

DDR4 vs. DDR5



Introduction to DDR Memory

DDR, or Double Data Rate, refers to a memory bus that can send signals or data on both sides of a clock cycle. If this clock cycle is externally coordinated it is known as synchronous, as in SDRAM – DRAM, or

Dynamic Random-Access Memory, being a type of volatile memory. Volatile memory loses its data when power is lost and, further, DRAM requires that its data is consistently refreshed, hence “dynamic.” Static RAM, or SRAM, is faster but costlier and is instead utilized for things such as microcontroller unit (MCU) cache on solid state drives (SSDs).

DRAM is most commonly used for the main system memory, or RAM, in a system, with speeds given by the JEDEC designation. For example, the designation will have DDR4 or DDR5 denoting the generation with the transfer rate following, such as 2133 or 4800. This transfer rate is given in millions of transfers per second, or megatransfers per second (MT/s), as opposed to clock rate or bandwidth. The MT/s value will be double the clock rate with true DDR, for example 3200 referring to 1600 MHz memory. As data is transferred by 8 bytes the bandwidth will be eight times this amount, as with PC4-25600 from DDR4-3200.

In a typical client system DRAM is the primary system memory that lies between the central processing unit (CPU) and its superfast cache, facilitating data operations in contrast to data stored in slower non-volatile media (NVM). This memory is arranged in one or more channels for parallelization and offers a certain amount of bandwidth, as found on other devices like graphical processing units (GPUs). Also of importance to performance is the latency of the memory, given by a multitude of values depending on the corresponding internal operation. Memory is managed by the CPU via an integrated memory controller (IMC) with other elements handled by the motherboard.

Architectural Differences

Although there are consistent differences between generations of DRAM – for example, a transition to lower voltages in order to improve efficiency – there are more changes between DDR4 and DDR5 than is typical. These changes are largely beneficial, as detailed in the next section. Fundamentally there are also changes required at the motherboard level, although CPUs – such as Intel’s Alder Lake ([read our blog](#)) – can use either DDR4 or DDR5 when paired with the correct hardware.

DDR5 has much higher latency, for example with column access strobe (CAS) latency, which is mitigated through higher transfer rates but also architectural differences. While both types of memory use a 288-pin

interface the notch position has been changed in order to prevent improper installation. DDR4 and DDR5 have other distinct differences – as listed in the table below – but DDR5 also institutes some novel features. While this includes more basic improvements for sideband access, memory training, pin definition, and temperature monitoring, there are some bigger changes as well.

For example, it has power management (via PMIC) on-module (dual in-line memory module, or DIMM) rather than through the motherboard which ostensibly lowers motherboard cost and complexity. This pulls from 5V on client machines and helps improve signal integrity as well. It also has on-die error correction (ODECC) which protects data on the modules from bit flips, distinct from channel or end-to-end protection. DDR5 also has the Same Bank Refresh Function (SBRF) which allows independent refreshing per bank group; this theoretically improves the ability to access memory.

DDR4 vs. DDR5 at a glance	
DDR4	DDR5
16Gb memory chips	64Gb memory chips
2133 MT/s baseline	4800 MT/s baseline
1600-3200 basic range	3200-6400 basic range
1x64-bit structure	2x32-bit structure
8-bit ECC (72-bit)	16-bit/2x8-bit ECC (80-bit)
8-byte burst (BL8)	16-byte burst (BL16)
1.2V nominal (VDD/VDDQ)	1.1V nominal (VDD/VDDQ)
XMP 2.0 (3 profiles)	XMP 3.0 (5 profiles + modes)
16 banks (4x4)	32 banks (4x8)
	On-die PMIC ODECC SPRF

Advantages & Disadvantages

Taking these changes one by one, the ability to use denser memory chips – and, in fact, stack multiple dies on top of each other, up to 8 per package for server usage – allows DDR5 to easily achieve 128GB (64Gbx16) per module in the client space. The higher clock and transfer rate means a doubling or more of bandwidth, particularly useful for content creation. While this comes at the cost of higher latency as mentioned above, the move to a dual 32-bit internal structure allows for more banks and therefore a doubling of simultaneous

open pages which reduces effective latency. In order to produce the same size payload as DDR4 with this structure DDR5 necessarily employs a higher burst length (BL), however.

While both DDR4 and DDR5 support additional optional ECC, that is channel ECC, with 8 bits of ECC per bus, the 2x32-bit structure of DDR5 takes the total width from 72-bit to 80-bit. This is not particularly relevant for client usage, however. DDR5 is also more efficient through using lower voltage, but more important for client usage is the improvement of extreme memory profile (XMP) profiling. Aside from allowing more total profiles, XMP 3.0 allows for manual editing of these profiles and additionally a dynamic memory boost function which operates like turbo mode on CPUs. These profiles are typically taken via serial presence detect (SPD).

Summary

DDR5's improvements are exciting for both client and server environments. The traditional improvements found with lower voltage and higher bandwidth are here, although with the typical increase in latency, but other architectural changes promise a bit more. The ability to use denser chips has increased capacity four-fold. The move to a 2x32-bit versus 1x64-bit structure allows for better memory access availability. Locating the PMIC on-module allows the memory to scale better with improved voltage regulation. ODECC further helps improve data integrity. XMP improvements also encourage extra flexibility, including potential power savings with dynamic profiles.

The first client platform to utilize DDR5, Alder Lake on the Z690 chipset, has already shown the impact of these improvements. The bandwidth in particular is a boon for content creation and other data-intensive tasks, but other improvements are already showing potential gains for other areas such as gaming. New generations of DRAM tend to start off slow and expensive and that is still the case here, particularly due to current supply issues. However, the potential is there and can already be realized by early adopters, although there will likely have to be changes within the industry to fully make use of DDR5's advantages. This paradigm shift of architecture opens the door to more efficient machines that are able to finally feed the multi-core future.

*All product and company names may be trademarks or registered trademarks of their respective holders.

Our SSD Solutions



**CA6 Series | PCIe®
Gen 4**

- Slim form factor— M.2 2280
- Random read/write up to 1000K/1000K IOPS
- Low latency
- LDPC technology



**CL4 Series | PCIe®
Gen 4**

- Slim form factor— M.2 2230/2242/2280
- Random read/write up to 450K/400K IOPS
- Low latency
- 256GB - 1TB

Please contact our [Solid State Storage Technology Corp. expert](#) for more information.

*Specifications and features are subject to change without prior notice. Images are samples only, not actual products.

Request Full Specs Sheets



ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

© 2022 Solid State Storage Technology Corporation. All rights reserved.

Learn more at www.ssstc.com

[Report abuse](#)

Created with  **mailchimp**