Micron® DDR5 SDRAM: New Features

By Randall Rooney and Neal Koyle

Introduction

This white paper is a follow-up to Micron's earlier DDR5 white paper titled, "Introducing Micron® DDR5 SDRAM: More Than a Generational Update," which highlighted key fifth-generation double data rate (DDR5) SDRAM features and functionality that delivers significant performance improvements over DDR4. In this paper, we provide further detail about improved performance, RAS and ease of implementation compared to prior generations of SDRAM.

Performance

Overall Bank Increase

When memory density increases, there is a need to expand the number of banks to account for the increased memory density. The DDR5 standard doubles the number of bank groups while leaving the number of banks per bank group the same. This improves the overall system efficiency by allowing more pages to be open at any given time and by increasing the statistical probability of high page-hit ratios.

Bank-group-to-bank-group interleaved timing accesses are shorter than bank-to-bank within a specific bank group access. These timing parameters have "long" timing definitions (^{t}CCD_L, ^{t}WTR_L, ^{t}RRD_L) and "short" timing definitions (^{t}CCD_S, ^{t}WTR_S, ^{t}RRD_S). The long timings refer to bank-to-bank within a bank group, while the short timings refer to accessing different bank groups (as illustrated in Figure 1). To give some perspective, ^{t}CCD_L can be nearly double ^{t}CCD_S. Increased bank groups mitigate internal timing constraints by increasing the probability that the short timings are in use.


Figure 1: DDR5 Bank/Bank Group Timings

Data Burst Length Increase

DDR5 SDRAM default burst length increases from BL8 (seen on DDR4) to BL16 and improves command/address and data bus efficiency. The same read or write CA bus transaction can now provide twice as much data on the data bus while limiting the exposure to IO/array timing constraints within the same bank. Reducing the commands required to access a given amount of data also improves the power profile for read and write accesses.

The burst length increase also reduces the number of IOs required to access the same 64B cache line data payload. The default burst length increase enables a dual sub-channel for the DDR5 DIMM architecture (shown in Figure 2), which increases overall channel concurrency, flexibility and count. For systems that utilize a 128B cache line data payload, DDR5 adds a burst length of 32 option specifically for x4-configured devices. This further improves the command/address, data bus efficiency and overall power profile.
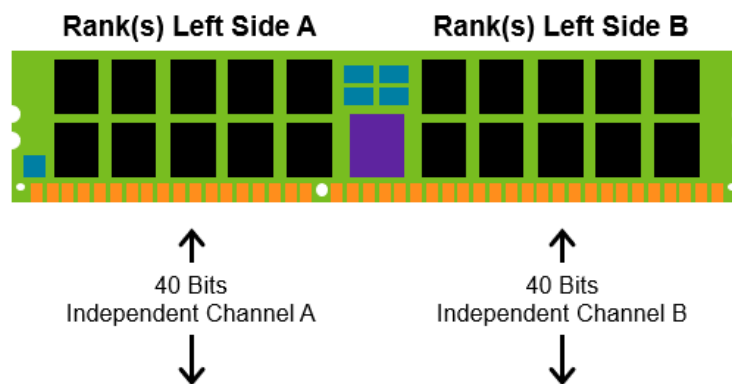

Figure: 2: Simplified DDR5 40-Pin Sub-Channel DIMM Example

Refresh Commands

In addition to the standard ALL-BANK REFRESH command (REFab) available on DDR5 and earlier DDR SDRAM products, DDR5 introduces a SAME-BANK REFRESH (REFsb) command. The REFsb command targets the same bank in all bank groups, as designated by bank bits via command/address bits when the REFsb command is issued.

REFRESH commands on SDRAM devices require that the banks targeted for refresh are idle (precharged, no data activity) prior to the command being issued, and the banks cannot resume subsequent write and read activity for the duration of the REFRESH command (timing parameter ^{t}RFC). REFRESH commands are issued at an average periodic interval (timing parameter ^{t}REFI). For REFab commands, the system must ensure all banks are idle prior to issuing the command, on an average of once every 3.9µs in "normal" refresh mode, with a duration of 295ns for a 16Gb DDR5 SDRAM device.

The performance benefit of the REFsb command is that only one bank in each bank group needs to be idle before issuing the command. The remaining 12 banks (for a 16Gb, x4/x8 device; blue cells in Figure 3) do not have to be idle when the REFsb command is issued, and the only timing constraint to non-refreshed banks is the same-bank-refresh-to-activate delay (timing parameter ^{t}REFSBRD). REFsb commands can only be issued in the fine granularity refresh (FGR) mode, meaning each bank must receive a REFRESH command every 1.95µs on average. The REFsb duration is only 130ns for a 16Gb DDR5 SDRAM device, which also reduces the system access lockout (^{t}RFCsb) to actively refreshing banks (red cells in Figure 3). A restriction when using REFsb is that each "same bank" must receive one REFsb command prior to that "same bank" being issued a second REFsb command, but the REFsb commands can be issued in any bank order.

Figure: 3: DDR5 REFsb Bank Mapping

Depending upon the read/write command ratio, simulations indicate a 6% to 9% increase in system performance throughput when using REFsb as compared to REFab, as shown in Figure 4. Furthermore, REFsb reduces the refresh impact to average idle latency from 11.2ns to 5.0ns, as highlighted in Table 1. Calculations are based on standard queuing theory and are applicable for a single bank with randomly driven data traffic.


Figure 4: DDR5 System Throughput Performance Improvement

Refresh Adder to Average Idle Latency	
REFab	11.2ns
REFsb	5.0ns

Table 1: DDR5 Average Idle Latency Adder

Performance Improvement

Taking the features above (2x banks, 2x bank groups, BL16, and same bank refresh) and simulating a 64B random access workload, substantial performance increases are realized compared to DDR4 dual-rank modules at 3200 MT/s, as shown in Figure 5. In this scenario we assume eight channels per system and 1DPC.


Figure 5: DDR5 Performance Improvement by Speed/Module Ranks

Reliability, Availability and Serviceability (RAS)

On-Die Error Correction Code (ECC)

RAS improvements like on-die ECC reduce the system error correction burden by performing correction during READ commands prior to outputting the data from the DDR5 device. DDR5 SDRAM ECC is implemented as single error correction (SEC), pairing 128 data bits with 8 parity bits to form a 136-bit codeword that is stored in the DRAM during a WRITE command. During subsequent READ commands to that address, a syndrome will be calculated based on the 136 bits, correcting any single-bit errors that may occur.

DDR5 designs implement the ECC with Hamming codes, where data bits are arranged into quadrants to align with system-level error correction coverage. Since 8 parity bits with 128 data bits do not allow for double-bit detection, the Hamming codes will "alias" a syndrome bit for two error bits within a pair of quadrants, into the other quadrants or into an "unused" data bit. This means that two error bits within quadrants 1 and/or 2 will alias to quadrants 3 or 4, or to an unused

data bit. Likewise, two error bits within quadrants 3 and/or 4 will alias to quadrants 1 or 2, or to an unused data bit. This allows the errors to still appear as a double bit fail to the system-level error correction.

An additional feature of the DDR5 SDRAM ECC is the error check and scrub (ECS) function. The ECS function is a read of internal data and the writing back of corrected data if an error occurred. ECS can be used as a manual function initiated by a Multi-Purpose Command (MPC), or the DDR5 SDRAM can run the ECS in automatic mode, where the DRAM schedules and performs the ECS commands as needed to complete a full scrub of the data bits in the array within the recommended 24-hour period. At the completion of a full-array scrub, the DDR5 reports the number of errors that were corrected during the scrub (once the error count exceeds a minimum fail threshold) and reports the row with the highest number of errors, which is also subject to a minimum threshold.

PPR Enhancements

Post-package repair (PPR) is broken into two separate repair features, hPPR (hard) and sPPR (soft), which may be better described as permanent repair (hPPR) and temporary repair (sPPR). hPPR is nonvolatile with power cycling and sPPR is not.

The DDR4 SDRAM definition had several rows per bank, called associated rows, which were required to be backed-up prior to an sPPR event. One key DDR5 enhancement for sPPR is the reduction of rows which need to be backed-up before performing an sPPR repair. DDR5 requires only one target row in the bank where the sPPR will occur. This minimizes the system time required to back up and store a large amount of information, typically in the range of ~2µs per row of data.

Another key feature added to PPR is the ability to track resource availability. At boot-up, each DRAM device will determine the availability of a PPR resource in each bank and then set a group of mode registers (MR54-57) to track this information. In the case where a multi-die 3DS stacked package is used, each die in the multi-die 3DS stacked package will be tracked via the same mode registers. This enhancement gives added visibility to RAS capabilities of the memory base of any given deployed system.

Implementation Simplification

Multi-Purpose Command (MPC)

Increased clock frequencies on the DDR5 device have created challenges for performing operations prior to the completion of initialization and training. To remedy those challenges, the Multi-Purpose Command (MPC) was developed to perform functions like interface initialization, training and periodic calibration. On earlier DDR SDRAM device generations, these functions were either individual commands or mode register settings.

The MPC function is initiated when the command/address bus contains the MPC instruction code, as well as eight operation bits, allowing for up to 256 different functions to be performed. The default MPC mode setting is for multiple clock cycles with CS_n LOW, where the command/address bus is held constant throughout the multicycle period. This allows for the CS training mode to be entered and completed via the MPC. Once CS_n is aligned with CK, the MPC can be switched to single-cycle mode for improved operational efficiency.

Aside from performing initialization, training and calibration functions, the MPC is used for resetting the DLL, switching between 2N and 1N command timings, PDA enumerate programming and ID selection, setting termination values for on-die termination, and configuring DLL and other array timings.

Conclusion

The DDR5 architecture will increase the value that Micron brings to the industry with our DRAM solutions. Significant performance improvements, made possible in part by increasing bank groups, burst lengths and same-bank refresh will help to meet the stringent requirements of next-generation systems and improve the total cost of ownership. The system RAS is improved with the DDR5 on-die error correction code and post-package repair enhancement. Finally, the implementation of the memory management is simplified with the use of the new MULTIPURPOSE command feature.

Excitement continues to build around the possibilities that DDR5 will offer for computing systems, and Micron is ready to engage with system architects to help them reach the full potential of this new product line.

micron.com