

## Advanced Memory Device Correction (AMDC) for Servers

Sudhanva Gurumurthi, RAS Architect, AMD

# Advanced Memory Device Correction (AMDC) for Servers

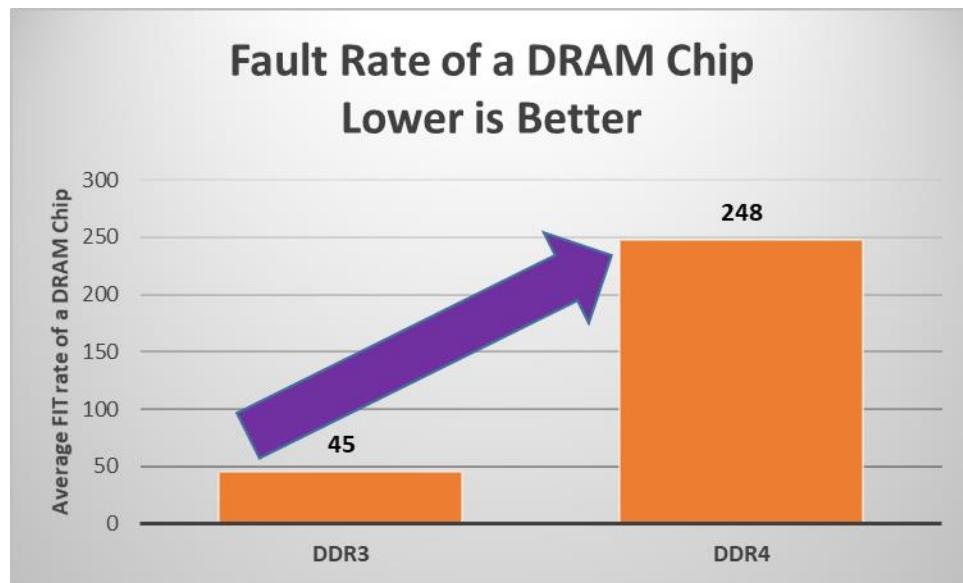
Memory Reliability, Availability, and Serviceability (RAS) is important in servers. AMD's studies of production data centers have shown that a variety of faults occur in DRAM, from those that impact a single bit of data, to those that impact multiple bits within a DRAM chip<sup>1,2,3</sup>. It is important to correct as many of these faults as possible in the hardware to minimize the impact on software running on these servers and reduce costs at the data center level. This article provides an overview of AMD's Advanced Memory Device Correction (AMDC) for server Dual-Inline Memory Modules (DIMMs).

## Executive Summary

DRAM chip reliability is ~5.5x worse in DDR4 compared to DDR3. AMD EPYC™ second-generation processors with AMDC are designed to mitigate this problem by correcting errors originating from a single DRAM chip on a x4 ECC DIMM.

## DRAM Reliability is Getting Worse

The rate at which faults occur in DRAM is the key measure of its reliability at the technology level. *Figure 1* shows the reliability of a JEDEC DDR4 DRAM chip, the latest commercially available server memory technology, compared to a previous generation DDR3 chip\*. The fault rate is expressed in Failures in Time (FIT)\*\*.



*Figure 1* - The graph shows the fault rate of a DRAM memory chip in DDR3 and DDR4.

We can observe that the DDR4 chip fault rate is approximately 5.5x the DDR3 fault rate.<sup>2,5</sup>

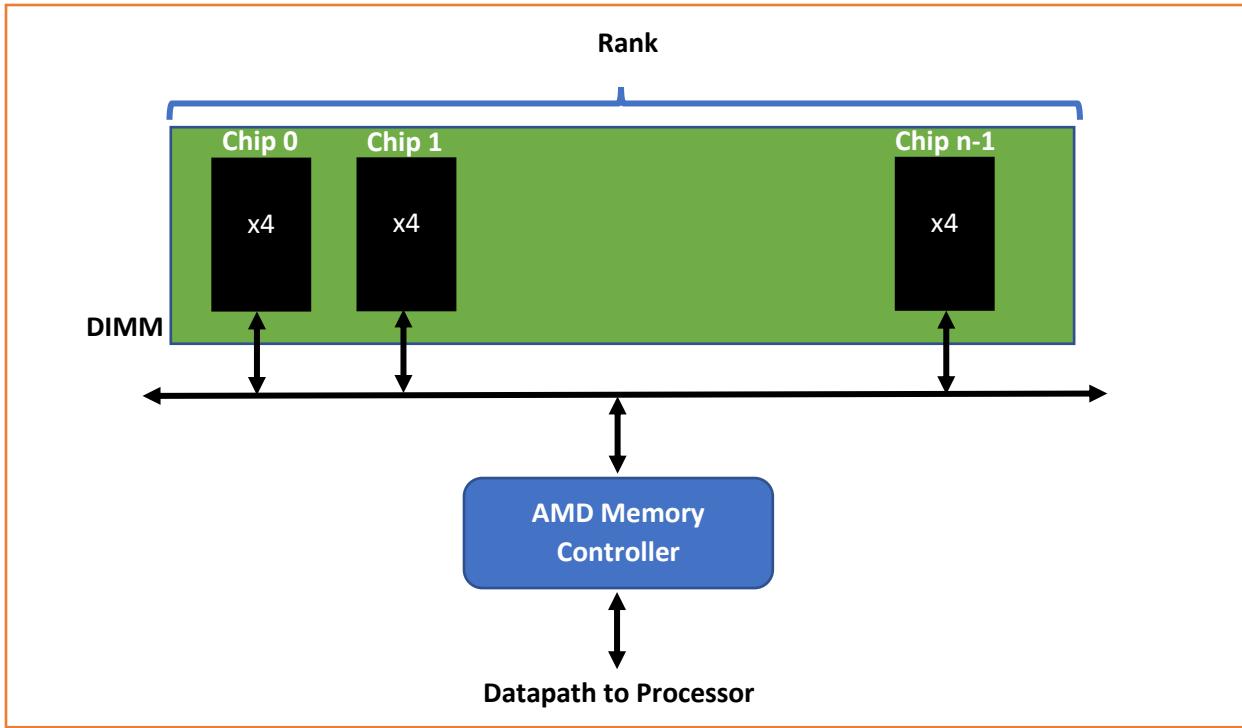
\* The DDR3 data is from a previously published study by AMD<sup>2</sup>. The DDR4 data is from an internal study carried out by AMD in a production data center<sup>5</sup>. The data shown in the graph for each memory generation is an average over DRAM parts from multiple vendors.

\*\* One FIT is defined as one failure every billion hours and is a standard reliability metric.

## Overview of Error Correcting Codes for DRAM

An illustration of a DIMM is shown in *Figure 2*. A DIMM consists of several DRAM memory chips. Each chip is part of a “rank”. All chips in a rank are accessed in parallel when reading from or writing to memory. Since each chip in a rank contributes part of the data required for a memory access, it is important to minimize errors in the data sent to the processor that were caused by faults in the DRAM chips. This requirement can be fulfilled by hardware using an Error Correcting Code (ECC).

An ECC is a mathematical function that takes the data bits as input and transforms the data into a new representation. This new representation of the data is then written to memory. When this transformed data is later read by the processor, this representation is fed through the same mathematical function. Based on the specific ECC algorithm, the function can detect and possibly correct any errors in the data sent to the processor.



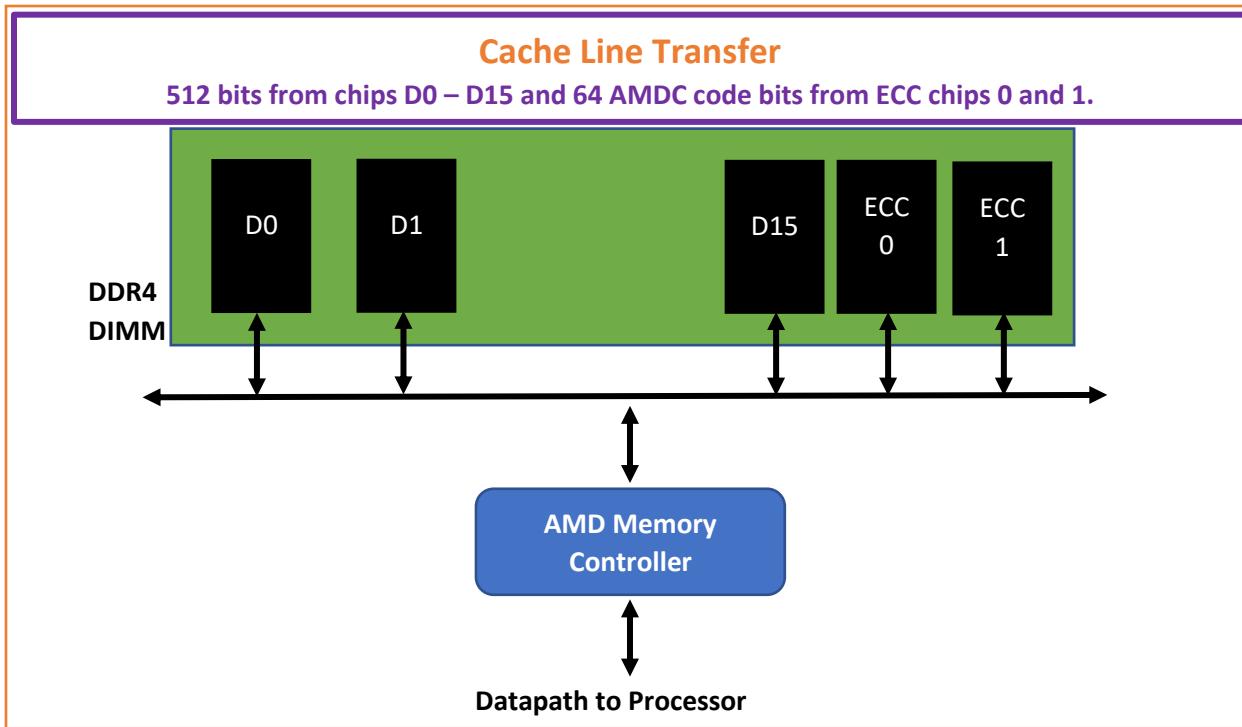
*Figure 2 - DIMM memory components and access. The DIMM shown here consists of one rank that has 'n' memory chips. Each chip is 4-bits wide. When the rank is accessed, each chip contributes 4 bits of data during a DRAM burst. The memory controller in the AMD processor implements ECC algorithms designed to detect and correct errors in the DRAM chips and the interface between memory and the controller before the data is sent to the processor.*

## AMDC Overview

AMD's server processors implement a memory ECC algorithm known as Advanced Memory Device Correction (AMDC). AMDC is designed to correct any number of faults that occur within a single DRAM chip in a rank. AMDC achieves this correction capability by using a type of ECC that allows large groups of bits to be corrected with negligible memory performance impact. For example, consider the scenario

where Chip 1 in *Figure 2* has a fault that can cause multiple bits of error if a memory request accesses the fault. AMDC is designed to correct those errors before the data is sent to the processor.

AMDC works with standard DDR4-compliant ECC RDIMMs. Such RDIMMs are composed of DRAM chips that contain data and ECC bits. *Figure 3* shows a DDR4 DIMM rank that uses AMDC. The rank is composed of x4 chips. A memory request accesses 16 data chips and 2 ECC chips per rank. For the DDR4 burst length of 8 (BL=8), 512 bits are transferred from the data chips D0-D15 and 64 AMDC code bits are transferred from the ECC chips. Thus, AMDC uses existing ECC chips on RDIMMs to provide the data protection.



*Figure 3 - A DDR4 ECC RDIMM with x4 data and ECC chips. AMDC uses the existing ECC chips on the RDIMM to provide memory protection.*

AMD has collected several billions of hours of DRAM data from production data centers and this field data provides insightful statistics on the type of faults that occur in real DRAM chips. The data has shown that there are a variety of fault patterns that occur within a DRAM chip<sup>1,2,3</sup>.

AMDC is designed to correct errors originating from any number of faults within a single DRAM chip, including those ranging from single-bit faults to those spanning single columns, rows, banks, and even multiple banks within a DRAM chip.

**The result is that AMD Advanced Memory Device Correction (AMDC) helps data centers limit uncorrectable memory faults, and minimize their hardware, IT, and maintenance costs.**

## References

<sup>1</sup> V. Sridharan, N. DeBardeleben, S. Blanchard, K. Ferreira, J. Stearley, J. Shalf, S. Gurumurthi, Memory Errors in Modern Systems: The Good, The Bad, and the Ugly, Proceedings of International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), March 2015.

<sup>2</sup> V. Sridharan, J. Stearley, N. DeBardeleben, S. Blanchard, S. Gurumurthi, Feng Shui of Supercomputer Memory - Positional Effects in DRAM and SRAM Faults, Supercomputing (SC), November 2013.

<sup>3</sup> V. Sridharan and D. Liberty, A study of DRAM failures in the field. International Conference on High Performance Computing, Networking, Storage and Analysis (SC), November 2012.

<sup>4</sup> L. Barroso, U. Hölzle, P. Ranganathan, The Datacenter as a Computer: Designing Warehouse-Scale Machines – Third Edition, Synthesis Lectures on Computer Architecture, Morgan and Claypool Publishers, 2019.

<sup>5</sup> Based on internal AMD study from July to November 2018 of 655 servers in an AMD data center, with 423 million DRAM hours and 86 million CPU hours.

## Disclaimer

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

AMD, the AMD Arrow logo, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

© 2020 Advanced Micro Devices, Inc. All rights reserved.

Published June 2020