JEDEC PUBLICATION

NEAR-TERM DRAM LEVEL ROWHAMMER MITIGATION

JEP300-1

MARCH 2021

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION

NOTICE

JEDEC standards and publications contain material that has been prepared, reviewed, and approved through the JEDEC Board of Directors level and subsequently reviewed and approved by the JEDEC legal counsel.

JEDEC standards and publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for use by those other than JEDEC members, whether the standard is to be used either domestically or internationally.

JEDEC standards and publications are adopted without regard to whether or not their adoption may involve patents or articles, materials, or processes. By such action JEDEC does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the JEDEC standards or publications.

The information included in JEDEC standards and publications represents a sound approach to product specification and application, principally from the solid state device manufacturer viewpoint. Within the JEDEC organization there are procedures whereby a JEDEC standard or publication may be further processed and ultimately become an ANSI standard.

No claims to be in conformance with this standard may be made unless all requirements stated in the standard are met.

Inquiries, comments, and suggestions relative to the content of this JEDEC standard or publication should be addressed to JEDEC at the address below, or refer to www.jedec.org under Standards and Documents for alternative contact information.

Published by
©JEDEC Solid State Technology Association 2021
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2108

PRICE: Contact JEDEC

Printed in the U.S.A.

NEAR-TERM DRAM LEVEL ROWHAMMER MITIGATION

Contents

Introduction

RAM process node transistor scaling for power and DRAM capacity has made DRAM cells more sensitive to disturbances or transient faults. This sensitivity becomes much worse if external stresses are applied in a meticulously manipulated sequence, such as Rowhammer. Rowhammer related papers have been written outside of JEDEC, but some assumptions used in those papers didn’t explain the problem very clearly or correctly, so the perception for this matter is not precisely understood within the industry. This publication defines the problem and recommends following mitigations to address such concerns across the DRAM industry or academia.

NEAR-TERM DRAM LEVEL ROWHAMMER MITIGATION

(From JEDEC Board Ballot JCB-21-05, formulated under the cognizance of the JC-42 Committee on Solid State Memories. Item 1866.01.)

1 Scope

Dynamic Random Access Memory (DRAM) vendors mitigate Rowhammer bit flipping through a combination of design, architecture and process choices to ensure that Rowhammer bit flips do not occur under ordinary workload usage conditions. On-DRAM mitigation, coupled with proper external support, can effectively protect the data integrity of the DRAM during targeted and/or malicious exploit attacks.

This publication recommends the use of Refresh Management (RFM) as a first option for DRAM data integrity protection.

If the DDR product does not support RFM, Refresh commands issued at regular intervals to satisfy tREFI is recommended. Postponing or/and Pulling-in the Refresh commands provides flexibility to systems' command scheduling for certain DDR family product. However, maintaining regular tREFI enables the DRAM to improve data preservation. Issuing Refresh commands at increased refresh rates can improve DRAM data protection for DRAM products without RFM.

2 Terms and Definitions

Aggressor Row(s): A row or rows that receive excessive activations during a Rowhammer attack in an attempt to disturb data in nearby DRAM cells.

RFM: Refresh Management supported by the DRAM.

Refresh pull-in and postpone: To allow for improved efficiency in scheduling, DRAM Refresh command can pull-in or postpone the refresh cycles.

Victim Row(s): A row or rows with cells that are affected by Rowhammer activity.

3 Problem Statement for Rowhammer

Memory access patterns that repeatedly access one or more rows (aggressor rows) may result in bit flips to the neighboring rows (victim rows). A Rowhammer attack pattern is one example of repeated access pattern. A Rowhammer attack may take advantage of unintended or undesirable bit flips, so the DRAM industry has taken steps to prevent these potential system security exploits through DRAM's design innovation. However, recently published papers have indicated that some DRAM design innovation may be insufficient for the DRAM to adequately protect against all Rowhammer attacks. This publication addresses such concerns in the near-term, while the whole industry (DRAM, system, software) works to address the Rowhammer concerns in the long-term.

4 DRAM Rowhammer Mitigation

A Rowhammer attack consists of repeated accesses to DRAM rows or row regions, and the industry has taken steps to prevent this type of exploit through DRAM design, architecture and process innovations, as well as by CPU and/or operating system's Rowhammer protection algorithms.

This publication highlights DRAM-level mitigation resolution to prevent from Rowhammer attack for legacy DRAMs, such as DDR4 and LPDDR3, and contemporary DRAMs, such as DDR5, LPDDR4/4x and LPDDR5. Refresh management (RFM), if applicable, will be a first option and refresh cycle optimization will also be another option. The role of the DRAM Rowhammer protection is to provide additional safety, but it cannot eliminate risk from all possible forms of attacks.

4.1 Refresh Management (RFM)

DDR, LPDDR and HBM product families adopted RFM on DDR5, LPDDR4/4X, LPDDR5, HBM3 and GDDR6 to support RFM. RFM's goal is to be an effective measure to handle Rowhammer attacks from the DRAM-level. Periods of high DRAM activity may require additional refresh commands to protect the integrity of the DRAM data. The DRAM will indicate the requirement, i.e., Rolling Accumulated Activate Initial Management Threshold (RAAIMT) and Rolling Accumulated Activate Maximum Management Threshold (RAAMMT), for additional RFM commands by setting the associated read only Mode Register (MR).

The expected DDR5 controller implementation of RFM monitors ACT commands issued per bank to the DRAM (refer to specific datasheets to confirm expected implementation for other product families). This activity can be monitored as a Rolling Accumulated ACT (RAA) count. Each ACT command will increment the RAA count by 1 for the individual bank receiving the ACT command.

When the RAA counter reaches a DRAM vendor specified RAAIMT, executing the Refresh Management (RFM) command allows additional time for the DRAM to manage refresh internally. For DDR5, the RFM operation can be initiated to all banks on the DRAM with the RFMab command, or to a single bank address in all bank groups with the RFMsb command. Other product families may support per bank operation with the RFMpb command.

4.2 Refresh Cycle Control (Devices w/o RFM support)

A Refresh command is required to be issued to a DRAM on average tREFI. To allow for improved efficiency in scheduling and switching between tasks, some flexibility in the absolute refresh interval is provided for postponing and pulling-in refresh commands. To maintain higher data integrity with malicious patterns like Rowhammer, it is recommended to issue Refresh command at the average refresh interval or use an increased refresh rate (e.g., double).

Elimination of pulled-in and postponed Refresh commands allows the DRAM the chance to consistently align internal address sampling and perform Rowhammer mitigation at consistent intervals. Additionally, some randomization when issuing of Refresh commands within the interval can also assist Rowhammer mitigation.

4.2 Refresh Cycle Control (Devices w/o RFM support) (cont'd)

Setting the DRAM to double the refresh rate (e.g., 2x for DDR4, 0.5x for LPDDR3) allows the DRAM to perform Rowhammer mitigation twice as often as the 1x mode. Additionally, doubling the refresh rate will refresh all rows twice as often as 1x mode. Both changes positively impact Rowhammer mitigation on the DRAM.

The following sections review existing refresh specifications.

4.2.1 DDR4 Refresh Specifications

A maximum of 8 Refresh commands can be postponed when DRAM is in 1X refresh mode and for 2X or 4X refresh mode, 16 or 32 Refresh commands can be postponed, respectively, during operation of the DDR4 SDRAM, meaning that at no point in time more than a total of 8,16, or 32 Refresh commands are allowed to be postponed for 1X,2X, or 4X Refresh mode, respectively. In a case where 8 Refresh commands are postponed in a row, the resulting maximum interval between the surrounding Refresh commands is limited to 9 × tREFI.

In 2X and 4X Refresh mode, the resulting maximum interval between the surrounding Refresh commands is limited to 17 x tREFI2 and 33 x tREFI4, respectively. A maximum of 8 additional Refresh commands can be issued in advance ("pulled in") in 1X refresh mode and for 2X or 4X refresh mode, 16 or 32 Refresh commands can be pulled in, respectively, with each one reducing the number of regular Refresh commands required later by one. Note that pulling in more than 8, 16, or 32 Refresh commands in advance, depending on Refresh mode, does not further reduce the number of regular Refresh commands required later, so that the resulting maximum interval between two surrounding Refresh commands is limited to 9 × tREFI , 17 x tREFI2 and 33 x tREFI4 respectively.

At any given time, a maximum of 16 REF (32REF in 2X, 64REF in 4X) commands can be issued within 2 x tREFI (4 x tREFI2 in 2X, 8 x tREFI4 in 4X).

4.2.2 LPDDR3 Refresh Specifications

An all bank refresh command needs to be issued to the LPDDR3 SDRAM regularly every tREFI (or more precisely tREFIM = tREFI x RM) interval. To allow for improved efficiency in scheduling and switching between tasks, some flexibility in the absolute refresh interval is provided for postponing and pulling-in refresh command. To maintain higher data integrity with malicious pattern like Rowhammer, it is recommended to apply the regular refresh interval or double the refresh rate.

A maximum of 8 Refresh commands can be postponed during operation of the LPDDR3 SDRAM, meaning that at no point in time more than a total of 8 Refresh commands are allowed to be postponed. In case that 8 Refresh commands are postponed in a row, the resulting maximum interval between the surrounding Refresh commands is limited to 9 × tREFI (9 x tREFIM = 9 x RM x tREFI) (see Figure 1).

A maximum of 8 additional Refresh commands can be issued in advance ("pulled in") , with each one reducing the number of regular Refresh commands required later by one. Note that pulling in more than 8, depending on Refresh mode, Refresh commands in advance does not further reduce the number of regular Refresh commands required later, so that the resulting maximum interval between two surrounding Refresh commands is limited to 9 × tREFI (9 x tREFIM = 9 x RM x tREFI).

4.2.2 LPDDR3 Refresh Specifications (cont'd)

At any given time, a maximum of 16 REF commands can be issued within 2 x tREFI (2 x tREFIM = 2 x RM x tREFI)

4.2.3 GDDR6 Refresh Specifications

GDDR device requires REFRESH cycles at an average periodic interval of tREFI. To allow for improved efficiency in scheduling and switching between tasks, some flexibility in the absolute refresh interval is provided for postponing and pulling-in refresh command. To maintain higher data integrity with a malicious pattern like Rowhammer, it is recommended to apply the regular refresh interval.

A maximum of 8 REFab commands can be postponed during operation of the device; at no point in time more than a total of 8 REFab commands are allowed to be postponed. In case that 8 REFab commands are postponed in a row, the resulting maximum interval between the surrounding REFab commands is limited to 9 × tREFI (see figures below).

A maximum of 8 additional REFab commands can be issued in advance ("pulled in"), with each one reducing the number of regular REFab commands required later by one. Note that pulling in more than 8 REFab commands in advance does not further reduce the number of regular REFab commands required later, so that the resulting maximum interval between two surrounding REFab commands is limited to 9 × tREFI (see figures below). At any given time, a maximum of 9 REFab commands can be issued within tREFI

4.2.4 HBM2 Refresh Specifications

HBM devices require REFRESH cycles at an average periodic interval of tREFI. To allow for improved efficiency in scheduling and switching between tasks, some flexibility in the absolute refresh interval is provided. To maintain higher data integrity with a malicious pattern like Rowhammer, it is recommended to apply the regular refresh interval or double the refresh rate.

A maximum of 8 REFRESH commands can be postponed during operation of the HBM device; at no point in time more than a total of 8 REFRESH commands are allowed to be postponed. In case that 8 REFRESH commands are postponed in a row, the resulting maximum interval between the surrounding REFRESH commands is limited to 9 × tREFI.

A maximum of 8 additional REFRESH commands can be issued in advance ("pulled in"), with each one reducing the number of regular REFRESH commands required later by one. Note that pulling in more than 8 REFRESH commands in advance does not further reduce the number of regular REFRESH commands required later, so that the resulting maximum interval between two surrounding REFRESH commands is limited to 9 × tREFI.

At any given time, a maximum of 16 REFRESH commands can be issued within 2 × tREFI.

4.2.5 DDR4, LPDDR3, GDDR6 and HBM2 Refresh Command Postpone/Pull-in Timing Diagram

The Refresh Command sections of the DDR4 SDRAM Standard (JESD79-4), LPDDR3 SDRAM Standard (JESD209-3), Graphic Double Data Rate 6 SGRAM Standard (JESD250), and HBM DRAM Standard (JESD235), include timing diagrams depicting how postpone or pull-in of the refresh command(s) can take place. Sample diagrams are included in Figure 1 for reference.

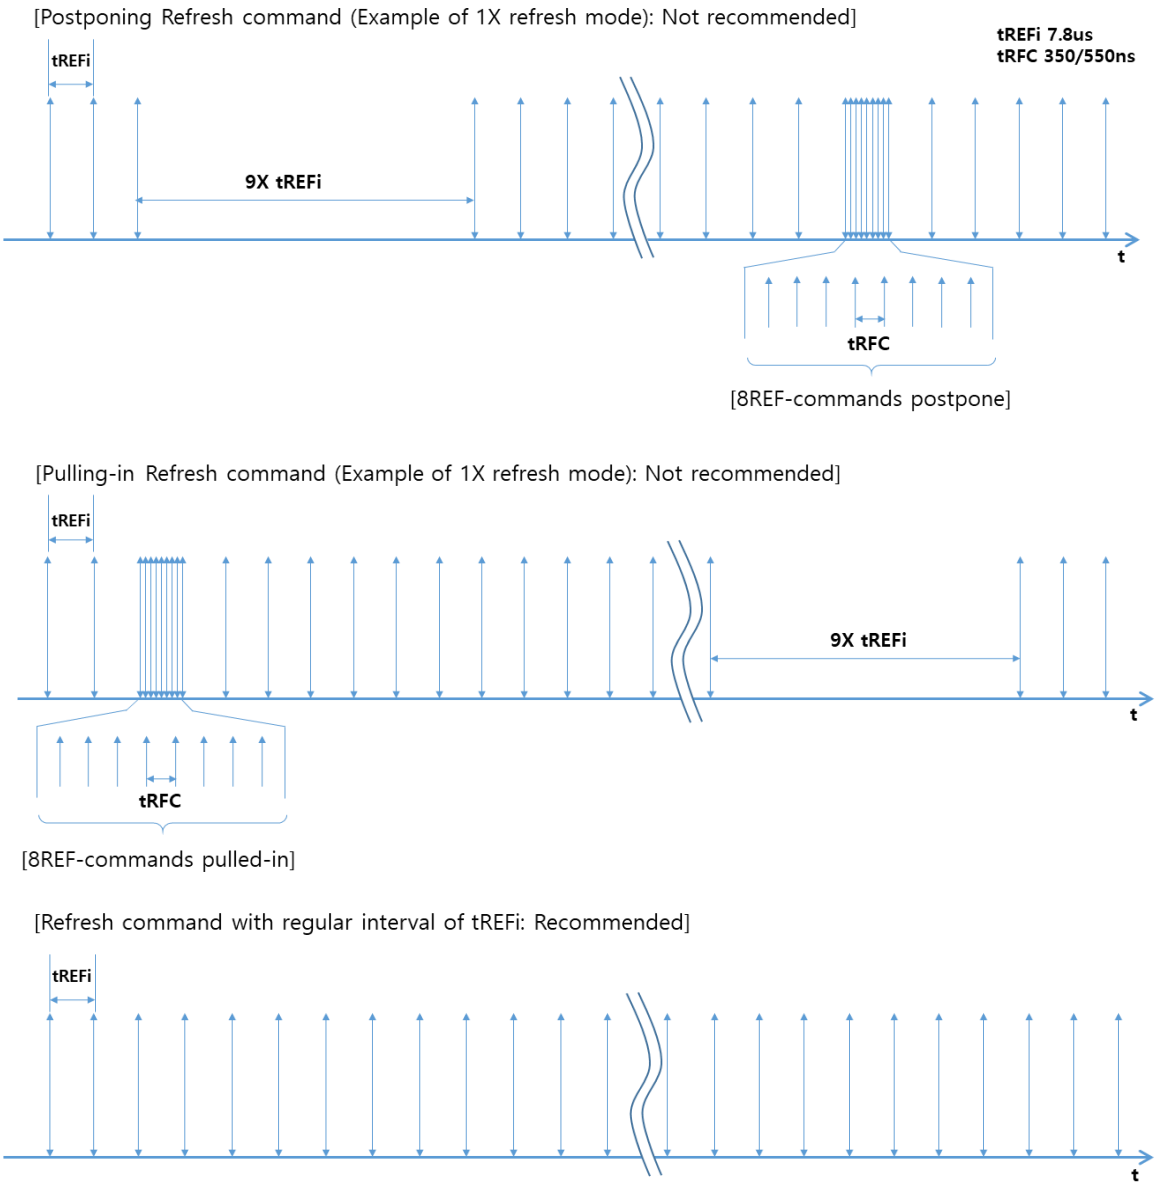

Figure 1 — Sample timing diagrams

Standard Improvement Form JEDEC _____

The purpose of this form is to provide the Technical Committees of JEDEC with input from the industry regarding usage of the subject standard. Individuals or companies are invited to submit comments to JEDEC. All comments will be collected and dispersed to the appropriate committee(s).

If you can provide input, please complete this form and return to:

JEDEC
Attn: Publications Department
3103 North 10th Street
Suite 240 South
Arlington, VA 22201-2107

Fax: 703.907.7583

1. I recommend changes to the following:

☐ Requirement, clause number _____

☐ Test method number _____ Clause number _____

The referenced clause number has proven to be:

☐ Unclear ☐ Too Rigid ☐ In Error

☐ Other _____

2. Recommendations for correction:

_____

_____

_____

_____

3. Other suggestions for document improvement:

_____

_____

_____

_____

Submitted by

Name: _____ Phone: _____

Company: _____ E-mail: _____

Address: _____

City/State/Zip: _____ Date: _____


JEDEC®