

Technical Note

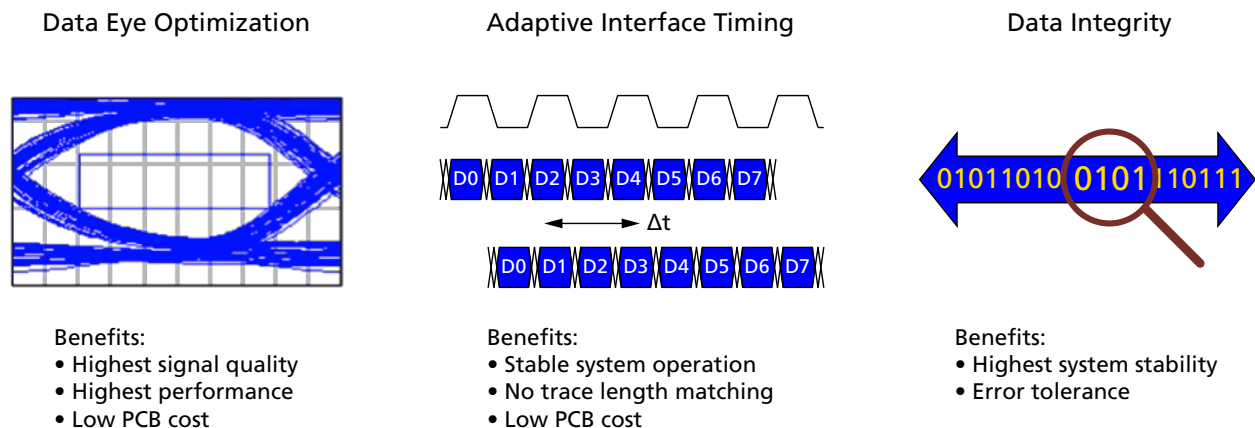
GDDR5 SGRAM Introduction

Introduction

This technical note describes the features and benefits of GDDR5 SGRAM. GDDR5 is the ideal DRAM device for graphics cards, game consoles, and high-performance computing. The device offers unprecedented memory bandwidth and low system implementation costs with the following key features:

- Data eye optimization by adapting I/O impedance and reference voltage to the system characteristics
- Efficient adaptation and tracking of interface timings
- Improved data integrity with hardware support for detecting and correcting transmission errors

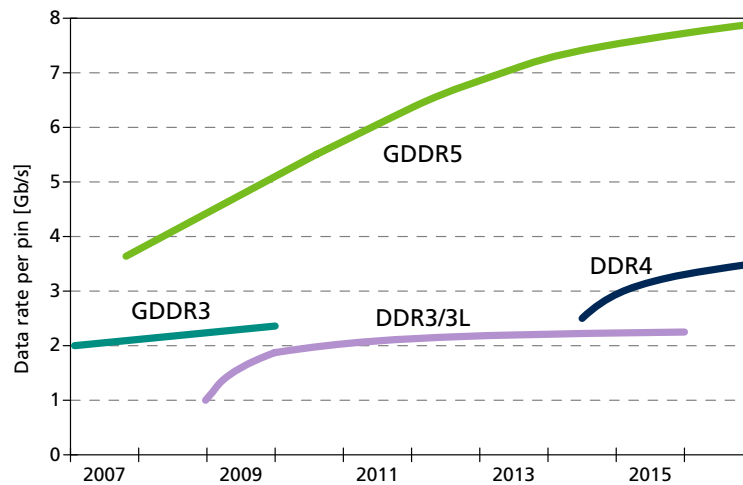
Figure 1: GDDR5 Key Features



The device has ultra-high bandwidth compared to other popular DRAM standards (see the figure below). When the device was introduced, GDDR5-based systems operated at 3.6 Gb/s. Since then, data rates have increased to over 6 Gb/s in mainstream applications and 7 Gb/s in high-end systems. For example, a single GDDR5 can read or write the data equivalent of five DVDs (4.7GB each) in less than a second when operating at 6 Gb/s per pin, or 24 GB/s per device.

Micron is working closely with enablers to increase the data rate, with the goal of reaching 8 Gb/s in the near future.

Figure 2: Data Rate Comparison



Interface and Clocking

GDDR5 Interface

GDDR5 combines reliable single-ended signaling with improvements to the clocking system that overcome the speed limitations in previous generations of graphics memory devices. These improvements enable the industry to constantly increase the data rates of GDDR5-based systems with each new product generation.

The device uses high-level termination for command, address, and data. This results in significant power savings compared to mid-level terminated systems. It operates from a 1.5V power supply for high-speed applications and a 1.35V power supply when maximum performance is not required.

The device interface is designed for systems with a 32-bit wide I/O memory channel, resulting in 32 bytes of data transferred per memory cycle. Systems can span from 64-bit wide I/O (two memory channels) for entry-level systems, to 512-bit wide I/O (16 memory channels) for high-end systems.

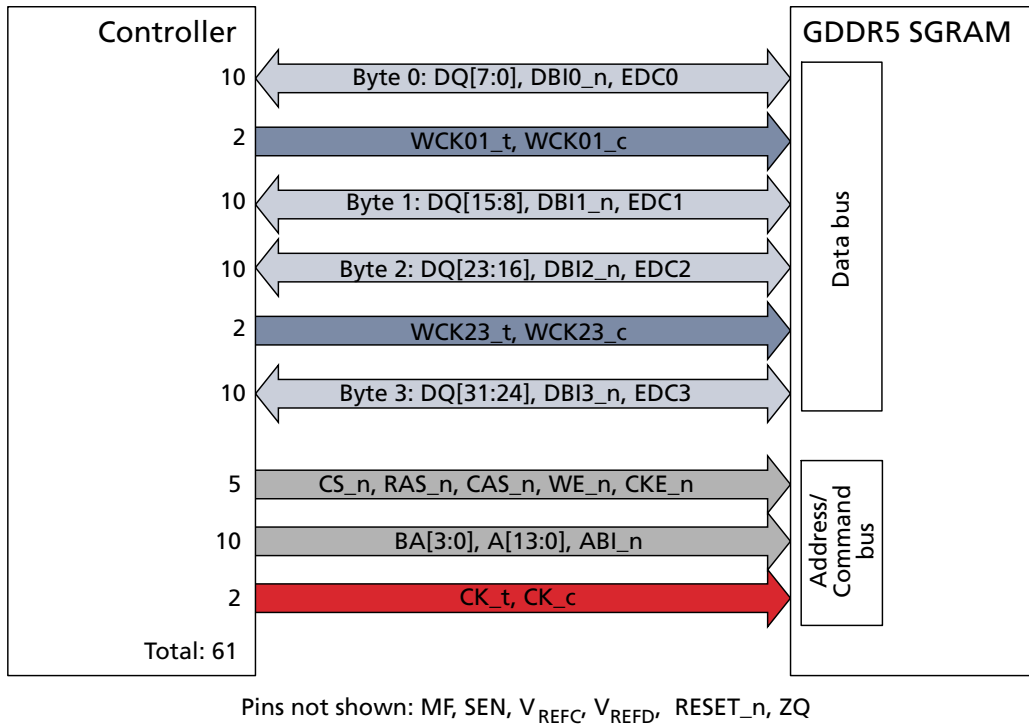
A single memory channel is comprised of 61 interface signals (see the figure below):

- One differential clock pair for command and addresses: CK_t and CK_c
- Five command inputs: RAS_n, CAS_n, WE_n, CS_n, and CKE_n
- Ten multiplexed address inputs: BA[3:0], A[13:0], and ABI_n
- One 32-bit wide data bus: Each byte is accompanied by one data bus inversion (DBI_n) and one error detection and correction (EDC)
- Two differential forwarded data clock pairs for bytes 0 and 1 (WCK01_t, WCK01_c) and bytes 2 and 3 (WCK23_t, WCK23_c)

The following pins are either pulled HIGH, pulled LOW, or are connected to other sources:

- Mirror function (MF)
- Scan enable (SEN)
- Input reference for command and address (V_{REFC})
- Data input reference (V_{REFD})
- Chip reset (RESET_n)
- Impedance reference (ZQ)

Figure 3: GDDR5 Interface



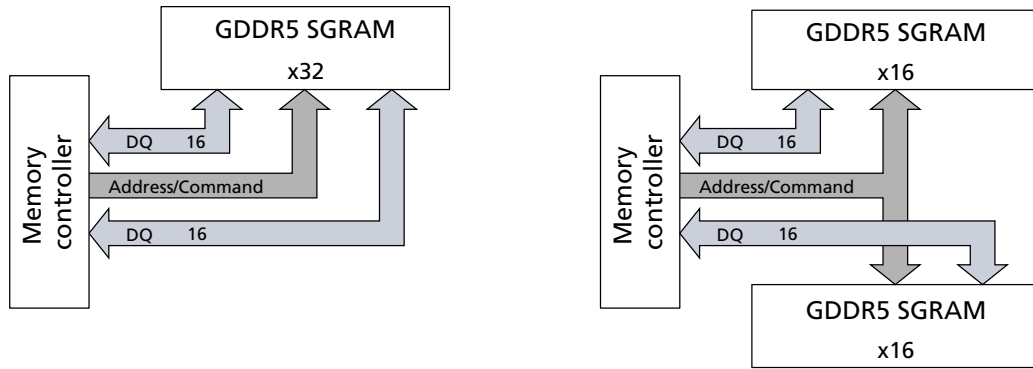
Normal (x32) and Clamshell (x16) Modes

Adding additional DIMMs to memory channels is the traditional way of increasing memory density in PC and server applications. However, these dual-rank configurations can lead to performance degradation resulting from dual-load signal topology. GDDR5 uses a single-loaded or point-to-point (P2P) data bus for the best performance.

GDDR5 devices are always directly soldered down on the PCB and are not mounted on a DIMM.

Each device supports x32 mode and a x16 clamshell mode, and the mode is set at power-up. In x16 mode, the data bus is split into two 16-bit wide buses that are routed separately to each device. Address and command pins are shared between the two devices to preserve the total I/O pin count at the controller. However, this point-to-two-point (P22P) topology does not decrease system performance because of the lower data rates of the address or command bus.

Figure 4: Normal (x32) and Clamshell (x16) Modes



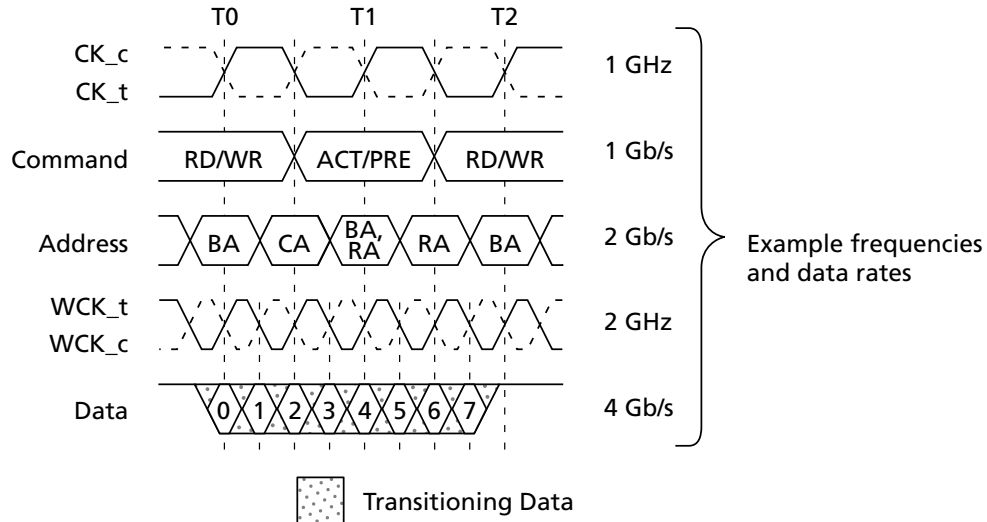
Clamshell mode essentially doubles the memory density on a x32 GDDR5 channel. The frame buffer size can be changed using the same component. For example, a 2Gb device can be used to build the following systems with a 256-bit wide memory bus:

- 2GB frame buffer using 8 devices configured to x32 mode
- 4GB frame buffer using 16 devices configured to x16 mode

Clocking and Data Rates

The figure below shows how the device runs off of two different clocks.

Figure 5: Clock Frequencies and Data Rates



Commands and addresses are referenced to the differential clock CK_t and CK_c. Commands are registered as SDR at every rising edge of CK_t. Addresses are registered as DDR at every rising edge of CK_t and CK_c.

Read and write data is referenced as DDR at every rising edge of a free-running differential forwarded clock (WCK_t and WCK_c). WCK_t and WCK_c replace the pulsed strobes (WDQS and RDQS) used in other devices, such as GDDR3, DDR3, or DDR4.

Clock frequencies and data rates are often confused with each other when referencing graphics card performance. Compared to the 2x-data rate and the CK clock relationship in DDR3, DDR4, and GDDR3, the 4x-relationship between the data rate and the CK clock is a key advantage for GDDR5. For example, a 1 GHz clock is equivalent to a 2 Gb/s data rate for a DDR3 or DDR4 compared to a 4 Gb/s data rate for GDDR5.

The lower command and address data rates were selected intentionally to allow a step-wise interface training at the target speed.

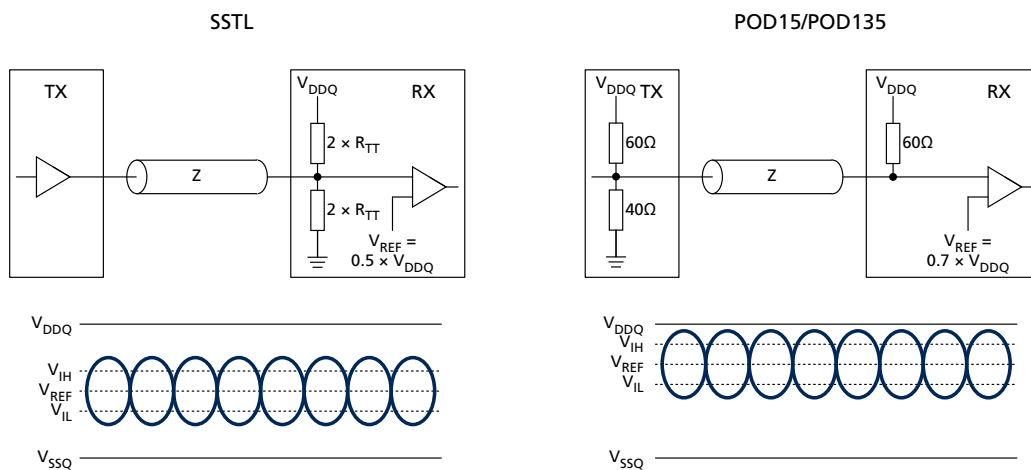
Considering the burst length of 8 and the CK and WCK frequency relationship (shown in Figure 5), each READ or WRITE burst takes two CK clock cycles. READ and WRITE commands are issued every second cycle for gapless READ or WRITE operations (see T0 and T2 in Figure 5). The intermediate command slot at T1 is used to open (ACTIVATE command) or close (PRECHARGE command) a page in the bank that is parallel with the seamless READ or WRITE operations.

High-Speed Signaling

Signaling Scheme and On-Die Termination

The figure below compares the pseudo open drain (POD) signaling scheme of GDDR5 with the stub series terminated logic (SSTL) scheme of DDR3. The POD driver uses a 40Ω (pull-down) or 60Ω (pull-up) impedance that drives into a 60Ω equivalent on-die terminator tied to V_{DDQ}. The benefit of the V_{DDQ} termination is that static power is only consumed when driving LOW. This helps reduce power consumption in the memory interface.

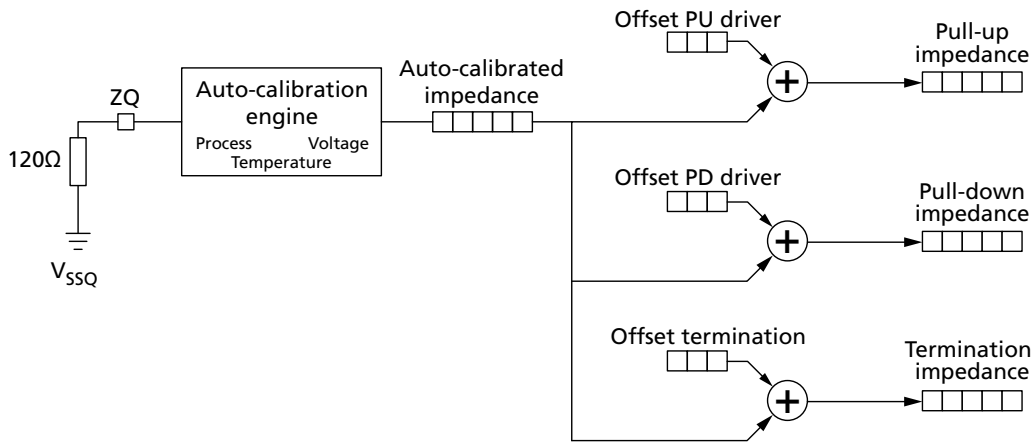
Figure 6: Signaling Schemes



Impedance Calibration and Offsets

Driver and terminator impedances are continually calibrated against an external precision resistor that is connected to the ZQ pin. This auto-calibration feature compensates for impedance variations that are a result of process, voltage, and temperature changes. A special memory controller command is not required (as it is for other DRAM devices) because the command is triggered internally and executed in the background. The calibrated driver and terminator impedance values can be adjusted (offset) to optimize the matching impedance in the system. This offset capability is provided separately for pull-down and pull-up driver strength, data termination, address/command termination, and WCK termination.

Figure 7: Impedance Offsets



V_{REFD} Options and Offsets

The data input reference voltage (V_{REF}) in Figure 6 may be supplied externally or generated internally. A more stable data eye typically results from using the internal V_{REFD} . V_{REFD} offset capability can vertically shift the write data eye when the eye opening is not symmetrical around the default V_{REFD} level. The optimum V_{REFD} offset is typically determined during system qualification, and then the value is programmed into the GDDR5 during power-up.

Data Bus Inversion and Address Bus Inversion

Data bus inversion (DBI_{dc}) reduces the DC power consumption and supply noise-induced jitter on data pins because the number of DQ lines driving a low level can be limited to four within a byte. DBI_{dc} is evaluated per byte.

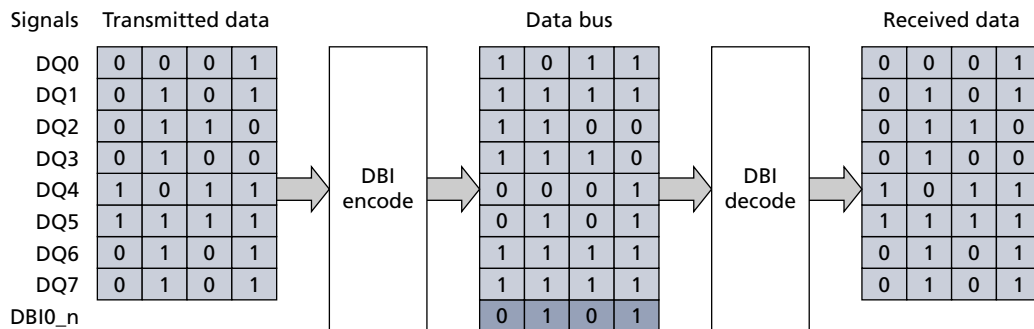
The DBI_n pins are bidirectional, active LOW, DDR signals. For WRITES, they are sampled by the device along with the DQ of the same byte. For READS, they are driven by the device along with the DQ of the same byte.

The transmitter (the controller for WRITES and the device for READS) decides whether to invert or not invert the data conveyed on the DQs. The receiver (the device for WRITES and the controller for READS) has to perform the reverse operation based on the DBI_n pin level.

The same function is also available via the address bus inversion (ABI) and is supported by the ABI_n signal.

The positive effect of DBI and ABI on the data eye width is generally accepted. Systems can achieve higher data rates by simply enabling DBI and ABI.

Figure 8: Data Bus Inversion



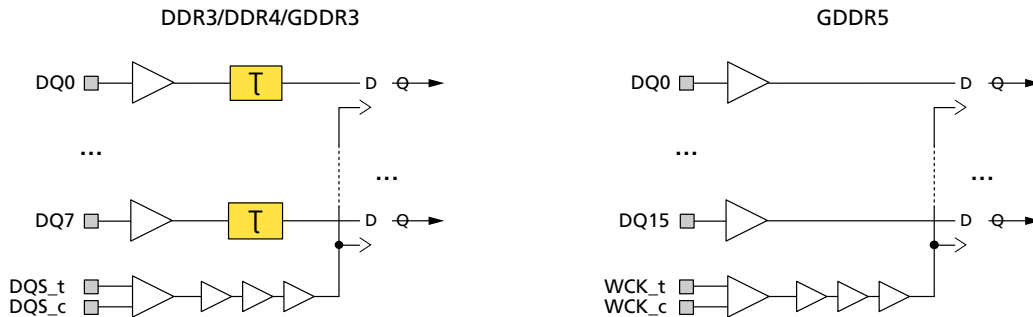
Write Data Latching and Clock Distribution

DDR3, DDR4, and GDDR3 devices latch write data using a data strobe (DQS) that is driven by the memory controller. The write data strobe is center-aligned with the write data to provide equal setup and hold times at the DRAM's receiver. The DRAM has to maintain this phase relationship. The phase relationship is achieved by adding delay elements in the latch's data path that match the clock path's insertion delay (see the block labeled "T" in the figure below). It is challenging to maintain accurate delay matching over the process, temperature, and voltage (PVT) variations. This scheme has proven to be effective with DDR3, DDR4, and GDDR3 devices, but it is considered inadequate for the data rates of GDDR5.

GDDR5 uses a scheme with direct latching data receivers and no delay matching between the data receiver and the WCK clock. The memory controller determines the optimum phase relationship between write data and the WCK clock for each data pin through data training (see Write Training).

The same differences apply to the read data path when you compare GDDR5 to older memory devices. Read data in DDR3, DDR4, and GDDR3 devices are edge-aligned with a data strobe. GDDR5 does not provide this kind of delay matching. As with write data, the memory controller determines the optimum phase for latching the read data. GDDR5 continuously drives a clock-like pattern on the EDC pins to the memory controller. The memory controller can use this pattern to adjust the internal strobe position.

Figure 9: Write Data Latching



PLL Characteristics

GDDR5 devices may be operated in PLL on or PLL off mode depending on system characteristics and operating frequency. PLL cancels a duty cycle error of the incoming WCK clock. It also suppresses high frequency WCK jitter that is above the PLL's bandwidth but tracks low frequency clock phase variation. The PLL bandwidth is programmable and is adjustable to system characteristics. However, a disadvantage of PLL is that it introduces jitter into the clock path.

The decision to use PLL on or PLL off mode is typically made during system qualification. In many cases, PLL off mode is preferred. It results in a fast frequency change procedure because PLL lock time does not have to be met. It also results in lower power consumption.


PCB Signal Routing

GDDR5 does not require delay matching for writes and reads. This is advantageous for the signal routing between the memory controller and the memory device.

The figure below compares the DDR3, DDR4, and GDDR3 signal routing topologies to GDDR5 signal routing topologies. The DDR3, DDR4, and GDDR3 signal routing attempts to achieve equal trace lengths for all signals that maintain the phase relationship between the data and the strobe, resulting in low pin-to-pin skew.

The data interface does not require this type of trace length matching. The skew between the data and the clock is compensated by the write and read data training. The advantage is a wider data eye resulting from a larger PCB area. This creates larger spacing between adjacent data lines and reduces cross talk and jitter.

Figure 10: PCB Routing with Unmatched and Matched Trace Length



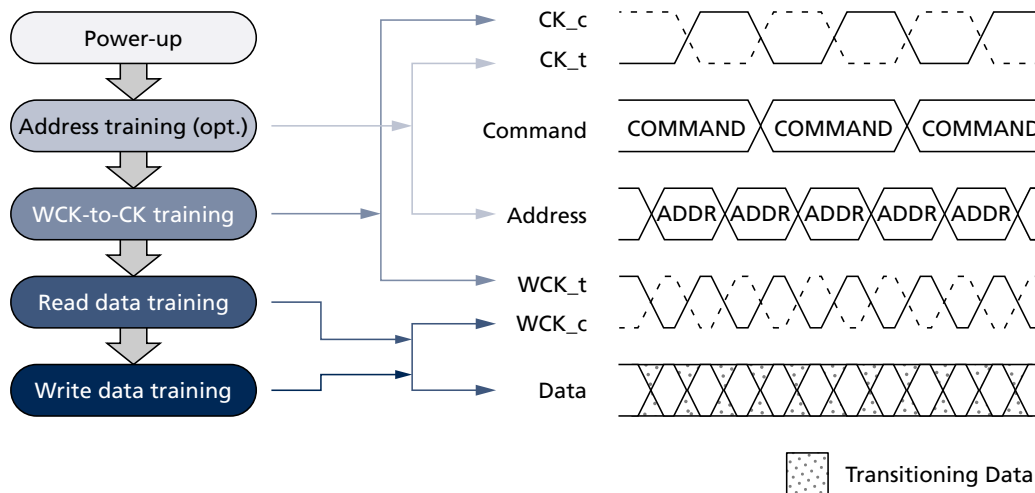
Adaptive Interface Training

GDDR5 provides hardware support for adaptive interface training. The purpose of this training is to ensure that the device operates with the widest timing margins on all signals.

All interface training is operated by the memory controller. The device assists the memory controller by offering several hardware features that result in fast and accurate training. The timing adjustments are made within the memory controller, not the DRAM.

If the steps in the figure below are followed in sequence, then all training steps can be performed at the application's maximum operating frequency.

Figure 11: Interface Training Sequence



Power-Up

The device configuration (x32 or x16 mode) and ODT for the address/command lines are set at power-up. When a stable CK clock is applied, the device is ready to receive commands. The command-pin timing has to be guaranteed by design and does not require training.

Address Training

Address training is optional and may be used to center the address input data eye.

Address training mode uses an internal bridge between the device's address inputs and DQ, DBI_n outputs. It also uses a special READ command for address capture. The address values registered coincident with this special READ command are asynchronously returned to the controller on the DQ and DBI_n pins. The controller compares the address pattern to the expected value and then adjusts the address transmit timing accordingly. The procedure may be repeated using different address pattern and interface timings. A WCK clock is not required for this special READ command during address training mode.

WCK2CK Training

WCK and CK clocks require a specific phase relationship that varies depending on the device. This phase relationship ensures a reliable phase-over of write data from the external WCK clock domain to the internal CK clock domain. Similarly, the same phase relationship ensures a reliable phase-over of read data from the internal CK clock domain, to the external WCK clock domain, and the output drivers. This helps to define READ and WRITE latencies between the device and the memory controller.

WCK2CK training is initiated by the controller. The controller sweeps the WCK clocks against the CK clock. The device responds by a static signal indicating an "early" or "late" clock phase. The optimum phase relationship is indicated by the transition from early to late phase.

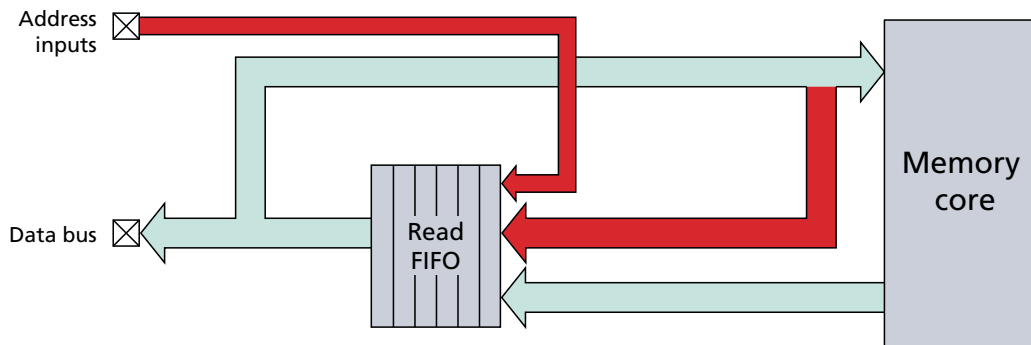
In most applications, the trained WCK2CK phase relationship provides sufficient margin to cover any drift that occurs during system operation. However, a new WCK2CK training is required if there are any frequency changes or changes in the PLL on/off mode.

Read Training

Read training enables the memory controller to find the data eye center (symbol training) and burst frame location (frame training) for each high-speed output of the device. Read training is the first step in aligning the data bus to the WCK clock. This involves two characteristics:

1. The alignment of the latching clock in the memory controller to the center of the read data bit (bit training).
2. The detection of burst boundaries out of a continuous read data stream (framing).

Read and write data training does not require access to the slower memory array. Specific training commands utilize the read FIFO that typically functions as temporary storage for read data. The figure below shows the data paths and additional paths for data training.

Figure 12: Read and Write Data Training Data Paths


Initially, the FIFO is preloaded with data that is safely transmitted over the previously trained address bus (LDFD command). Once the FIFO is preloaded, special READ commands that return the FIFO data to the controller are repeatedly issued. Then, the controller sweeps its clock phase until the data is correctly sampled.

Write Training

Write training enables the memory controller to find the data eye center (symbol training) and burst frame location (frame training) for each high-speed input of the DRAM.

Write training is the final step in aligning the data bus to the WCK clock. It includes the same characteristics of read training:

1. The alignment of the latching clock in the DRAM to the center of the write data bit (bit training).
2. The detection of burst boundaries out of a continuous write data stream (framing).

Knowing that the read path has been trained before, the controller writes and reads data to and from the read FIFO and sweeps the write data phase until the data is written correctly. After write training, all data eyes are expected to be centered and the device is ready for normal operation.

Continuous Tracking

Due to GDDR5's high data rates, even small changes in supply voltage or temperature gradually shift the write and read data eye position away from the trained optimum. This shift makes transmission errors more likely. The controller is able to observe and compensate this data eye drift by monitoring the EDC pin, which can be programmed to send a clock-like pattern (EDC hold pattern) continuously to the controller. This is known as clock and data recovery (CDR).

To re-center the data eye, the memory controller repeats write training and read training at regular intervals. GDDR5 allows this training in parallel with an ongoing regular REFRESH operation; a period in other DRAM devices, the data bus is idle. Carefully implemented training during refresh does not result in lower performance.

High-End and Low-Cost Systems

The amount and accuracy of training depends on the target data rates and system characteristics. A high-end graphics card will require all training steps with the highest possible accuracy to tweak the data rate to maximum levels. This includes a per-bit training on the data lines that will cancel out any differences in signal flight times in the individual data lines.

Systems that do not require the highest data rates may skip address training and perform per-byte training or use more coarse resolution in the timing adjustment. At lower data rates, minor differences in signal flight times or minor training inaccuracies may be acceptable. This usually results in cost-effective and power-optimized memory controller design.

Data Integrity

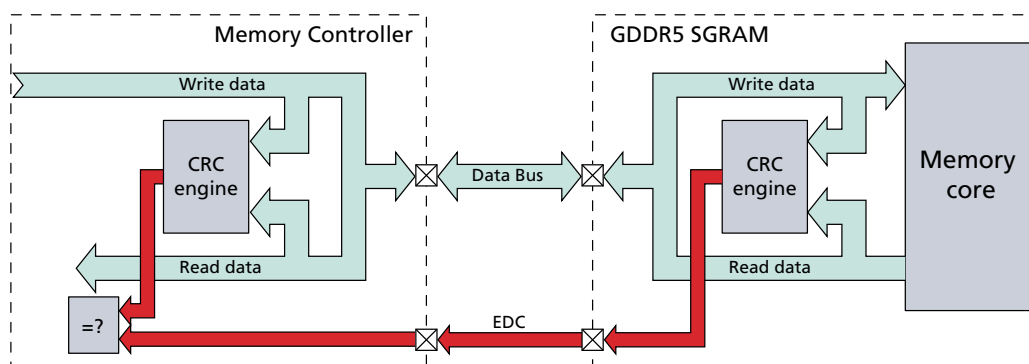
GDDR5 SGRAM manifold hardware features and training algorithms ensure reliable operation at very low bit error rates (BER). However, some critical applications require BER to be significantly lower than consumer applications. The device addresses this requirement in two ways:

- Securing the signal integrity of the high-speed I/Os by adding redundancy
- Securing partial WRITE operations by using a safe path for conveying the write data mask

Error Detection and Correction

GDDR5 supports error detection and correction (EDC) on its bidirectional DQ and DBI_n lines using a cyclic redundancy check algorithm (CRC-8). This algorithm is widely accepted in high-speed communication networks. The algorithm detects all single and double bit errors.

Figure 13: Error Detection and Correction



GDDR5 calculates the CRC checksum for each READ or WRITE burst and returns the checksum to the controller on the dedicated EDC pin. The controller performs the same CRC calculation: If both checksums do not match, the controller assumes that there is a transmission error, and it is designed to repeat the command that has the error.

The procedure is asymmetric. Only the controller performs the CRC check and takes corrective actions. The DRAM executes the command, regardless of whether there is a CRC error or not.

This EDC feature can be used as a data eye drift indicator and trigger a retraining. However, a safer procedure is scheduling a retraining on a regular basis and using the EDC capability as an additional safeguard.

Write Data Mask

DRAM devices support partial WRITE operations when individual bytes may be masked. These WRITE operations are the equivalent to READ MODIFY WRITE operations but consume less memory bandwidth.

The data mask (DM) information is usually conveyed on an extra data mask pin that is associated with each data byte. The disadvantage of this scheme is that bit errors on the DM signal are not recoverable. Therefore, a masked byte may be mistakenly overwritten if the DM signal is flipped.

The EDC feature does not solve this issue because the failure would be detected by the controller after the actual write. Therefore, the GDDR5 device implements a safer scheme that transmits the masking information via the slower address bus. Special WRITE commands support single- and double-byte mask granularity.

Memory Core

Memory Organization

GDDR5 uses an $8n$ -prefetch architecture to achieve high-speed operation. With $8n$ -prefetch architecture, the internal data bus to and from the memory core is eight times as wide as the I/O interface but is operated at only one-eighth of the I/O data rate.

Table 1: Addressing Scheme

Address	2Gb		4Gb		8Gb	
Memory organization	64Mb x32	128Mb x16	128Mb x32	256Mb x16	256Mb x32	512Mb x16
Row address	A[12:0]	A[12:0]	A[13:0]	A[13:0]	A[13:0]	A[13:0]
Column addresses	A[5:0]	A[6:0]	A[5:0]	A[6:0]	A[6:0]	A[7:0]
Bank address	BA[3:0]	BA[3:0]	BA[3:0]	BA[3:0]	BA[3:0]	BA[3:0]
Bank groups	4	4	4	4	4	4
Page size	2KB	2KB	2KB	2KB	4KB	4KB

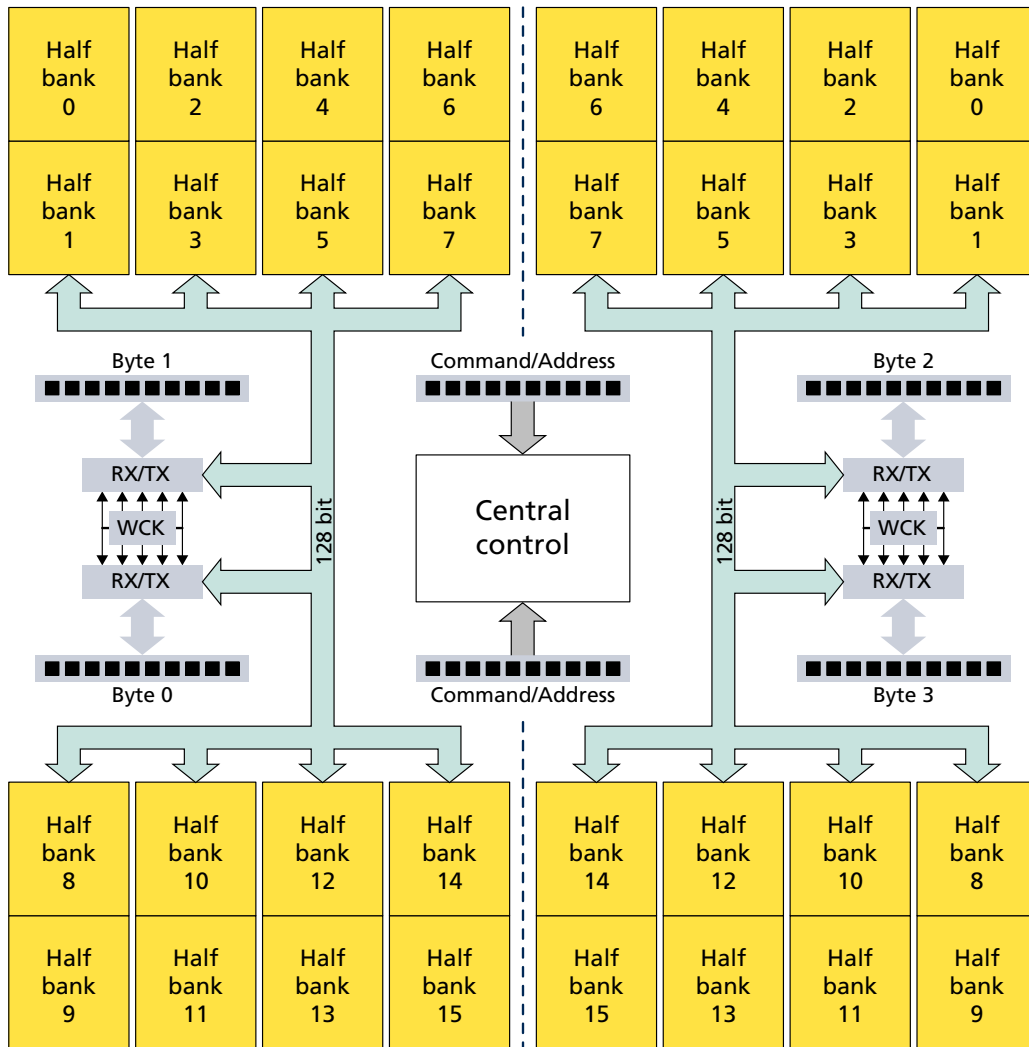
Outer Data, Inner Control Architecture

The outer data, inner control (ODIC) chip architecture is reflected by the ballout:

- The 32-bit data interface is physically split into four bytes. One byte is located in each quadrant of the package: Bytes 0 and 1 and bytes 2 and 3 each share WCK clocks; both sections are physically separated with no data lines that cross the chip center.
- The address, command, CK clock, and other control signals are located in the die center.

The advantages of ODIC architecture are shorter internal WCK clock trees and high-speed data lines (see the figure below), resulting in extremely low on-die jitter and excellent device supply noise immunity.

Figure 14: ODIC Architecture



Memory Core Speed and Bank Groups

GDDR5's high-speed memory core is another characteristic that contributes to its superior performance. A 750 MHz memory core is required to transfer eight data words per READ/WRITE command at 6 Gb/s within two CK clock cycles. DDR3 and DDR4 memory cores typically operate at speeds of 200 MHz–250 MHz.

DDR4's higher data rates (compared to DDR3) are not the result of a faster memory core; they are due to the introduction of bank groups that require seamless accesses be directed to different banks or bank groups. This bank group restriction typically results

in a performance loss caused by higher latencies or delayed READ or WRITE commands.

Micron's current GDDR5 devices do not require the use of bank groups.

Power-Saving Features

GDDR5 features and device operation enable lower power consumption. To estimate the potential power savings, consider the power saving of the device and the interface.


Supply Voltage

The device operates from a 1.5V supply voltage. It also supports 1.35V, at a slightly lower data rate. Micron is working with its customers to further reduce supply voltage.

Dynamic Voltage Scaling

The device supply voltage (DVS) can be changed on-the-fly between 1.5V and 1.35V, and the system's power consumption can be scaled to the actual system workload. The voltage transition occurs when the DRAM is in self refresh mode (see the figure below). The voltage transition duration is determined by the characteristics of the voltage regulator and the onboard buffer caps.

Figure 15: Dynamic Voltage Scaling




Dynamic Frequency Scaling

The device can operate over a wide frequency range, starting at 200 Mb/s. While 400 Mb/s is sufficient for displaying static images from a web browser or e-mail client, a data rate of 1.5 Gb/s may be required for HD video playback, and high end gaming applications may require the maximum data rate.

The memory system's power consumption depends on the clock frequency. Micron recommends scaling the clock frequency to the actual required memory bandwidth.

Figure 16: Supply Current vs. Data Rate



On-Die Termination

The signal lines are typically terminated with an impedance of 60Ω. At lower data rates, it might be possible to achieve stable operation by using a termination of 120Ω or by completely disabling on-die termination (ODT). In both cases, system power is reduced. The device allows independent control of the ODT value for address/command and data.

WRITE Latency

WRITE latency is the delay between a WRITE command and the start of a WRITE burst. When the latency is set to small values (for example, WL = 3), the input receivers remain enabled. When set to large values (WL = 6 or 7), the input receivers turn on for the duration of a WRITE burst only. Power savings with larger WL values is possible because WRITE bursts only account for a small percentage of the overall memory transactions. The performance penalty of a higher WRITE latency is negligible.

Power-Down and Self Refresh

To save power during idle states, the device supports power-down and self refresh modes.

Power-down disables the input buffers and internal clock trees, while the external CK and WCK clocks remain active to keep the DRAM's PLL and internal synchronization logic in a locked state. Power-down supports a fast exit to quickly react to a new memory request.

The self refresh state retains stored information without external interaction. Exiting from self refresh takes longer than exiting from power-down because the CK and WCK clocks need to be re-synchronized and the PLL must re-lock. The device also supports temperature-compensated self refresh mode that further reduces the power consumption at lower operating temperatures.

Conclusion

Offering ultra-high bandwidth, improved data integrity compared to older DRAM devices, and manifold features to control power consumption, GDDR5 SGRAM is the ideal device for graphics cards, game consoles, and high-performance computing systems.

8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900
www.micron.com/productsupport Customer Comment Line: 800-932-4992
Micron and the Micron logo are trademarks of Micron Technology, Inc.
All other trademarks are the property of their respective owners.