



Description of the Row Hammer Feature on the FS2800 DDR Detective ®

Detecting a Known Failure mechanism in DDR3 Memory

Row Hammer is a name given to a recently discovered but not well publicized failure mechanism that plagues nearly all cloud computing servers on the planet. In order to detect this obscure but potentially frequent failure a new tool was designed to count occurrences of memory ACTIVATE Commands within a 64ms time period. If this occurs then the system is highly vulnerably to undetected data corruption.

Description of the Row Hammer feature on the FS2800 DDR Detective

Barbara Aichinger FuturePlus Systems Corporation, Bedford, NH

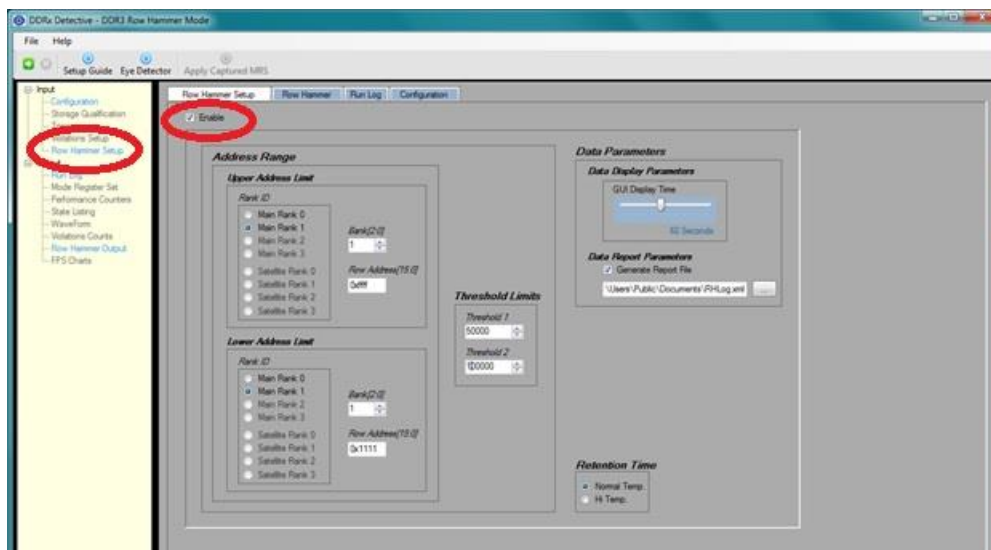
Barb.Aichinger@FuturePlus.com 603-472-5905

Introduction

The FS2800 is a DDR3/DDR4/LPDDR3 protocol analyzer that connects to the Address, Command and Control bus to analyze all aspects of the DDR bus protocol. The onboard FPGA can be reconfigured for multiple uses. Among them the detection of excessive ACTIVATE commands to a single row address for DDR3 which has been shown to create bit flips on physically adjacent rows. This failure mode has been coined 'Row Hammer' since a particular row is being 'hammered' with ACTIVATE commands within the 64 ms retention period.

How to load the Row Hammer feature

After selecting DDR3, configuring the probe and running Eye Detector the user can select *Row Hammer Setup* from the Input/Output pane of the FS2800 Probe Manager. From here the user selects ENABLE as shown below. The FS2800 will now be loaded with a dedicated FPGA image to allow for thousands of counters to help count the number of ACTIVATE commands to unique row addresses.



How the Row Hammer Detection Units work

The DDR Detective utilizes 2400 counters divided into 2, 1200 counter Row Hammer Detection Units to track every ACTIVATE by row address that occur within an interval of 64mS, which relates to the minimum retention period of a single location in a synchronous DRAM. These counters are reset and reassigned at the end of each retention period interval. These counters run continuously and never miss an ACTIVATE command. If the number of unique addresses exceeds the available counters a remainder count is incremented to indicate that the traffic is highly variable. In addition the retention period can be lowered to 32ms by selecting Hi Temp as this is the retention period for high temperature

operation. The retention period of the tool is fixed in that it is not a 'rolling' window. To help compensate for this the tool actually has a duplicate set of counters offset by ½ the retention period to give better coverage.

There are 2 variable threshold limits that the user can set to flag Row Addresses that cross those limits. The output graphic shows the rows that have seen ACTIVATE commands exceed the threshold. Blue for Threshold 1 and red for Threshold 2. Threshold1 and Threshold2 are defaulted to 100K and 300K respectively and can be adjusted. The exact address of the row address that crosses these thresholds is also listed in the right hand side scroll pane.

The Row Hammer setup page allows the user to set an Address Range across slots, Ranks and Banks for which the tool will only count ACTIVATE commands to rows in this range. By narrowing down the range it helps keep the counters from overflowing and can give more accurate results.

An XML based report file for all ACTIVATES that cross the thresholds can be generated.

The tool will give a trigger out 7200 cycles following an ACTIVATE command which reaches Threshold 2. The size of the Threshold counters is 21 bits.

Interesting Row Hammer Metrics

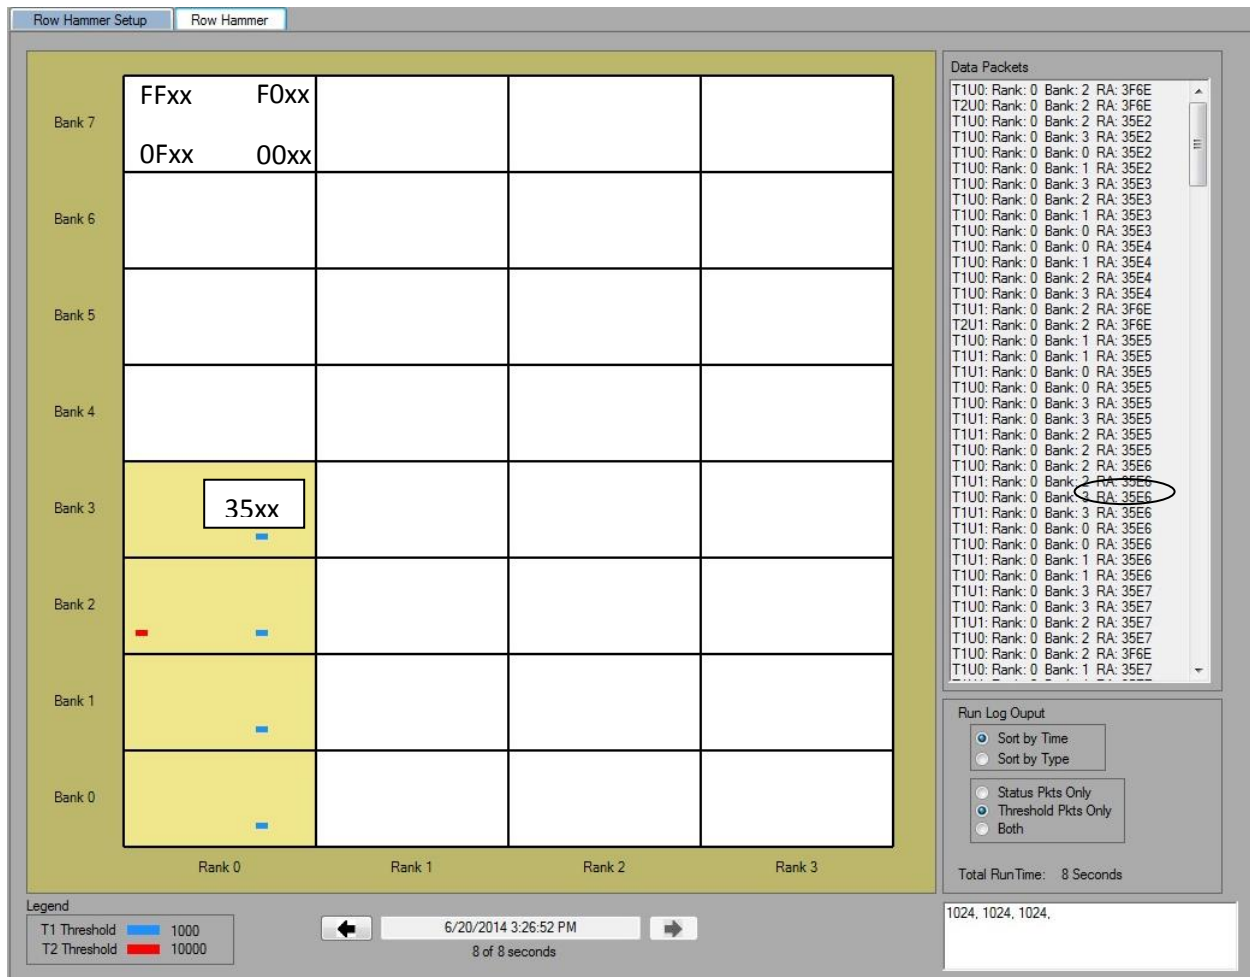
64 mS	is the DDR SDRAM retention rate tRET (below 85 degree C) is the Row Hammer Detect Unit measurement time (standard Temp) is about a 16 th of a second (15.62 intervals per second) is 68 million tCK cycles at DDR3-2133
43 nSec	is the minimum time between DDR3 Activates to the same bank is the sum of DDR3 tRASmin (ACT-> PRECH) and tRP (PRECH-> ACT) it includes time for some reads and writes
1.5 M	is the max number of Activates to a single bank in a 64 mS retention period
12 M	is the max number of Activates to a single DDR3 Rank in a 64 mS retention period
300 K	is a Row Hammer event if a single page gets activated this many times in 64 mS
21 bits	Counts to 2 Million: 2,097,151 and is the size of the counters
3600	is the number of SYSCK cycles of latency for a single RHDU (Row Hammer Detection Unit)
40	is the theoretical max number of 300K-activated pages in a single 64 mS tRET period

The Row Hammer Output Screen

The Row Hammer output graphic shows a mapping by Bank and Rank of each Row Address location that had a total number of ACTIVATES occur over the defined Thresholds (T1 & T2). This display is updated every second.

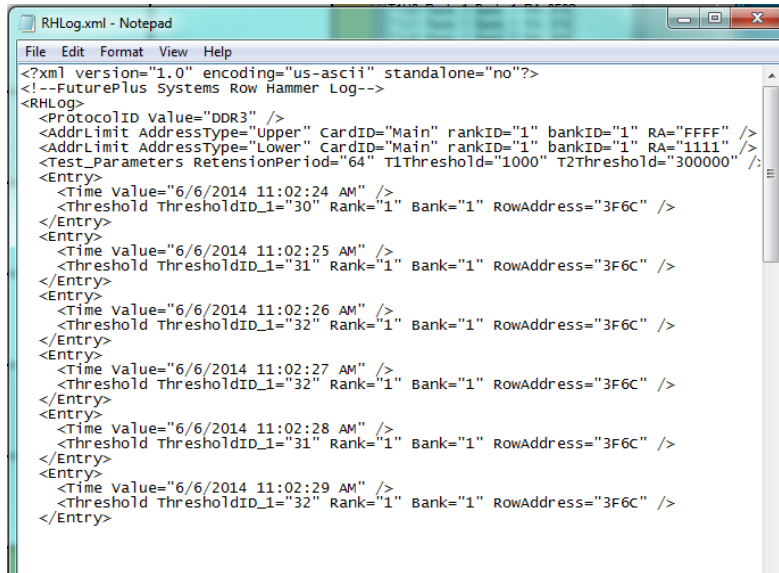
The Row Hammer output graphic shows Ranks in columns and Banks in Rows, which form cells. The specific Row Address is represented in that cell area by the first 2 nibbles (last 2 are don't cares) of the Address as shown below. If the exact Address is needed it can be seen in the report window on the right, or from a generated report.

The output graphic is refreshed every second. After the tool is stopped each previous or successive 1 second period can be paged through using the arrow buttons below the output graphic. A maximum of a 120 seconds worth of traffic can be stored.



Each Row Hammer Detection Unit also reports status at the end of each retention cycle. This status would indicate how many ACTIVATES were not counted (remainder count) and if an overflow did not occur how many counters remained unused.

An xml output report, if checked in the setup screen, can also be generated that writes out to a file all of the row addresses that exceed the threshold for each Row Hammer detection unit.



```
<?xml version="1.0" encoding="us-ascii" standalone="no"?>
<!--FuturePlus Systems Row Hammer Log-->
<RHLog>
  <ProtocolID value="DDR3" />
  <AddrLimit AddrType="Upper" CardID="Main" rankID="1" bankID="1" RA="FFFF" />
  <AddrLimit AddrType="Lower" CardID="Main" rankID="1" bankID="1" RA="1111" />
  <Test_Parameters RetentionPeriod="64" T1Threshold="1000" T2Threshold="300000" />
  <Entry>
    <Time value="6/6/2014 11:02:24 AM" />
    <Threshold ThresholdID_1="30" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
  <Entry>
    <Time value="6/6/2014 11:02:25 AM" />
    <Threshold ThresholdID_1="31" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
  <Entry>
    <Time value="6/6/2014 11:02:26 AM" />
    <Threshold ThresholdID_1="32" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
  <Entry>
    <Time value="6/6/2014 11:02:27 AM" />
    <Threshold ThresholdID_1="32" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
  <Entry>
    <Time value="6/6/2014 11:02:28 AM" />
    <Threshold ThresholdID_1="31" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
  <Entry>
    <Time value="6/6/2014 11:02:29 AM" />
    <Threshold ThresholdID_1="32" Rank="1" Bank="1" RowAddress="3F6C" />
  </Entry>
</RHLog>
```

Summary

This known failure mechanism in DDR3 called *Row Hammer* can lead to undetected data corruption and certainly reliability issues. It has been reported that the main culprit for this phenomenon is when software spins on a semaphore. The Row Hammer Detection feature of the FS2800 can help identify if this problem is occurring in your data center. Critical applications should certainly test for this known failure as the results of undetected data corruption could be catastrophic.