The Known Failure Mechanism in DDR3 Memory referred to as Row Hammer

Recently publicized, this server memory problem can cause undetected data corruption and undermine the world's confidence in Cloud Computing

9/4/2014
FuturePlus Systems Corporation
Barbara P Aichinger

The Known Failure Mechanism in DDR3 memory called "Row Hammer"

Barbara Aichinger FuturePlus Systems Corporation barb.aichinger@futureplus.com 603-472-5905

Computer Architecture relies on three basic building blocks. The CPU or central processing unit, the I/O, Input and Output and the Memory. Without these 3 basic building blocks computers as we know them do not exist. When it comes to the memory the dominate technology is DRAM or Dynamic Random Access Memory. The latest and most prevalent version of that is called DDR3 which stands for the 3rd generation of Double Data Rate Memory. In the quest to get memories smaller and faster memory vendors have had to make trade offs. One of these is very small physical geometries. These small geometries put memory cells very close together and as such one memory cell's charge can leak into an adjacent one causing a bit flip. It has come to the attention of the industry that this is indeed happening under certain conditions. Very simply the problem occurs when the memory controller under command of the software causes an ACTIVATE command to a single row address repetitively. If the physically adjacent rows have not been ACTIVATED or Refreshed recently the charge from the over ACTIVATED row leaks into the dormant adjacent rows and causes a bit to flip. This failure mechanism has been coined 'Row Hammer' as a row of memory cells are being 'hammered' with ACTIVATE commands. Once this failure occurs a Refresh command from the Memory Controller solidifies the error into the memory cell. Current understanding is that the charge leakage does not damage the physical the memory cell which makes repeated memory tests to try to find the failing device useless.

DDR3 memory is pervasive and used in nearly all cloud server systems today. It is used in embedded applications and in military applications. Most critical applications *do use* error detection and in some cases error correction techniques on the DDR3 memory. However these techniques are not a 100% guarantee for error free operation. In the case of multiple bit errors on a single transfer many error detection techniques fall short. Our dependence on DDR3 memory and this known failure mechanism should be a wake up call for the industry. So far the workaround for this is to double the refresh rate to the memory. This is an attempt to 'charge up' the dormant memory cells so that they do not fall victim to adjacent rows that might become 'hammered'. This hits performance and increases power consumption and the problem is not going away. This workaround just reduces the statistical probability.

Why does this happen?

The memory controller's job is to read and write information to and from the memory under program control. If the software running does repeated accesses to a single location the memory controller will generate excessive ACTIVATE commands. Currently there is nothing in the memory controller design to prevent this from happening. Software often uses repetitive accesses to check to see if a task has been completed. This is a very common occurrence in software architecture and referred to as a Semaphore. Several tasks or threads will communicate with each other using a shared location in the memory. Thus they all need to repeatedly access these shared locations in order to communicate.

The use of Semaphore's cause software to access a single DDR3 memory location repetitively

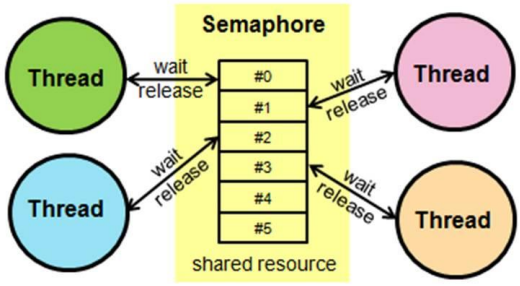

Figure 1 The use of Semaphores can cause failures in DDR3 Memory

Is this problem real?

Although curiously void from the JEDEC[1] meeting minutes we do see evidence of the problem being mentioned in the press and on the internet. Teledyn LeCroy announced an upgrade to their Kibra Memory Analyzer that counts the number of ACTIVATE commands to individual ROW addresses as they appear it the tools trace buffer. FuturePlus Systems has also announced a new feature to its DDR Detective® product to detect the occurrence of Row Hammer using a new technique that employs thousands of counters[2]. Electronic Design has recently mentioned this phenomenon[3] and IBM has a field update to its firmware to try to deal with it.[4] A search of recent patent applications reveals that in January of 2014 Intel submitted two patent applications that deal with Row Hammer. The first is a technique to detect excessive ACTIVATES to a single row address.

Row Hammer Condition Monitoring: US 20140006704 A1. A system monitors data accesses to specific rows of memory to determine if a row hammer condition exists. The system can monitor accessed rows of memory to determine if the number of accesses to any of the rows exceeds a threshold associated with risk of data corruption on a row of memory physically adjacent to the row with high access. Based on the monitoring, a memory controller can determine if the number of accesses to a row exceeds the threshold, and indicate address information for the row whose access count reaches the threshold.

The second Intel patent application deals with a targeted row refresh. That is if the memory controller sees the excessive ACTIVATE commands, below the error threshold, it can tell the DRAM the address of that Hammered Row and the DRAM can refresh and restore the charge to the physically adjacent rows to avoid the problem.

---

[1] www.JEDEC.com is the industry standard group that governs DDR memory.
[2] www.futureplus.com/images/FS2800/Description_of_the_Row_Hammer_feature_on_the_FS2800_DDR_Detective.pdf
[3] http://electronicdesign.com/embedded/achieve-reliability-availability-and-serviceability-memory-interfaces
[4] http://www-947.ibm.com/support/entry/portal/docdisplay?lndocid=migr-5092492

Row hammer refresh command: WO 2014004748 A1. A memory controller issues a targeted refresh command. A specific row of a memory device can be the target of repeated accesses. When the row is accessed repeatedly within a time threshold (also referred to as "hammered" or a "row hammer event"), physically adjacent row (a "victim" row) may experience data corruption. The memory controller receives an indication of a row hammer event, identifies the row associated with the row hammer event, and sends one or more commands to the memory device to cause the memory device to perform a targeted refresh that will refresh the victim row.

Samsung has divulged the issue in a recent investors presentation touting that its DDR4 "in-DRAM" solution is "most efficient for Row Hammer operation".[5]

So clearly the problem is real but clearly not well known.

What can be done?

The problem with all the 'solutions' that are being proposed is that they require significant changes if not replacement of the DRAM and/or memory controller. The only partial solution that can be implemented today is the doubling of the Refresh rate, and as mentioned earlier, this is a power and performance hit and only reduces the statistical probability of the failure. Millions of DDR3 based systems are in service today and these systems run the risk of these failures and all of this is unbeknownst to the users of these systems.

The problem has been showing up in some Data Centers. A few have resorted to clearing memory then prior to allocating a page of memory the memory is read and checked for all 0's. If any bits are 1's those pages are retired and not used.

Error correction techniques built into the DDR3 standard such as ECC are expensive to implement, add additional latency to every Read and Write transaction and will only correct a single bit error and only detect a double bit error. Anything beyond two bits of error in a 64 bit transaction will go undetected. Thus in many ways ECC is a false sense of security if users feel that this will save them.

The best defense for this problem is to test to see if your design or data center is experiencing these excessive ACTIVATE commands to a single row addresses.

Testing for the Row Hammer failure mode

Currently the best method of detecting excessive ACTIVATE commands to a single row address is available from FuturePlus Systems. This solution uses a DIMM interposer that snoops the bus traffic between the memory controller and the memory DIMM. This technique requires no special software to be run on the server or motherboard. Thus the problem can be detected without adding any other variables. The unit that connects to the DIMM Interposer watches all of the memory transactions to and from the Memory Controller and counts the ACTIVATE commands to unique row addresses on the DIMM being monitored. If a threshold is crossed (usually 100,000) within a certain time period the test

[5] http://aod.teletogether.com/sec/20140519/SAMSUNG_Investors_Forum_2014_session_1.pdf

equipment will indicate this. This equipment is physically small and can easily be deployed in a data center or lab. [6]

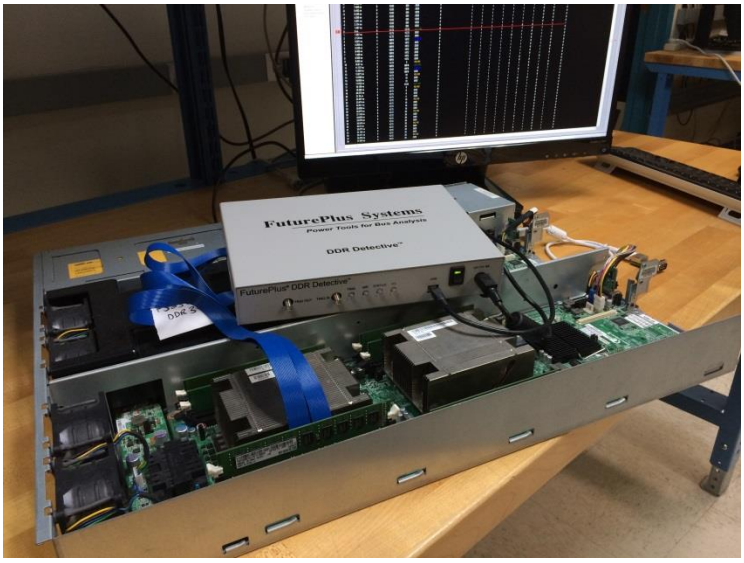

Figure 2 Row Hammer detection equipment installed in an OCP Server

FuturePlus has also produced a short video explaining the row hammer problem.

Summary

This known failure mechanism in DDR3 called *Row Hammer* can lead to undetected data corruption and reliability issues. It has been reported that the main culprit for this phenomenon is when software spins on a semaphore. Critical applications should certainly test for this known failure as the results of undetected data corruption could be catastrophic.

About the Author

Barbara P. Aichinger

Barbara P. Aichinger holds a Bachelors Degree in Electrical Engineering from the University of Akron, Ohio and Masters Degree in Electrical Engineering from the University of Massachusetts. She is a co-founder of FuturePlus Systems and is currently the Vice President of New Business Development. Barbara has almost 30 years of experience with computer architecture and system design that includes both hardware and software. She is widely published in the area of computer test and most recently has authored several presentations on *Memory Errors in the Data Center*. In her spare time Barbara enjoys tennis and is a member of the New Hampshire United States Tennis Association board. She is married and has three children. Her LinkedIn profile can be found at http://www.linkedin.com/pub/barbara-aichinger/3/3a/463/

[6] http://www.futureplus.com/DDR-Detective-Standalone/summary-2800.html

About FuturePlus Systems:
FuturePlus Systems is an innovator in the Test and Measurements industry and has been in business since 1991. The company has a global customer base and has offices in Bedford, New Hampshire and Colorado Springs, Colorado. The *DDR Detective®* is the latest in *never been done before* products that FuturePlus prides themselves on. The company is privately owned and can be found on the web at www.FuturePlus.com.