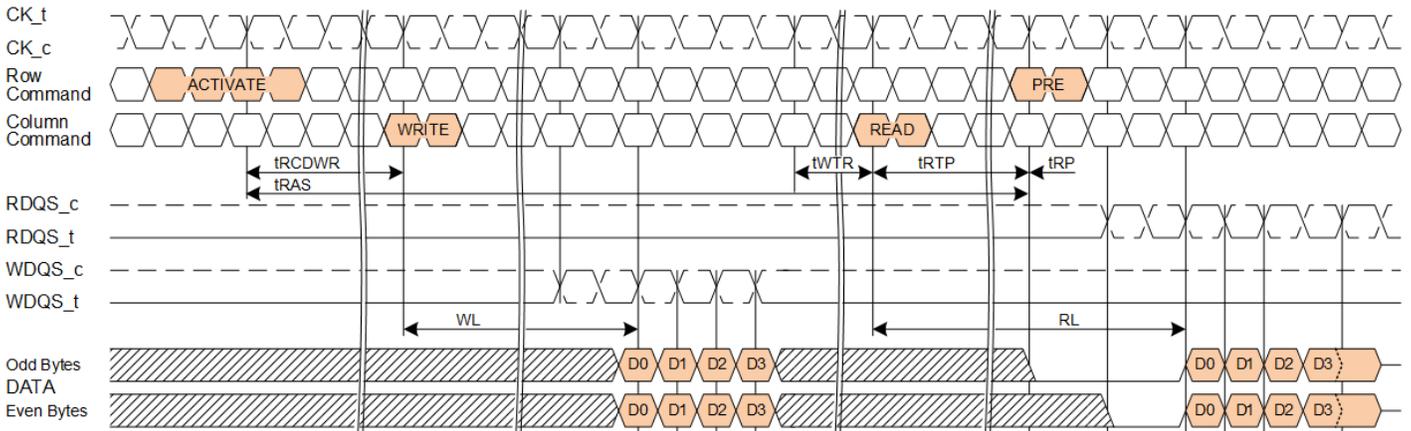


Figure 9: Basic Write and Read Operation

Low Power Modes

HBM2E offers the same low power modes as many other DRAMs to reduce the power consumption when the application temporarily does not need to access the DRAM:

- Power-down mode disables the command decoder and internal clocking. This mode is well suited for short periods of no activity. The DRAM can resume full operation instantaneously, and the memory array is not refreshed during power-down.
- Self-refresh is the lowest power state where the DRAM internally manages the refresh of the memory cells without intervention from the memory controller.

Package Dimensions and Micro Bump Matrix

Figure 10 shows the mechanical outline of the HBM2E DRAM. The cube size is 11mm x 10mm (max).

The bump matrix is organized in 300 rows and 68 columns located in the center of the device, with a total of 6,303 micro bumps. The staggered (diagonal) arrangement of the bumps, with an approximately 55µm diagonal pitch, allows for a larger bump diameter than possible with an orthogonal arrangement and the same bump pitch.

All I/O bumps are kept in the densely populated PHY region on the right-hand side of the bump matrix. The host-side PHY region is expected to be a copy of the DRAM-side PHY matrix, allowing straight parallel routing of all signal traces on the interposer with equal trace length. The location of the bump matrix can be seen in Figure 1.

A more detailed view of this PHY region is shown in Figure 11. Power and ground supply bumps are placed between the I/O signals, but also spread all over the entire bump matrix for a solid power delivery to the device. The picture also reveals gaps in the bump matrix outside the I/O region. These gaps may be used by the DRAM vendor to place probe pads for testing the device.

The signal bumps of the PHY region are grouped in AWORDS and DWORDS:

- All command and address inputs of a channel, along with the CK clock, are located within the AWORD.
- All data I/Os, including WDQS and RDQS, are located in the DWORDS. One channel comprises four DWORDS, and each DWORD comprises 32 DQs plus one WDQS and RDQS pair each.

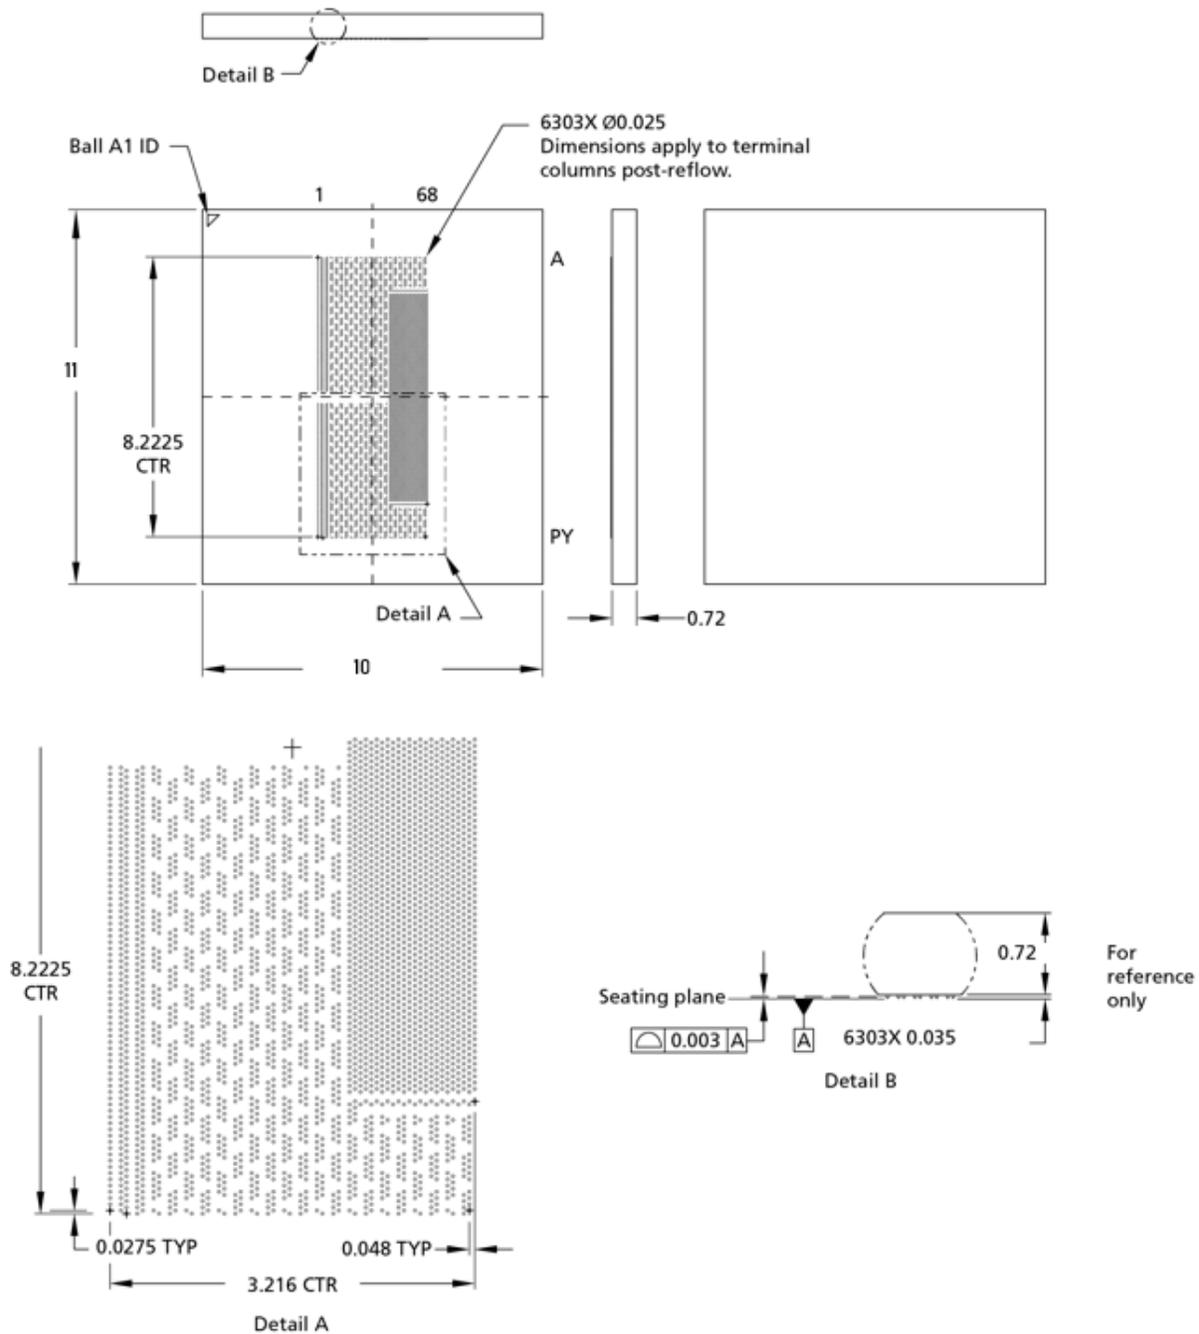


Figure 10: Micron HBM2E Mechanical Outline

- The signals in the so-called MIDSTACK region are common to all channels. These include the chip reset, signals for temperature reporting and the IEEE 1500 test port.

The color codes in Figure 11 illustrate the geographical arrangement of the eight channels within the PHY region:

- Channels A, B, E and F are located above the central so-called MIDSTACK region, while channels C, D, G and H are below it.
- The four channels each are placed in a 2 x 2 matrix with the AWORDS in the center, two DWORDS above and the other 2 DWORDS below the AWORD.

- This symmetrical and interleaved placement corresponds well with the vertical distribution of channels in the HBM2E stack, as shown in Figure 4.

DWORD0 Channel E	DWORD0 Channel A
DWORD0 Channel F	DWORD0 Channel B
DWORD1 Channel E	DWORD1 Channel A
DWORD1 Channel F	DWORD1 Channel B
AWORD Channel E	AWORD Channel A
AWORD Channel F	AWORD Channel B
DWORD2 Channel E	DWORD2 Channel A
DWORD2 Channel F	DWORD2 Channel B
DWORD3 Channel E	DWORD3 Channel A
DWORD3 Channel F	DWORD3 Channel B
MIDSTACK (Reset, IEEE 1500 Port, Temperature)	
DWORD0 Channel G	DWORD0 Channel C
DWORD1 Channel H	DWORD0 Channel D
DWORD0 Channel G	DWORD1 Channel C
DWORD1 Channel H	DWORD1 Channel D
AWORD Channel G	AWORD Channel C
AWORD Channel H	AWORD Channel D
DWORD2 Channel G	DWORD2 Channel C
DWORD2 Channel H	DWORD2 Channel D
DWORD3 Channel G	DWORD3 Channel C
DWORD3 Channel H	DWORD3 Channel D

Figure 11: HBM2E PHY Region

Test and Debug Features

The HBM2E DRAM includes a large and powerful set of features to test, repair and debug. These features help to verify the HBM2E DRAM's internal operation and connectivity to the host ASIC after the HBM2E has been assembled in the SiP. They also provide a way to repair the device if a failure is detected after assembly.

Deploying these test and repair capabilities is critical for a high assembly yield. To emphasize the importance, let's compare HBM2E assembly in a SiP with the traditional assembly of discrete DRAM components (or any other component) on a PCB or a memory module (DIMM). In the case of a solder joint defect found post reflow, the defective component is replaced in a process called rework, which consists of removing the component, cleaning the PCB's surface, and attaching a replacement component with a solder step.

Such rework is not possible on a silicon interposer as used in the HBM2E. Without on-die repair capabilities, a single solder joint fail on either the host ASIC or DRAM side would not only require the defective HBM2E DRAM to be thrown away, but the entire SiP with host ASIC and several HBM2E DRAMs would have to be scrapped.

IEEE 1500 Test Port

An IEEE 1500-compliant test port controls all test and repair functions in the HBM2E DRAM. The IEEE 1500 test port concept is derived from the IEEE 1149.1 boundary scan standard and targets embedded macros in a larger system. The central control (test access port, or TAP) resides within the host ASIC. The HBM2E DRAM receives pre-decoded control signals for serially shifting data in and out (SHIFTWR), capturing internal states (CAPTUREWR) and updating internal states with new data (UPDATEWR) as

illustrated in Figure 12. The related signal bumps are in the MIDSTACK region of the bump matrix and common to all channels.

The test can also provide the following functions:

- An instruction register (WIR) can serially be loaded with the specific test instruction. The instruction codes include bits that determine whether an instruction is executed by all channels or single channels only. The AND gates shown in the lower half of the figure turn on/off the internal capture, shift and update control signals for channels that are not selected.
- A bypass register (WBY) consists of a single-step shift register.
- A boundary scan register (WBR) does boundary scan testing (explained below).
- Several data registers (WDR) enable internal test and repair operations (described below).

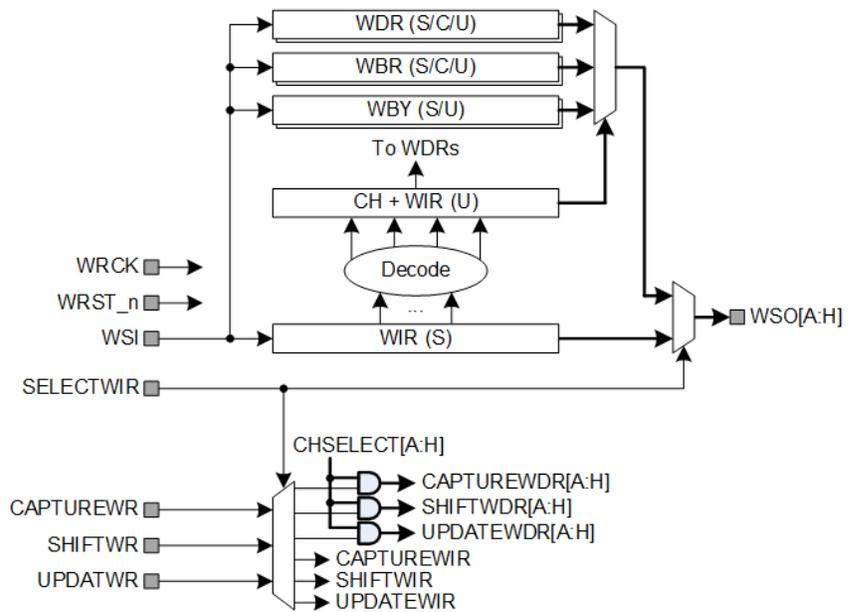


Figure 12: IEEE 1500 Test Port

As a special implementation of the IEEE 1500 standard, the HBM2E DRAM’s test port provides a dedicated serial out (WSO) for each channel, allowing the host ASIC to capture test results from each channel in parallel and thus reduce the overall test time.

Boundary Scan

Boundary scan is a test feature in widespread use across the entire industry for decades. It verifies the connectivity between components through a standardized test interface. The details of boundary scan are specified in the IEEE standard 1149.1.

The boundary scan instruction of the HBM2E DRAM allows for testing the interconnect between host ASIC and DRAM without operating the DRAM core itself. Logic 0 and 1 data are serially shifted in on the transmitting side (host ASIC or DRAM), and the received data are latched and then serially shifted out on the receiving side.

Loopback Test Modes

While boundary scan testing is adequate to find static (hard) fails, other interface-related fails might be detected only when the interface is operated at target speed. Loopback test modes are supported by the HBM2E DRAM for such tests; they exercise the I/O circuitry but not the memory array.

The loopback test modes use the primary input and output data latches of the mission mode interface but allow the latches to be configured as linear feedback shift register (LFSR) or multiple input shift register (MSIR), as shown in Figure 13.

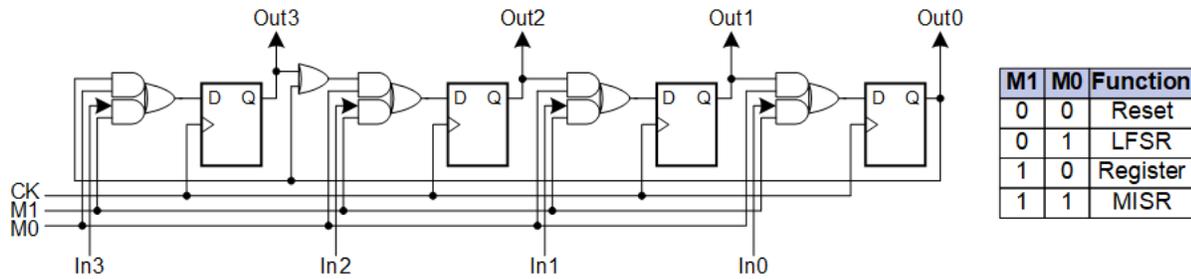


Figure 13: Linear Feedback Shift Register (LFSR) and Multiple Input Shift Register (MISR)

The loopback test modes work as follows:

- LFSR mode is intended for testing the read direction. The LFSR is first initialized with a starting value (seed) and then generates a pseudo-random pattern sequence that is transmitted with READ commands, following the same protocol and timings as for normal reads, as shown in Figure 9. The polynomial used by the LFSR is standardized, so the host ASIC can compute the expected read data and compare against the received data on each cycle.
- MISR mode is intended for testing the write direction. The MISR is again initialized with a starting value (seed). The host ASIC then issues one or more WRITE commands and sends random write data following the same protocol and timings as for normal writes (again, see Figure 9). The MISR compresses the write data into a signature. This signature can then be read back by a subsequent READ command or via the IEEE 1500 test port and compared against a precomputed expected signature.
- MISR mode helps detect whether a fail has occurred on any of the data bits over multiple write cycles, but it does not enable identification of a single faulty DQ lane. For that purpose a third test function, called LFSR compare mode, is provided; in this mode, the LFSR is used to compute expected write data. With each WRITE command, the host generates write data by an LFSR using the same polynomial. The HBM2E DRAM compares the received data against the expected write data given by the LFSR. In case of a mismatch, an internal bit (“sticky error flag”) associated with each DQ is set. At the end of the test, all error bits can be read via the IEEE 1500 test port to be checked.
- MISR and LFSR compare modes are also supported for the command/address interface.

Interconnect Redundancy Remapping (Lane Repair)

If an interconnect-related failure has been detected by use of boundary scan or loopback test mode, the host may initiate a repair of the broken lane. The HBM2E DRAM supports interconnect lane remapping to recover functionality from a broken or somehow nonfunctional signal connection between the host ASIC and the DRAM.

The bump matrix includes a limited number of redundant or spare I/Os that are normally off and only activated when a broken connection needs to be repaired. One redundant input is provided for the row (R) and column (C) command/address inputs each, and one redundant I/O is provided for each data byte (or per two data bytes if the DBI functionality must be preserved).

An example of a lane repair is illustrated in Figure 14. Note that any lane repair must be programmed identically on both the host ASIC and DRAM sides. In this example, the DQ0 lane is intact and logical DQ0 data are transmitted and received as normal via the DQ0 lane. The DQ1 lane, however, is broken; the DQ1 logical data are remapped to the DQ2 output latch and driver on the transmitting end, transmitted and received via the DQ2 lane, and mapped back to the logical DQ1 at the receiving side after the input latch. The same remapping to its neighbor I/O and lane occurs with other DQs, and the last DQ is remapped to the spare I/O and lane. With this scheme, the remapping occurs at the interconnect level only and does not affect the logic operation of the device.

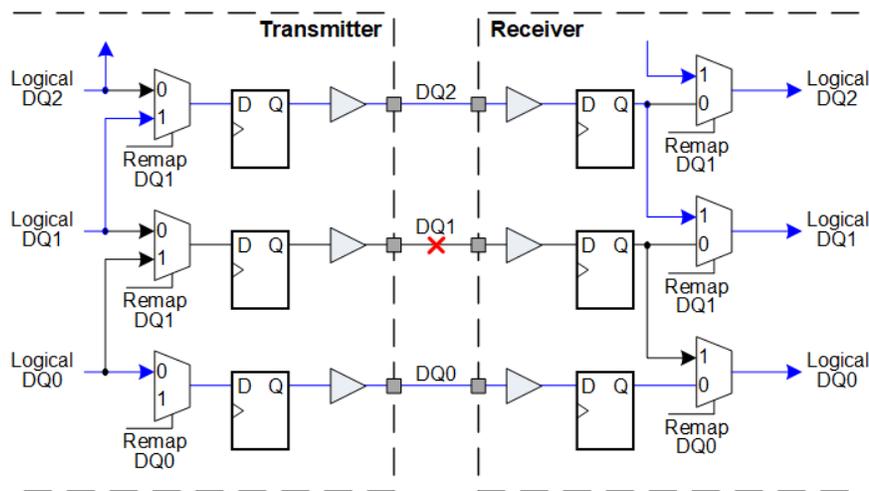


Figure 14: Example Lane Repair

The lane repair is programmed via the IEEE 1500 test port. There is an option provided to perform a temporary (volatile) remapping to validate the repair (soft lane repair) and another option to blow a fuse and make this remapping permanent (hard lane repair).

Memory Array Test and Repair

Embedded in the HBM2E DRAM is a very sophisticated and highly programmable engine for testing the memory array (MBIST). The test coverage of this MBIST is like the tests done at Micron before the final HBM2E DRAM is shipped to the customer.

This MBIST enables the customer and its ecosystem to thoroughly test the memory array after assembly before the final product is shipped to the end customer. MBIST is controlled via the IEEE 1500 test port and provides the test results upon completion.

If MBIST detects a failure, the host may initiate a repair of the failing cell. For that purpose, the memory array is equipped with spare rows and can be activated via the IEEE 1500 test port. Like with lane repairs, there is an option to perform a temporary (volatile) repair to validate the repair (soft repair) and another option to blow a fuse and make this repair permanent (hard repair).

Other IEEE 1500 Test Port Functions

The IEEE 1500 test port offers a few more functions that are useful in operating the HBM2E DRAM:

- The device ID function delivers several types of hard-coded data like density, manufacturer code, part number and a unique serial number to the host. The host can read these data at power-up and configure the HBM2E DRAM based on this information or use these data for tracking purposes.
- The temperature function reports the DRAM's junction temperature and allows the host to determine whether the device is safely operating within the allowed temperature range. This temperature read occurs on the IEEE 1500 test port so can be performed frequently without impacting the memory bandwidth.
- The mode register dump/set function configures the HBM2E DRAM's programmable features, but more importantly, it reads back the programmed value (e.g., for debug purposes).

Summary

This tech brief describes the HBM2E DRAM's 3D architecture and its integration with the host ASIC in a SiP. We offered best practices on how to operate an HBM2E DRAM, including some programmable options. For a high-yielding integration of the HBM2E DRAM in the SiP, this brief provides a description of essential embedded test and debug features.

HBM2E DRAM is the highest bandwidth DRAM device in Micron's Ultra-Bandwidth Solutions portfolio. HBM2E combines ultra-high memory bandwidth and high energy efficiency with an extremely small footprint enabled by the vertical or 3D integration of four or eight DRAM memory layers in a single stacked device. An even higher density integration with 12 DRAM layers in the HBM2E memory is planned for future devices.

Learn More

Investigate Micron's full portfolio of high-bandwidth solutions at micron.com/hbs. Check out HBM2E memory at micron.com/hbm2e. For the latest developments on Micron's memory innovation, follow us on [@MicronTech](#) on [Twitter](#) or [LinkedIn](#).

micron.com

© 2021 Micron Technology, Inc. All rights reserved. All information herein is provided on an "AS IS" basis without warranties of any kind, including any implied warranties, warranties of merchantability or warranties of fitness for a particular purpose. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Rev. A 04/2021 CCM004-676576390-11532