Technische Universität
Braunschweig

INSTITUT FÜR DATENTECHNIK
UND KOMMUNIKATIONSNETZE

TN-IDA-RAD-12/13

Heavy Ion SEE Test of 2-Gbit and 4-Gbit DDR3 SDRAM Devices
Test report

M. Herrmann, K. Grürmann, F. Gliem

ESTEC Contract No. 4000101358/10/NL/AF

Radiation hard memory, Radiation testing of candidate memory devices for Laplace mission

Technical officer: V. Ferlet-Cavrois

Final issue: November 24, 2014

1 Abstract

From January 9 to January 13, 2012, and from April 16 to April 20, 2012, we performed heavy ion test campaigns with DDR3 SDRAM and NAND Flash devices at RADEF, Jyväskylä, Finland. This document reports on the findings for DDR3 SDRAM devices.

2 Test coverage

2.1 General remarks

SDRAM is designed for use in ground mass memories and is also used in space mass memories, where they compete with single level cell (SLC) NAND-Flash. The usual Reed-Solomon error correction (RS-ECC) is very effective against sparsely distributed single event upsets (SEU), which therefore are of minor concern.

Therefore, the data integrity of SDRAM and NAND-Flash based space mass memories is determined mainly by radiation induced breakdown of the device functionality, in case of SDRAM several types of SEFI, in particular device SEFI, and in case of NAND-Flash several types of SEFI including device SEFI and in particular destructive failures (DF).

Device SEFIs cause a transient or persistent blocking of the device functionality, but no physical destruction. Device SEFIs can be resolved by corrective actions, sometimes by reset measures and always by power cycling.

Power cycling destroys the data contents of SDRAMs, in contrast to Flash.

NAND-Flash is prone to destructive failure, in contrast to SDRAM. The main reason for this is the on-chip high voltage (≈ 20 V) generation and distribution.

Both NAND-Flash and SDRAM combine the storage array with an on-chip processor, which controls various operation sequences. The existence of the on-chip processor gives rise to the implementation of several modes of device operation and of device communication with the memory controller.

Examples for the tested DDR3 SDRAM devices are: auto-refresh vs. self-refresh mode, burst lengths, DLL on or off, and various timing parameters.

In consequence, it is not possible to achieve full test coverage over all operational mode combinations of the respective device within the given beam time restrictions. Instead, only a subset of operational conditions can be tested, which is believed to be more or less representative for future device applications.

The synchronous high data rate communication with the memory controller implies a tight coupling between the SDRAM device and its external controller. Details of the interaction between device and controller can cause significant modifications of the error image.

If, such as in this campaign, a survey of the devices on the market is intended, the test coverage necessarily is even more restricted. What the survey can deliver is an assessment of the general usability of these devices in the space radiation environment, in particular in comparison to the competing NAND-Flash.

The gained test data show that state of the art commercial DDR3 SDRAM are strong candidates for the implementation of space mass memories. The parts of the different manufacturers are fabricated according to a common JEDEC specification of their functionality to be interchangeable for ground applications. Nevertheless we see differences in their SEE response, which reflect differences in the chip design.

In view of any real application we do strongly recommend to test the device chosen for that application again under the exact conditions of the respective application, in particular with respect to its specific operational mode and in combination with the intended memory controller.

2.2 Specific remarks

As already mentioned, the SEE performance of the SDRAM devices is mainly determined by their sensitivity to SEFIs, especially to device SEFIs. Proper device operation depends on the writable contents of several on-chip control registers. Corruption of the regular contents causes severe disruptions of the device functionality, which appear as SEFI in the retrieved data. A recommended countermeasure is *software conditioning*, namely a periodic rewrite of the control registers. In this test, the rewrite was performed every row for the read and write/read mode tests, or every second for the storage mode tests. By comparison with test runs without software conditioning, we proved the effectiveness also in our case.

Accordingly, we tried to get test runs with software conditioning over the full LET scale, because we assume that the advantage by software conditioning will be exploited in space applications. As expected, software conditioning only has a minor influence on the SEU cross section. But it provides a significant improvement of the SEFI cross sections.

Main attention was given to the 4-Gbit devices of Elpida and Samsung, and also to the 2-Gbit devices of Samsung, revision D, and Micron. Unfortunately, Elpida went out of business. Samsung is the major player. Micron is another important supplier using a different chip design. Additional tests were performed for the 2-Gbit Hynix and the 2-Gbit Nanya device.

Removal of device SEFIs by power cycling destroys the stored data. Therefore, it is of interest to circumvent this “ultima ratio” by less intrusive reset measures. Respective tests were performed and show substantial differences between the investigated DUT types; however, the results have poor statistics due to the low number of observed SEFI events.

3 Test setup

3.1 DUTs

We tested 7 devices from 5 manufacturers, as shown in table 1. All tested devices had been thinned from an initial thickness of approximately 300 µm to 60 µm. For the Samsung 2-Gbit parts, two die revisions were tested: die revision D is produced in 35-nm technology, while die revision B is produced in the previous generation technology (≈50 nm).

Table 1: tested DUTs

Manufacturer	part number	Capacity	Lot code	Date code	Samples	Notes
Elpida	J4208BASE-DJ-F	4 Gbit	0WPEY00	1044	3	
Samsung	K4B**4**G0846B-HCH9	4 Gbit	GMD02390	1125	3	
Samsung	K4B**2**G0846**B**-HCH9	2 Gbit	GLJ423AC	0949	2	
Samsung	K4B**2**G0846**D**-HCH9	2 Gbit	GEB701GES	1113	8	
Hynix	H5TQ2G83BFR-H9CR	2 Gbit	DTK89426H3	1046	3	
Micron	MT41J256M8HX-15E:D	2 Gbit	BY8F416.21 BYFCM91.21	1116 1006	1 4	
Nanya	NT5CB256M8BN-CG	2 Gbit	01340100GP	1026	1	

In the January test, all of the parts in table 1 were tested. In the April test, only 4-Gbit Samsung, 2-Gbit Samsung (revision D) and 2-Gbit Micron parts were tested.

Table 2 shows the die area per bit and a photo for the different parts that were tested.

Table 2: die area per bit for different DDR3 parts

Manufacturer	Capacity	Comments	Die area per bit [cm²]	Photo
Elpida	4 Gbit		$2.2 \cdot 10^{-10}$	
Samsung	4 Gbit		$1.6 \cdot 10^{-10}$	Not available[1]
Samsung	2 Gbit	Revision B	$3.9 \cdot 10^{-10}$	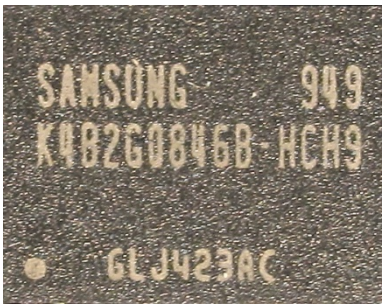
Samsung	2 Gbit	Revision D	$1.6 \cdot 10^{-10}$	
Hynix	2 Gbit		$2.8 \cdot 10^{-10}$	

Micron	2-Gbit	$3.5 \cdot 10^{-10}$	  Or similar[2]
Nanya	2-Gbit	$3.4 \cdot 10^{-10}$	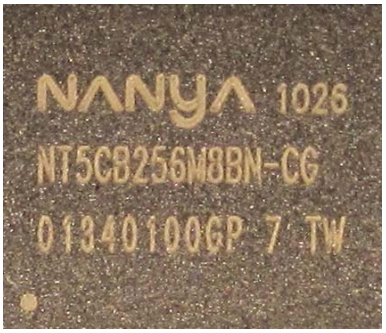

[1] No photos are available of this part because all available samples with the same lot code have been thinned.

[2] The 4-character text on the left side of the package may be different. Details are not available because the samples have been thinned.

3.1.1 DUT preparation

DDR3 silicon dies are encapsulated in flip-chip plastic ball grid array (BGA) packages in order to reduce the lead inductance (figure 1). Before irradiation, and because of the limited range of the used heavy ions, the plastic has to be removed down to the back surface of the die, and the die may have to be thinned, depending on the available ion range.

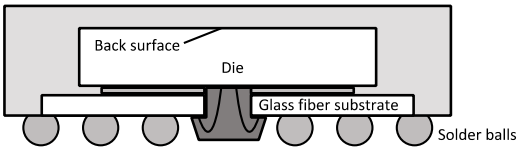

Figure 1: exemplary cross section of a DDR3 device

The thinning of the devices was performed by Institut für Werkzeugmaschinen und Fertigungstechnik (IWF), TU Braunschweig, for some of the devices, and Fraunhofer-Institut für Angewandte Optik und Feinmechanik (IOF) for the rest of the devices. The IWF process is described as follows:

> *[The] plastic above the back surface of the die and a large part of the plastic rim around the edges of the die are removed by wet chemical etching. Thereafter the die is found to be not exactly flat but warped. The surface height varies by up to 40 µm. The*

height contour is measured, and the rotating blade tool used for grinding follows the measured contour.

The IOF thinning process is described by IOF as follows:

[The method] enforces the die in an exactly flat position. This supports grinding and polishing without the need for contour following of the abrasive tool. At first, the layer thickness of the plastic encapsulation and the silicon die to be grinded is determined, based on the cross section of a sample chip. The devices are bilaterally puttied between glass plates. Planarity and parallelism of the chip surface during the mechanical processing is guaranteed by fixing the devices between the glass plates. The deployed machine used (DISCO DAG 810) enables a thickness tolerance of 5 µm.

In the first grinding work step, plastic layer and glass plate were grinded down to the silicon surface of the DDD3 chip. After reaching the silicon surface, the target thickness of the silicon die is entered to the machine and, after that, the silicon die as well as the plastic encapsulation are thinned. Inevitably, the grain size of the tool must be carefully adapted for the silicon material. If the grain size is too rough, microscratches due to the high mechanical stress lead to imperfections in the functionality up to the destruction of the DDR3 chip.

Figure 2 shows a mechanical cross section of a DDR3 device.

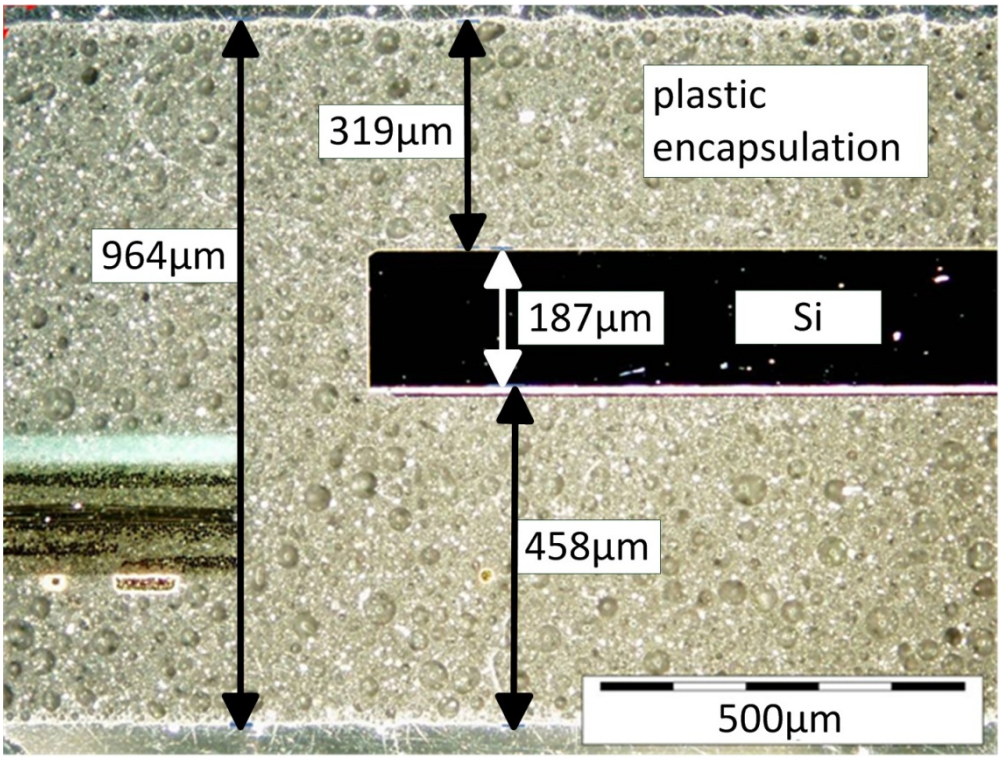

Figure 2: mechanical cross section of a DDR3 device (IOF/IDA)

3.2 Test facility

The tests were performed at the Radiation Effects Facility (RADEF) in the Accelerator Laboratory of the University of Jyväskylä, Finland. The used 9.3 MeV/amu ion cocktail is described in table 3. The depth of 55 µm assumes a die thickness of 60 µm and a depth of the active layer of 5 µm beneath the back surface of the thinned die.

Table 3: RADEF 9.3 MeV/amu ion cocktail

Ion	Energy [MeV]	Range [µm]	LET [MeV cm^2 mg^{-1}]		
			Bragg peak	Surface	55µm
$^{15}N^{4+}$	139	202	5.9	1.83	2.16
$^{20}Ne^{6+}$	186	146	9.0	3.63	4.50
$^{30}Si^{8+}$	278	130	14.0	6.40	*not used*
$^{40}Ar^{12+}$	372	118	19.6	10.2	12.94
$^{56}Fe^{15+}$	523	97	29.3	18.5	25.26
$^{82}Kr^{22+}$	768	94	41.0	32.2	39.83
$^{131}Xe^{35+}$	1217	89	69.2	60.0	68.69

3.3 Test bench

The test bench, RTMC6 (figures 3 and 4), is capable of operating one DDR3 DUT in x4 or x8 configuration at a clock frequency of up to 400 MHz. It is based on a Xilinx ML605 evaluation board, which contains a Xilinx Virtex6 FPGA.

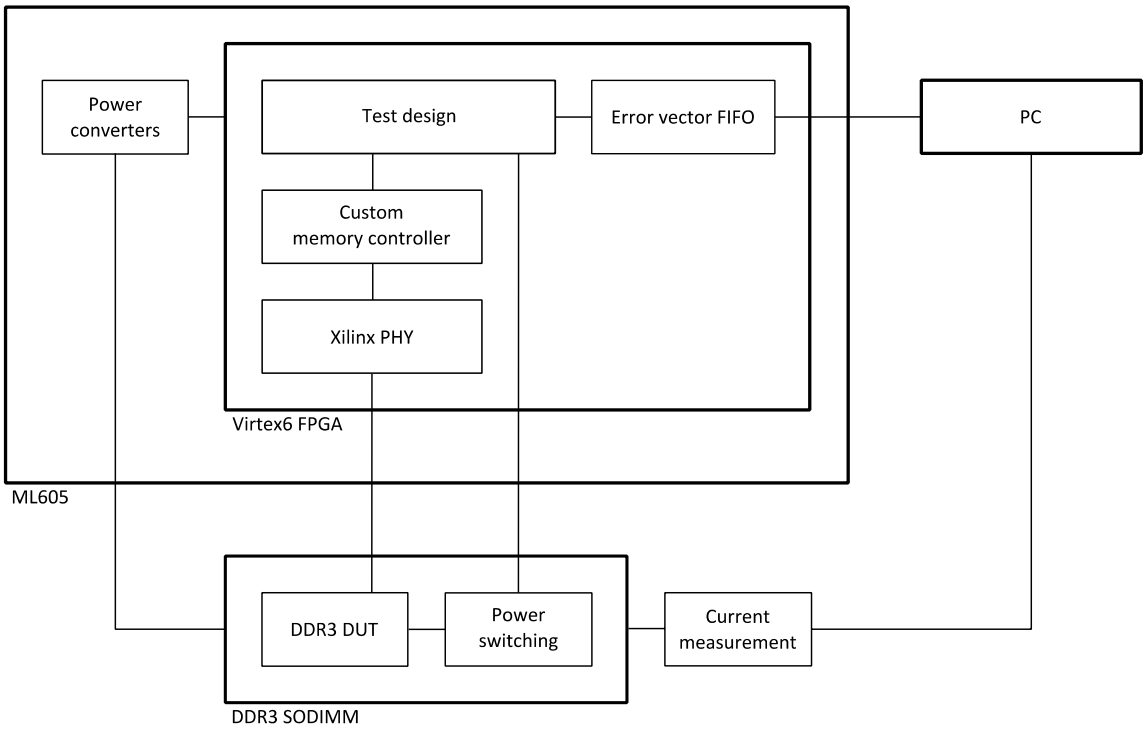

Figure 3: an overview of the RTMC6 test bench (simplified)

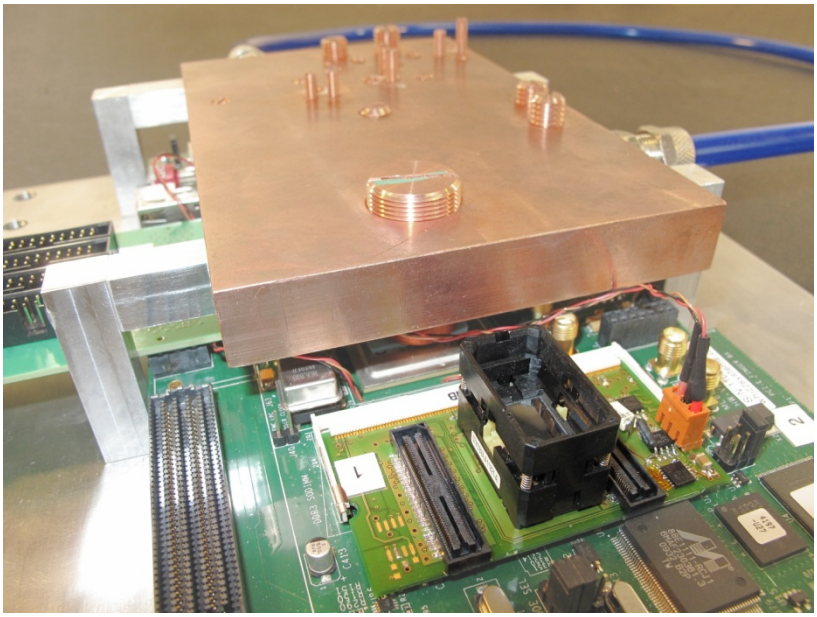

Figure 4: the RTMC6 head station with an opened DDR3 device in the ZIF socket

The ML605 is equipped with a small outline dual inline memory module (SODIMM) socket designed for commercially available memory modules. Since thinned devices cannot be soldered, we developed a DUT adapter which connects to the SODIMM slot and contains a zero insertion force (ZIF) FBGA socket suitable for DDR3 SDRAM devices. The DUT adapter also contains circuitry for switching power to the DUT.

The FPGA contains a custom test design which writes a constant, counting or pseudo-random pattern to the DUT, reads the data from the DUT and compares it to the original pattern. If the data is different from the pattern, an error vector is generated and transmitted to a PC via a high speed USB connection. Since the DUT has a higher data transfer rate than the USB connection, error vectors have to be buffered in a FIFO in order to be able to handle large runs of consecutive errors without slowing down the test. If the error record FIFO overflows due to too many errors, error vectors are discarded.

The DUT is controlled by a custom memory controller. This memory controller provides fine-grained control of the DUT and allows performing operations such as writing the mode registers, resetting the DLL of the DUT or calibrating the termination resistance at arbitrary times. It interfaces with Xilinx' DDR3 PHY (the lowest DUT interface design layer).

On the PC, an error map is displayed for preliminary visual analysis, along with a total error count and various statistics. The error vectors are stored on a hard disk for offline analysis.

3.4 Test sequence

Several test modes are available in order to test the behavior of the device in different situations:

- Storage mode: before irradiation, the pattern is written into the DUT and read in order to determine hard errors. After irradiation, the contents of the device are read and compared to the pattern.
- Read mode: before irradiation, the pattern is written into the DUT and read in order to determine hard errors. During irradiation, the device is continuously read.

- Write/read mode: before irradiation, the pattern is written into the DUT and read in order to determine hard errors. During irradiation, the device is continuously written and read.

3.5 Software conditioning

All test modes can be performed with and without *software conditioning*[1]. Software conditioning consists in a configurable set of operations performed at regular time intervals (storage mode) or after a regular number of read/written pages (other modes). The available operations are:

- Rewriting the mode registers, containing the DUT configuration
- Resetting the DLL internal to the DUT
- Recalibrating the data line termination resistance (ZQ calibration)

3.6 Error classification

Several error classes are distinguished, according to the overview shown in figure 5:

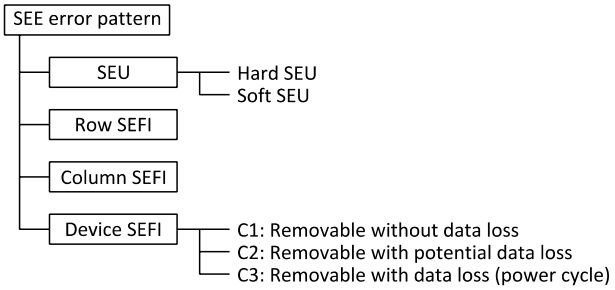

Figure 5: error classification overview

- SEUs are isolated single-bit or multi-bit errors. An SEU is called a *hard SEU* if it cannot be removed by writing the cell again. All other SEUs are called *soft SEUs*. In contrast to soft SEUs, a hard SEU cannot be removed by scrubbing.
- Row SEFIs and column SEFIs are errors that corrupt a single row or column, respectively. The row or column may be corrupted completely or in part.
- Device SEFIs are errors that corrupt a whole device or an extended region of a device. Some device SEFIs can be removed by one of the operations that are also used for software conditioning. These operations do not result in data loss. Other device SEFIs can only be removed by resetting the DUT. The specification does not guaranteed data retention in this case, but in practice, no data loss has been observed. Some device SEFIs can only be removed by power cycling the DUT, which always causes data loss.

An example error map, containing SEUs, row SEFIs and column SEFIs is shown in figure 6.

[1] For historic reasons, software conditioning has sometimes, and inaccurately, been referred to as “reinitialization”.

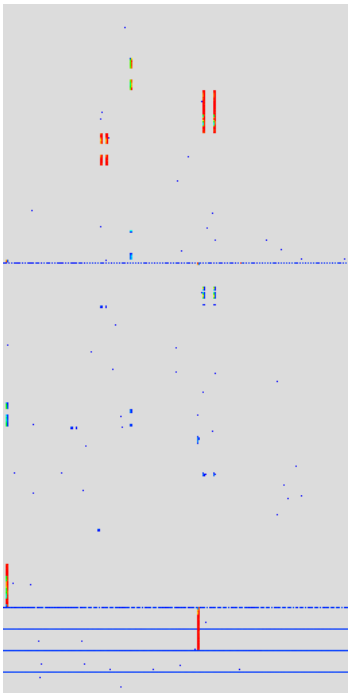

Figure 6: an error map showing row SEFIs, column SEFIs and SEUs

4 Test procedures and test results

In the January test, the influence of software conditioning was examined. In the April test, all tests were performed without software conditioning due to the limited available test time.

A pseudo-random pattern was used for all tests.

4.1 Cross sections

The complete cross section data is hard to visualize, as discussed in appendix A. We therefore present an overview and some conclusions in this section. The complete cross section data for all parts, LETs, test modes, error classes, test campaigns, and with/without software conditioning, is shown in appendix A.

For SEU cross section charts, the approximate die area per bit (see table 2) is also shown.

Not all cross section test data is complete; values may be missing for a variety of reasons:

- Due to test time constraints, it is not possible to exhaustively test every part at every LET and in every test mode, with and without software conditioning. Therefore, some runs had to be skipped in order to get broader results. Preferably, runs were skipped for parts that are considered less relevant, for example because they are outdated.
- For tests in storage mode, if a device SEFI occurs, no SEU, row SEFI, and column SEFI cross sections are available because they are hidden by the device SEFI. Generally, multiple runs were performed in such a case, but sometimes, every run resulted in a device SEFI.
- For some parts tested in 1/2012, only a single measurement was performed in 4/2012 in order to verify the agreement of the measurement. Generally, the results are in good agreement. The only exception is the 4-Gbit Samsung part in write/read mode, where, for unknown reasons, for some of the LETs, the DUT measured in 4/2012 delivered a

significantly higher SEU cross section than the DUT measured in 1/2012. The full results of the comparison are shown in appendix A (section 6.4).

4.1.1 Influence of test mode and overall part comparison

Generally, the results in the different modes are very similar (the full data is shown in appendix A, section 6.1).

We approximate the cross section by a Weibull function

.

The parameters are the *limit cross section* (or *saturation cross section*) , the *threshold LET* , and the parameters. is the cumulative density function of the Weibull distribution:

Figures 7 to 11 show Weibull approximations for the SEU, row SEFI and device SEFI cross sections, respectively, for the worst case test results over all test modes (storage mode, read mode, and write/read mode). The approximations were calculated using OMERE 3.4.0.0. The approximation parameters are listed in appendix C. For the column SEFI cross section (figure 9), the actual worst-case values are shown instead of an approximation, since most part types did not show any column SEFIs at all. For the device SEFI cross section with software conditioning (figure 11), not enough data was available for the 2-Gbit Hynix and Micron parts.

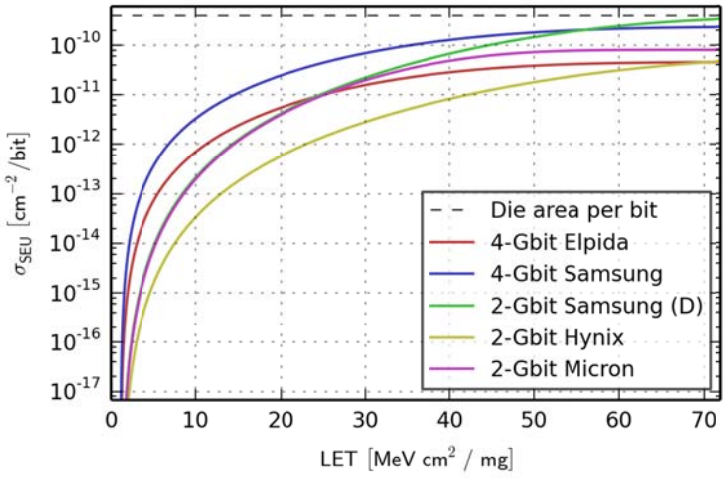

Figure 7: SEU cross section approximation, worst case over all test modes

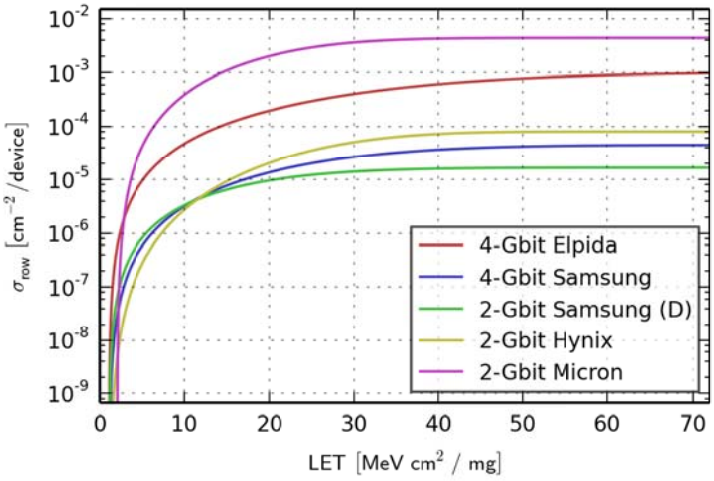

Figure 8: Row SEFI cross section approximation, worst case over all test modes

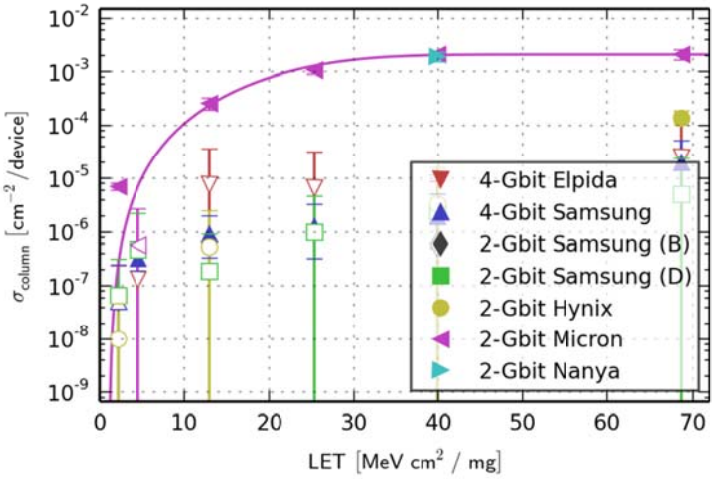

Figure 9: Column SEFI cross section, worst case over all test modes

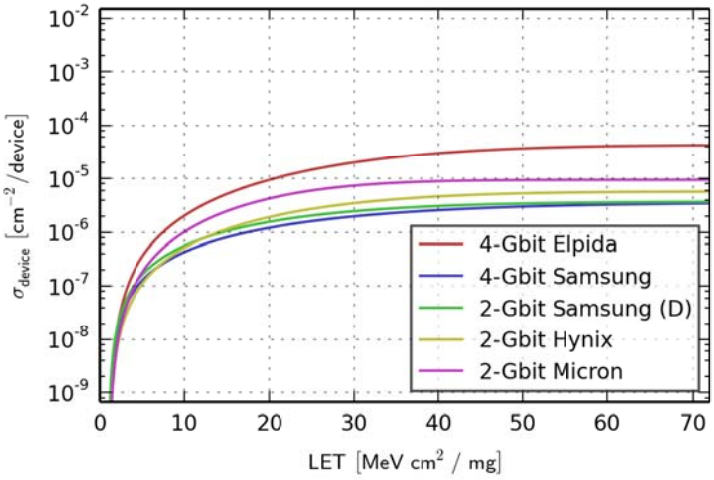

Figure 10: Device SEFI cross section approximation, worst case over all test modes, software conditioning off

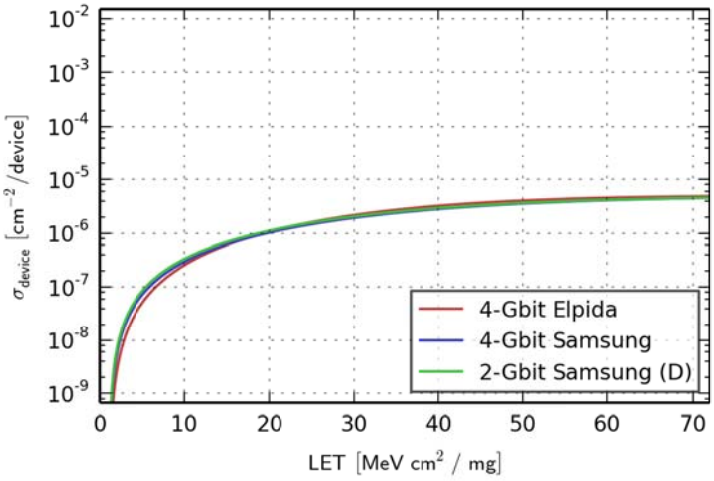

Figure 11: Device SEFI cross section approximation, worst case over all test modes, software conditioning on

Generally, the results for read mode and write/read mode are very similar to those for storage mode. This is plausible because dynamic memories perform a periodic refresh operation (every 7.8 µs in the case of DDR3), which opens and precharges banks, similar to what happens during a write or read operation. Notable exceptions are:

- For the 4-Gbit and 2-Gbit (revision D) Samsung parts, the SEU cross sections are significantly higher in write/read mode than in the other modes (sections 6.1.2 and 6.1.4)
- For the 2-Gbit Micron part, the device SEFI cross sections are slightly lower in storage mode than in the other modes, but mostly within the error bars (section 6.1.27)

For row SEFIs and column SEFIs, the 2-Gbit Micron part has a much higher cross section than the other parts.

Column SEFIs are very rare, being exhibited by only the 2-Gbit Micron and Nanya parts, and occasionally the 4-Gbit Samsung part.

For the 2-Gbit Nanya device, little data is available. For the available data, its sensitivity is similar to that of the 2-Gbit Micron device. Furthermore, the Samsung 4-Gbit and 2-Gbit (revision D) parts are similar in general.

Using the data from figures 7 to 10, we can now rank the tested part types, according to their sensitivity to different error classes. Table 4 shows this ranking for the arbitrarily chosen LET value of . For the 2-Gbit Samsung (revision B) part and the 2-Gbit Nanya part, not enough data is available. It is worth noting that the relative importance of the different error class sensitivities is application-specific. Furthermore, this comparison only includes the SEE performance; depending on the application, the TID performance has to be considered as well, see [1, 2, 3, 4].

Table 4: Ranking of part types according to their sensitivity at an LET of ; 1 is the best and 5 is the worst.

Part type	SEU	Row SEFI	Column SEFI	Device SEFI
4-Gbit Elpida	3	4	1	2
4-Gbit Samsung	4	2	3	1
2-Gbit Samsung (B)	Not enough data			

2-Gbit Samsung (D)	2	1	1	1
2-Gbit Hynix	1	3	2	1
2-Gbit Micron	2	5	4	2
2-Gbit Nanya	Not enough data			

4.1.2 Effects of software conditioning

Software conditioning was performed after every row for read and write/read mode tests, and once per second for storage mode tests.

Software conditioning has no effect on the SEU cross section for any part (see section 6.3).

The row SEFI cross section is slightly decreased by software conditioning for the 2-Gbit (revision D) and 4-Gbit Samsung devices (sections 6.3.11 and 6.3.9). The same is true for the 2-Gbit Micron (section 6.3.13) and Nanya (section 6.3.14) parts, but these devices were only tested at one single LET.

Column SEFIs are completely removed by software conditioning for the 4-Gbit Samsung part (section 6.3.16; column SEFIs were only observed in write/read mode). The 2-Gbit Hynix part only had column SEFIs at the maximum LET, where software conditioning was effective (section 6.3.19). For the 2-Gbit Micron (section 6.3.20) and Nanya (section 6.3.21) parts, the column SEFI cross section was severely reduced, but these devices were only tested at one single LET.

The device SEFI cross section is reduced severely by software conditioning for the 4-Gbit Elpida part (section 6.3.22). For the 4-Gbit Samsung part, the cross section is not significantly reduced (section 6.3.23), but note that the device SEFI cross section for the Samsung part *without* software conditioning is already about the same as for the Elpida part *with* software conditioning. For all other parts, the results are inconclusive.

The full comparison of tests with and without software conditioning is shown in section 6.3.

4.2 SEFI mitigation

Seeing that software conditioning can reduce the device SEFI cross section, the question arises whether software conditioning can also remove a device SEFI condition after it occurred, as opposed to preventing it.

In order to assess this question, we irradiated various parts at different LET in different modes without software conditioning until a device SEFI or a sufficient number of row/column SEFIs had occurred. We then stopped the irradiation and performed different software conditioning measures in varying order, writing and reading the device after every measure, until the SEFI was no longer present.

The measures we used are (cf. section 3.5):

- Rewriting one of the four available mode registers (MR0, MR1, MR2, MR3). The mode register values follow from the mode of operation (timing parameters etc.), which does not change for our controller.
- Performing a long ZQ calibration (ZQCL). Note that a short ZQ calibration is performed every 128 milliseconds at all times by the controller, as suggested by the standard ([5], section 5.5.1).

- Resetting the DLL of the DDR3 DUT. This is achieved by writing bit 8 in MR0. Since it is not possible to access individual mode register bits, this always includes rewriting MR0.
- Resetting the DUT (DR). This is followed by a full initialization sequence ([5], section 3.3.2)
- Resetting the PHY of the memory controller (PR). In addition to a DUT reset and the associated initialization sequence, this causes a recalibration of the line delays.
- Power cycling the DUT (PC) for 100 ms. This also necessitates a PHY reset and DUT reinitialization.

The measures can be classified in three groups, according to their impact on the stored data, as follows:

- Group A: MR0, MR1, MR2, MR3, ZQCL, and DLL reset. These measures may be performed by the controller at arbitrary times (subject to timing constraints) without affecting the data stored in the DUT ([5], section 3.4.1).
- Group B: DUT reset and PHY reset. These measures are not guaranteed to retain data. In practice, however, not data loss was observed. Note that during PHY initialization (performed after PHY reset), the PHY overwrites the first eight memory words with a test pattern (0xFF00AA5555AA9966) for performing line length calibration.
- Group C: power cycle. This measure invariably destroyed all data stored in the DUT.

We typically performed the measures on order from lowest to highest expected probability of success, in order to be able to determine the ineffectiveness of as many measures as possible without prematurely removing the SEFI condition. A typical sequence was therefore: ZQCL, MR3, MR2, MR1, MR0, DLL reset, device reset, PHY reset, power cycle. This order was not strictly adhered to in order to test different measures in different situations. Sometimes, all Group A measures were performed together in a sequence similar to what is performed during device initialization.

It is worth noting that in some cases, different SEFIs seemed to be present in the DUT at the same time. It is possible that one measure removes some of the errors, but not others.

An example is run 1074: a 4-Gbit Elpida DUT was Nitrogen in write/read mode. After a fluence of $1.6 \cdot 10^4$ cm^{-2}, a device SEFI occurred and the irradiation was stopped. ZQ calibration, rewriting mode register 3, rewriting mode register 2 and rewriting mode register 1 did not have any effect (the DUT was written and read after each measure). Rewriting mode register 0 removed the device SEFI, and a row SEFI remained. DLL reset (again writing and reading afterwards) did not have any effect. Finally, a device reset removed the row SEFI.

All runs used in this section showed persistent SEFIs; that is, they were present even after rewriting the device. SEFIs that could be removed by rewriting the device were not examined.

4.2.1 Individual measure effectiveness

Figures 12 to 20 show the effectiveness of the individual measures, i. e. the number of times a measure was effective in resolving a particular SEFI class, divided by the number of times a measure was attempted in the presence of such a SEFI. “All group A” describes all measures of group A performed together, without considering the individual measures.

ZQ calibration never caused a change. Note, however, that for some parts, it was never performed individualy.

For most parts, **rewriting MR0** removed the SEFI condition for in some cases, but not in others.

For most parts, **rewriting MR1** removed the SEFI condition in some cases, but not in others.

For the 2-Gbit Hynix and 2-Gbit Micron parts, **rewriting MR2** removed the SEFI condition in some cases, but not in others. It did not have any effect for any other part. In one case with a 2-Gbit Micron part, additional column SEFIs were present after rewriting MR2.

For the 2-Gbit Hynix part, **rewriting MR3** removed the SEFI condition in some cases, but not in others. It did not have any effect for any other part.

DLL reset rarely had an effect on a SEFI. Furthermore, since a DLL reset always includes rewriting MR0, it is hard to tell whether a change was caused by the DLL reset or the mode register rewrite, unless MR0 was unsuccessfully rewritten before.

For all parts, **resetting the DUT** removed the SEFI in about half of the cases where it was attempted.

Since a **PHY reset** always includes a device reset, it was only attempted after a device reset had failed. It removed the SEFI in most, but not all cases.

In all cases where a **power cycle** was performed, it removed all SEFIs.

4.2.1.1 After device SEFI

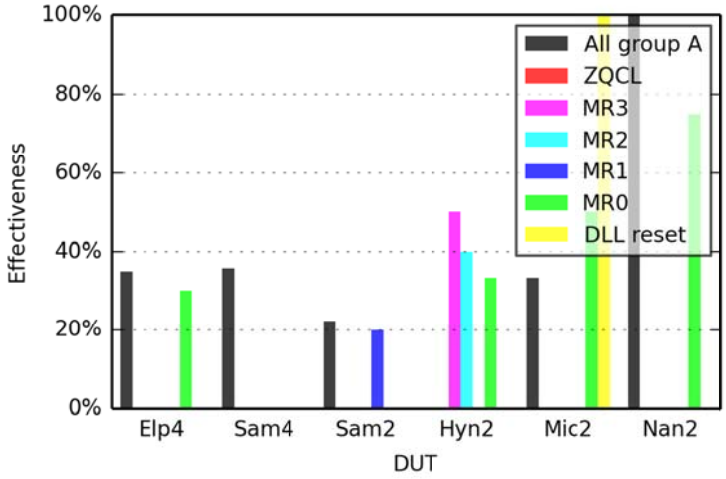

Figure 12 – device SEFI mitigation, measure effectiveness, group A

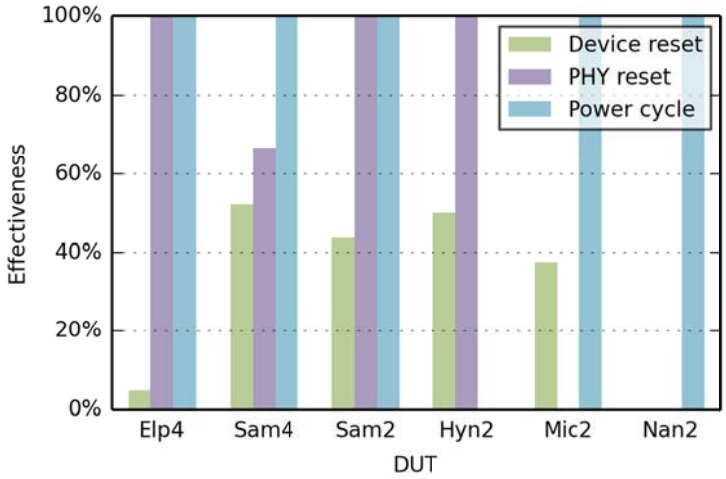

Figure 13 – device SEFI mitigation, measure effectiveness, groups B and C

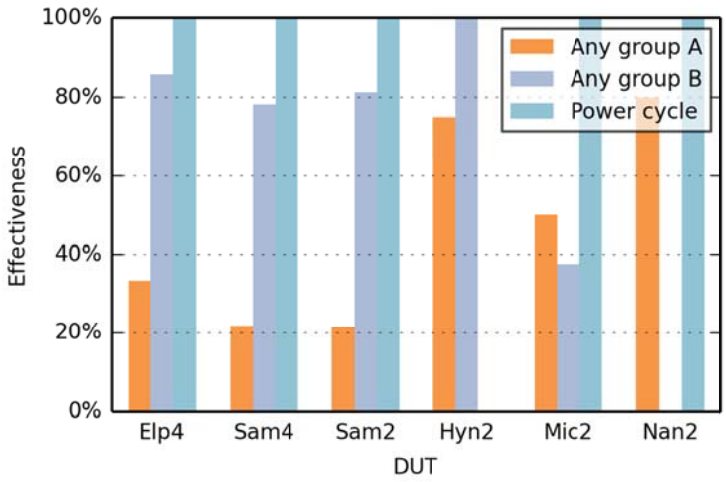

Figure 14 – device SEFI mitigation, measure effectiveness, by group

4.2.1.2 After row SEFI

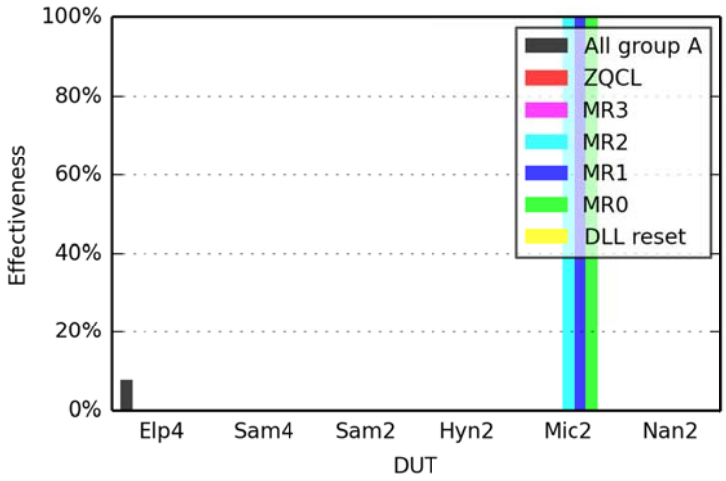

Figure 15 – row SEFI mitigation, measure effectiveness, group A

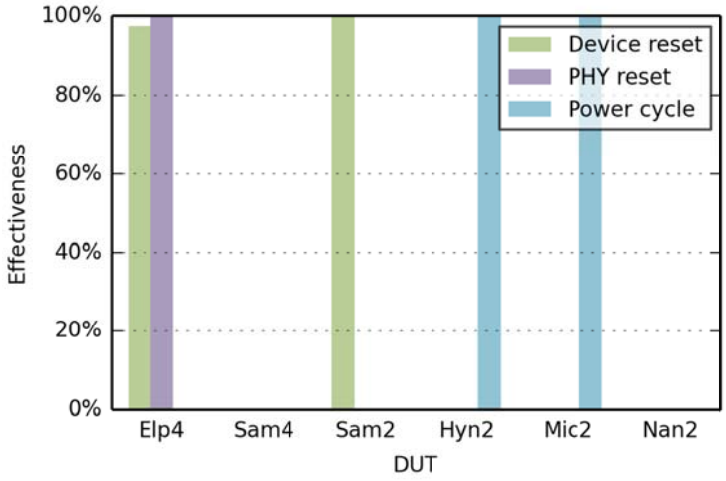

Figure 16 – row SEFI mitigation, measure effectiveness, groups B and C

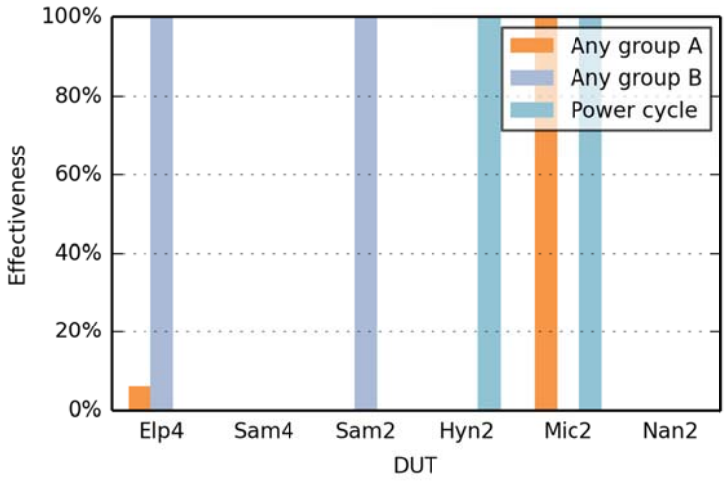

Figure 17 – row SEFI mitigation, measure effectiveness, by group

4.2.1.3 After column SEFI

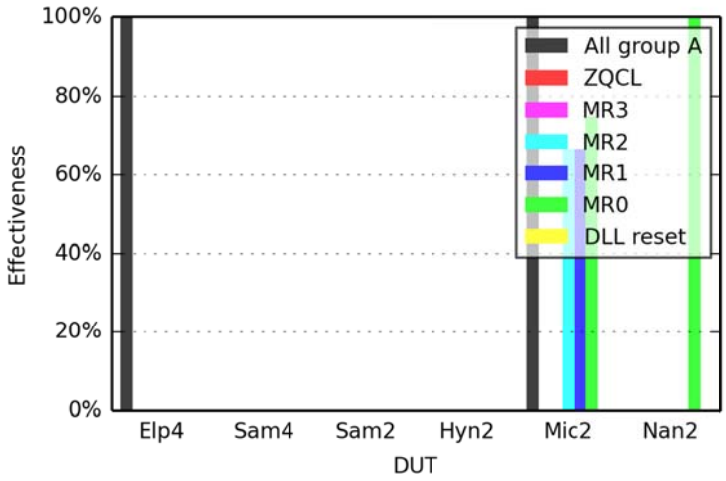

Figure 18 – column SEFI mitigation, measure effectiveness, group A

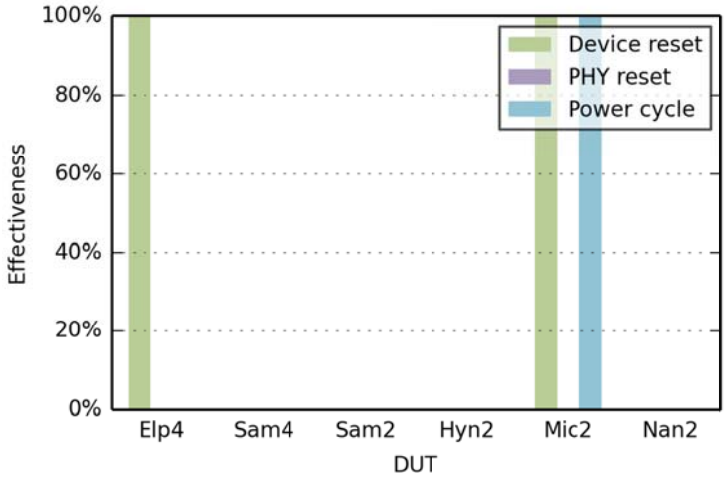

Figure 19 – column SEFI mitigation, measure effectiveness, groups B and C

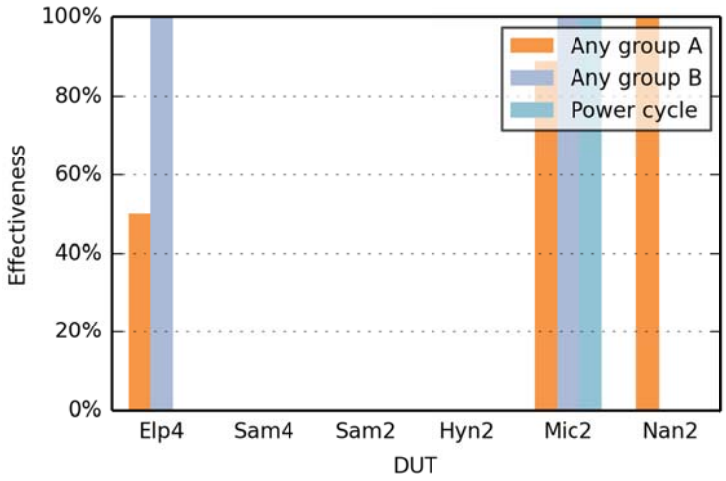

Figure 20 – column SEFI mitigation, measure effectiveness, by group

4.2.2 Required measure

In order to obtain a result that is relevant for a practical application, we determined the effectiveness of the three groups of measures, as defined in section 4.2. Since the timing impact of the individual measures is low, and the data impact of all measures of a given group is identical, it is reasonable to assume that all measures of a group can always be performed together. The interesting result is then the effectiveness of the individual groups, e. g. what percentage of SEFIs can be removed with guaranteed data retention (group A).

Given a group of SEFI mitigation tests (e. g. for a given part at a given LET), we can, for each group g (with g in [A, B, C]), count the number of times that a measure from this group has been performed, n_g, as well as the number of times that any measure from this group was effective, e_g. The effectiveness of this group (or the probability of success) is then given by $p_g = n_g / e_g$. Note in particular that $p_C = 1$.

The value we're interested in is the *required group*. This is the "first" group (the one with the lowest data impact) that is sufficient to remove a SEFI. This is motivated by the fact that each group includes all measures from the previous group and the data impact is monotonically increasing from group to group.

The percentage of cases where a group g is the required group, r_G is given by the effectiveness of this group, multiplied by the percentage of cases where all previous groups where ineffective:

$$r_A = p_A = p_A$$
$$r_B = p_B \cdot (1 - p_A) = p_B \cdot (1 - r_A)$$
$$r_C = p_C \cdot (1 - p_A) \cdot (1 - p_B) = p_C \cdot (1 - r_A - r_B)$$

With $p_C = 1$, we find that $r_A + r_B + r_C = 1$, meaning that any SEFI can be removed by one of the measures (but note that these measures may incur data loss).

Note that in some cases, no errors of a given error class were observed in any run, or a measure in a given group g was never attempted for this error class ($n_g = e_g = 0$). In this case, the effectiveness of the group (given by $p_g = n_g / e_g$) is undefined. Furthermore, the percentage of this group being the required group (r_g) is also undefined, as are all such values for subsequent groups.

This section shows the percentage of required measures for mitigation of a given SEFI class. Note that for some DUTs, we did not encounter any usable row or column SEFIs, or for the few SEFIs we did encounter, we did not attempt any of the measures from group A. Due to the reasons outlined in the preceding paragraph, no data is shown in this case.

4.2.2.1 Any ion

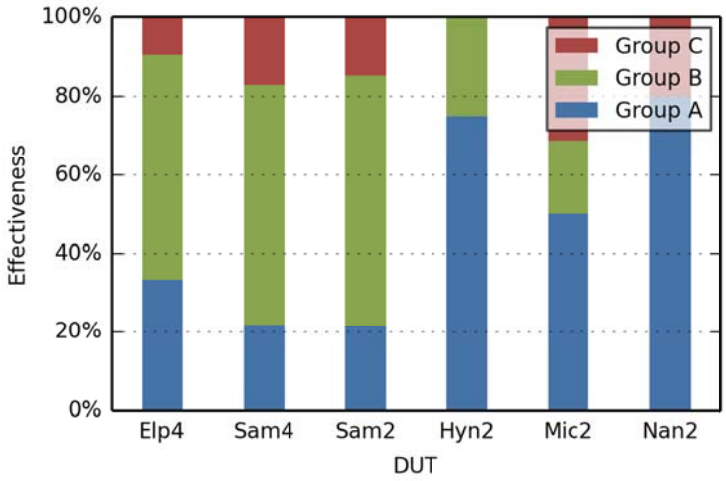

Figure 21 – device SEFI mitigation, required measure for different parts, by group

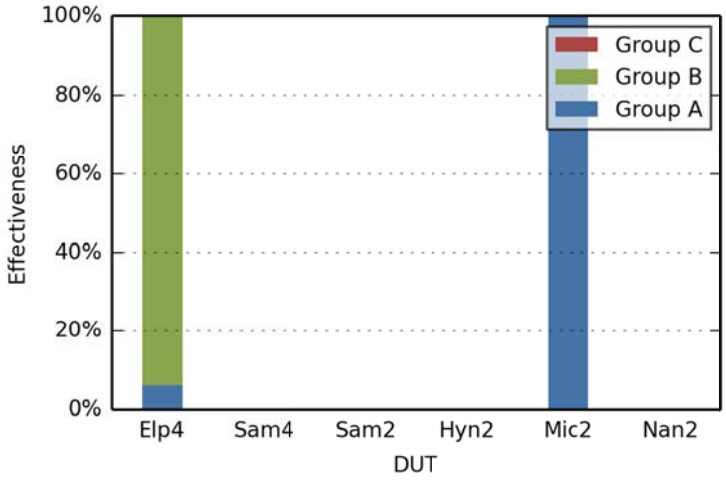

Figure 22 – row SEFI mitigation, required measure for different parts, by group

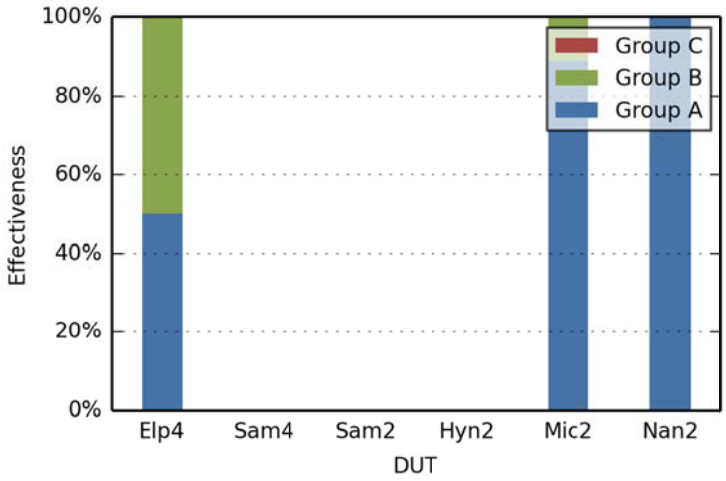

Figure 23 – column SEFI mitigation, required measure for different parts, by group

4.2.2.2 Ion comparison

In order to assess the dependence of SEFI mitigation measures on LET, in this section, we present the same data as in section 4.2.2, but separately for each ion. The amount of available data for a single ion may be very low.

For device SEFIs, useful data is only available for 4-Gbit Elpida, 4-Gbit Samsung, 2-Gbit Samsung and 2-Gbit Hynix parts.

For row SEFIs, useful data is only available for 4-Gbit Elpida.

For all other combinations of error class and part, data was either not available at all, or only for one ion, making the ion comparison pointless.

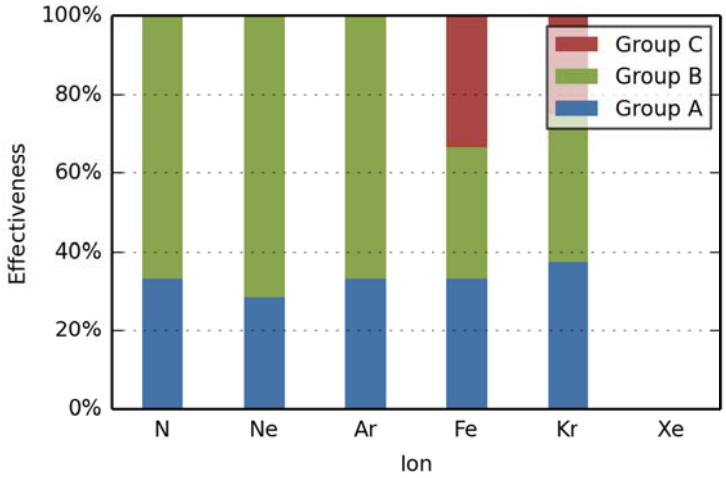

Figure 24 – device SEFI mitigation, 4-Gbit Elpida, required measure for different ions, by group

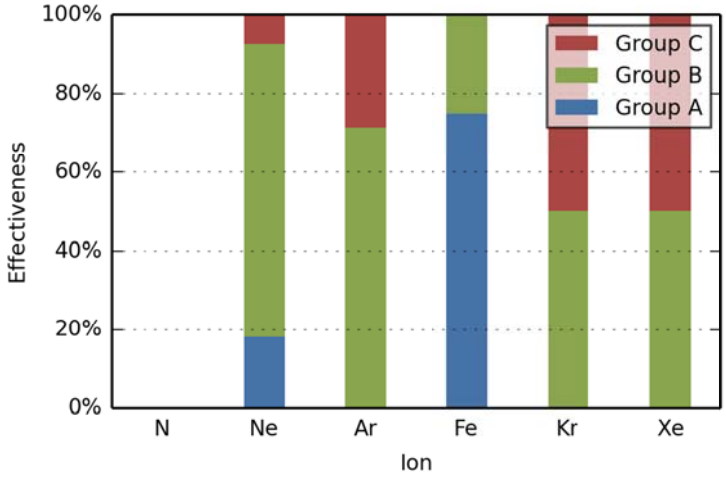

Figure 25 – device SEFI mitigation, 4-Gbit Samsung, required measure for different ions, by group

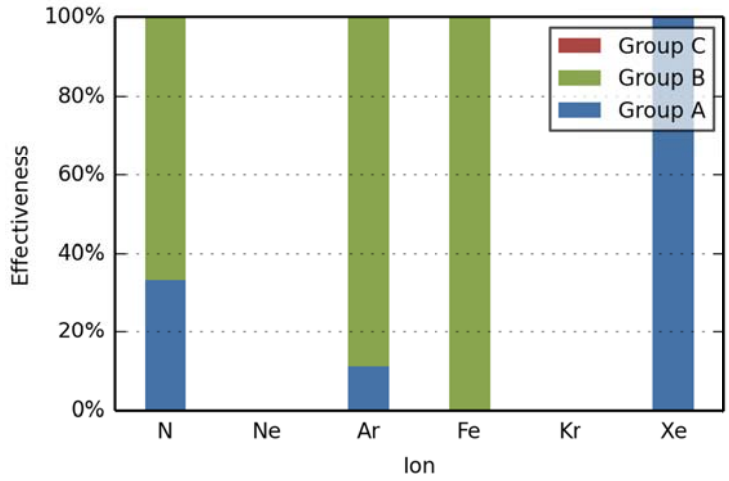

Figure 26 – device SEFI mitigation, 2-Gbit Samsung, required measure for different ions, by group

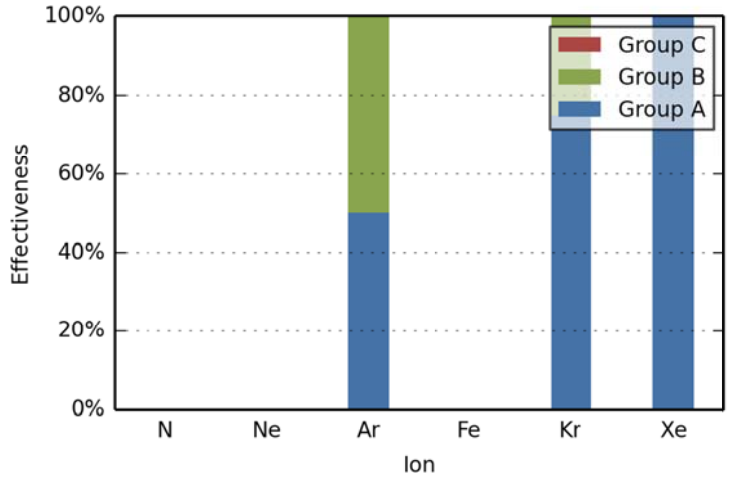

Figure 27 – device SEFI mitigation, 2-Gbit Hynix, required measure for different ions, by group

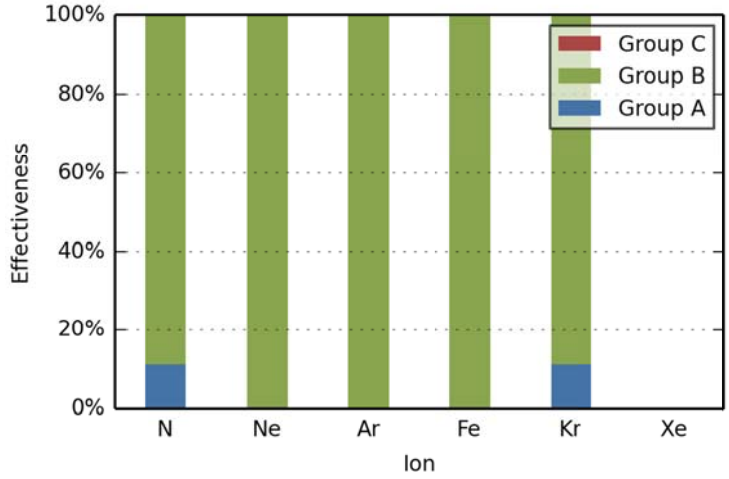

Figure 28 – row SEFI mitigation, 4-Gbit Elpida, required measure for different ions, by group

4.2.3 Discussion

With DDR3 SDRAM, it is not possible to read the mode register contents. It is therefore not possible to determine whether a mode register write removed a SEFI because the mode regis-

ter value was change by radiation, or because the mode register write triggered some internal process which caused the SEFI to disappear.

Since irradiation is very time consuming, and the number of possible combinations of measures is very high, the number of runs is not sufficient to gain authoritative numbers, or to derive cross sections for the effectiveness of measures. It is possible that, for a given device, a given measure has only been attempted once. The results presented in this section can therefore only give a first orientation.

Group A measures, which preserve data, are ineffective in many cases. Group B measures are effective in most cases, but endanger the data. Group C (power cycle) is always effective, but it does not preserve the data. Since group A and group B measures are not always effective, an application may have to perform a power cycle, and must therefore be prepared for data loss.

4.3 Hard SEUs

Hard SEUs (also called *stuck bits*) are SEUs that cannot be removed by writing the device. In other words, hard SEUs are SEUs that are present even if the device has not been irradiated between writing and reading.

For each run, pre-existing hard SEUs will also be included in the (total) SEU count determined after irradiation. In order to exclude these hard SEUs from the number of errors created during a run, the number of hard SEUs has been determined before each run by writing and immediately reading the device without intermediate irradiation.

In order to determine the number of hard SEUs created in a test run, we subtract the number of hard SEUs before the run from the number of hard SEUs after the run. The number of hard SEUs was not determined immediately after the run; instead, the number of hard SEUs before the next test run with the same DUT was used (note that this neglects the annealing of hard SEUs that occurs between a test run and the next one).

The cross sections of hard SEUs, determined in the 1/2012 test, are shown in figure 29. The cross sections for all SEUs (hard and soft) are shown as open symbols for comparison. Figure 30 shows the ratio of hard SEUs to total SEUs.

The ratio of hard SEUs to total SEUs is approximately constant over the whole LET range. It is slightly higher for the Samsung part than for the Elpida part. The threshold LET of hard SEUs is about the same as that of all SEUs for both parts.

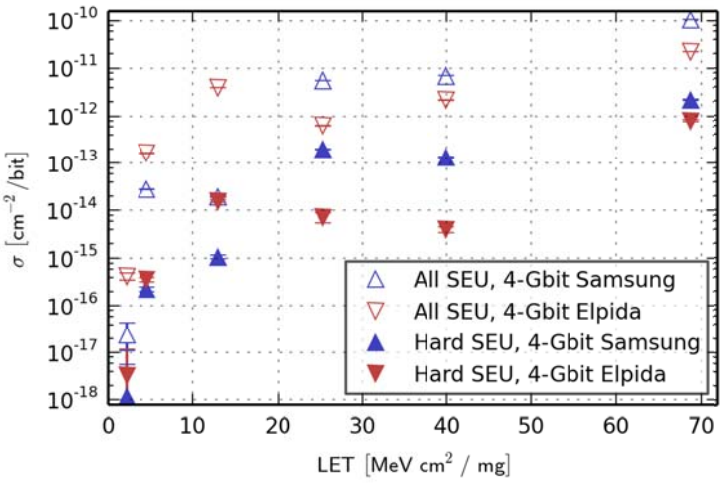

Figure 29 – Hard SEU cross sections, 1/2012. The cross sections for all SEUs (hard and soft) are shown as open symbols for comparison.

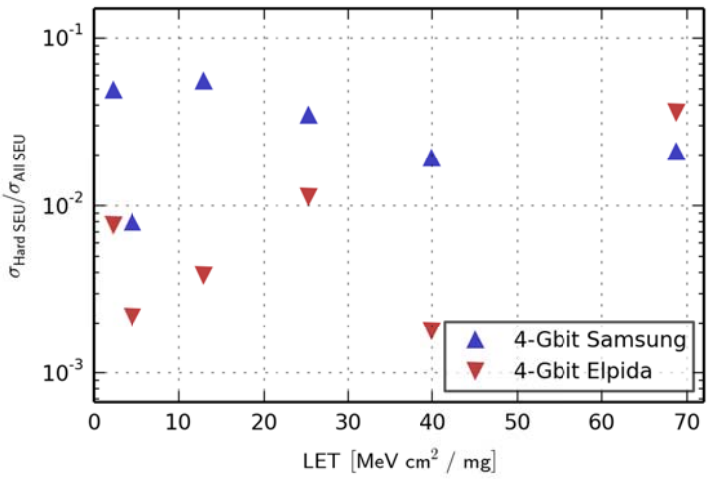

Figure 30 – Ratio of hard SEUs to total SEUs, 1/2012

Note that when plotting, for all runs, the number of hard SEUs created in a run versus the fluence of that run, it may appear that runs with higher fluence create fewer hard SEUs. This is caused by the fact that a higher fluence is typically used with lower LET ions, where fewer hard SEUs are created.

Hard SEUs may be explained by trapped charge that causes a shift of the cell transistor's threshold voltage, thereby permanently opening the cell transistor.

4.4 Hard SEU annealing

Annealing of hard SEUs was studied in the 4/2012 test. In order to determine the influence of LET on the annealing behavior, we used a "fresh" DUT (i. e., one that had not been irradiated before) for each ion. The only part type for which a significant number of fresh DUTs (5) were available was the 2-Gbit Samsung (rev. D) part. All annealing measurements were therefore performed with this part, using a different, fresh, DUT, for each ion. An exception was the part used for neon (AFSam1Db), which had been irradiated with krypton in the 1/2012 test. All but one of the hard SEUs from that test were determined to have annealed in the meantime.

The DUTs were first irradiated with sufficient fluence to generate a sufficient number of hard SEUs[2]. The number of hard SEUs was determined repeatedly in 5-minute intervals for 15 to 35 minutes (depending on the ion).

The number of hard SEUs was determined by writing and immediately reading the device. This is in contrast to the procedure we used for NAND Flash tests, where we did not repeatedly write the device. There are two reasons for this procedure: first, in a dynamic memory, the array is constantly read from and written to by the refresh operation. Second, when the power to the DUT is removed, the memory has to be rewritten anyway.

Previous experience with NAND Flash DUTs showed that opening the vacuum chamber can change the measured number of hard SEUs, probably due to the influence of DUT temperature. The chamber was therefore left closed for these measurements.

The number of hard SEUs was determined again on the last day of the test campaign and two weeks later, after the DUTs had been shipped back to the lab at Braunschweig. The DUT irradiated with Krypton, being the only one that still contained a significant number of hard SEUs after two weeks, was examined again 9 months after the irradiation and 13 months after the irradiation.

Note that the 5-minute-interval readings were performed without power cycling the DUT. After that, and before the reading at the end of the test campaign, the DUT was removed from the test device, powering it down.

Figure 31 shows the raw number of hard SEUs versus time on a log-log scale. Note that the value for xenon after 13 months is slightly higher than after 9 months. This may be due to a higher ambient temperature.

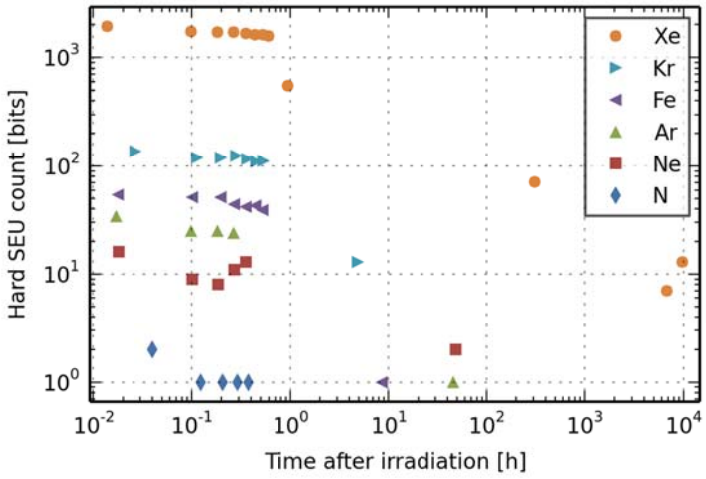

Figure 31 – Hard SEU annealing, 2-Gbit Samsung (rev. D), 4/2012

It has been observed that the number of hard SEUs strongly depends on the DUT temperature. It is therefore advisable to control or measure the DUT temperature for all measurements where the DUT may contain hard SEUs.

5 References

[2] In most cases, this run also caused a device SEFI.

[1] M. Herrmann, K. Grürmann und F. Gliem, „TN-IDA-RAD-14/3 - In-situ and unbiased TID Test of 4-Gbit DDR3 SDRAM Devices," 2014.

[2] M. Herrmann, K. Grürmann und F. Gliem, „TN-IDA-RAD-13/12 - Unbiased preselection TID Test of 4-Gbit DDR3 SDRAM Devices; Approach for in-situ test," 2013.

[3] M. Herrmann, K. Grürmann und F. Gliem, „TN-IDA-RAD-13/4 - TID Test of 4 Gbit DDR3 SDRAM Devices," 2013.

[4] M. Herrmann, K. Grürmann und F. Gliem, „TN-IDA-RAD-13/9B - Unbiased TID Test of 2-Gbit and 4-Gbit DDR3 SDRAM Devices; Proposal for Revision of the Test Strategy," 2013.

[5] JEDEC, *JESD79-3E - DDR3 SDRAM Specification.*

6 Appendix A: Full cross section charts

The following charts show the cross section versus LET. The cross sections are determined by dividing the number of errors by the fluence. If multiple runs were performed with the same parameters, they are combined by dividing the total number of errors from all runs by the total fluence from all runs. Open symbols represent cases where no errors of the respective class were observed. In this case, the value is determined by dividing 1 by the fluence. The result represents an upper limit for the cross section.

The error bars indicate the usual 5% Poisson error bounds. In cases where no errors were observed, the error bars extend to zero in the negative direction, which corresponds to negative infinity in the logarithmic plot.

Each measurement is characterized by six attributes:

- The part number (7 parts; see table 1)
- The LET (6 ions; see table 3)
- The test mode (read, write/read, or storage; see section 3.4)
- Whether software conditioning was used (see section 3.4)
- The error class (SEU, row SEFI, column SEFI, or device SEFI; see section 3.6)
- The test campaign (1/2012 or 4/2012)

Together, these attributes form a six-dimensional parameter space containing more than 2000 points. Such an amount of data is hard to visualize in its entirety. The following subsections therefore show different views of the configuration space, focusing on different aspects of the results.

Note that in each subsection, the cross section axis is scaled identically for all charts describing the same error class. This allows easy visual comparison between the results for different test modes and test campaigns as well as the effect of software conditioning.

Because all measurements with the same ion are taken at the same LET, all such measurements (up to 7) would be displayed at the exact same horizontal position in the charts, making it virtually impossible to distinguish the measurements for different parts from each other. The LET values shown in the charts are therefore slightly changed in order not to place points at the same position. This does not represent an actual change in LET.

The SEU cross section charts also show the approximate die area per bit. This value was obtained by measuring the total die area and dividing it by the number of bits in the device, not taking into account the amount of the die area used for the periphery (control logic, sense amplifiers, I/O etc.) The measured values for the die area per bit are shown in table 2. Note that the value is not identical for the different parts. The value shown in the chart is $4 \cdot 10^{-10}$, which is close to the largest value for all DUTs.

6.1 Test mode comparison

The charts in this section show the cross sections for all test modes, for a given error class, part, software conditioning and test campaign. Note that no tests with software conditioning were performed in the 4/2012 campaign.

This representation shows the influence of the test mode on the cross section.

Generally, the results for read mode and write/read mode are very similar to those for storage mode. This is plausible because dynamic memories perform a periodic refresh operation (every 7.8 µs in the case of DDR3), which opens and precharges banks, similar to what happens during a write or read operation. Notable exceptions are:

- For the 4-Gbit and 2-Gbit (revision D) Samsung parts, the SEU cross sections are significantly higher in write/read mode than in the other modes (sections 6.1.2 and 6.1.4)
- For the 2-Gbit Micron part, the device SEFI cross sections are slightly lower in storage mode than in the other modes, but mostly within the error bars (section 6.1.27)

6.1.1 SEU – 4-Gbit Elpida

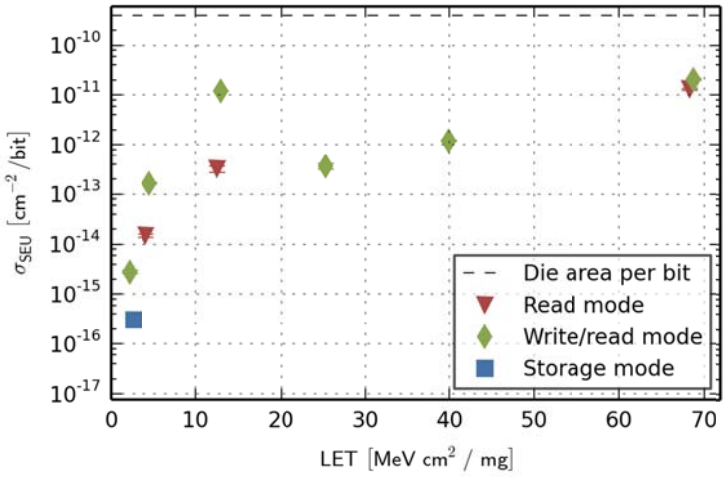

Figure 32 – SEU cross sections, 4-Gbit Elpida, without software conditioning, 1/2012

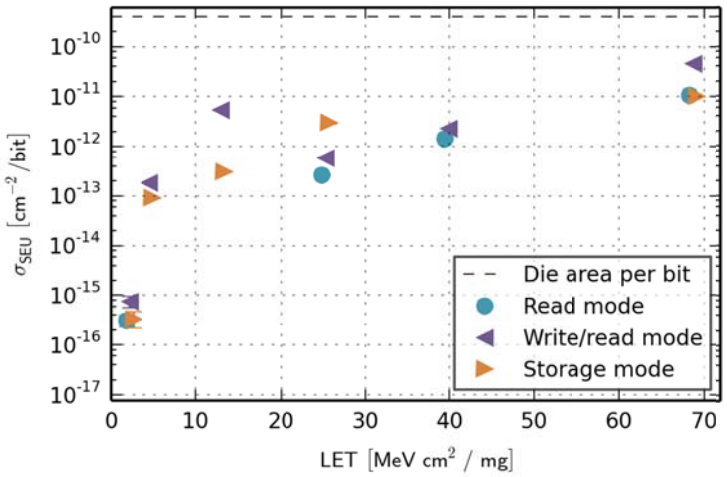

Figure 33 – SEU cross sections, 4-Gbit Elpida, with software conditioning, 1/2012

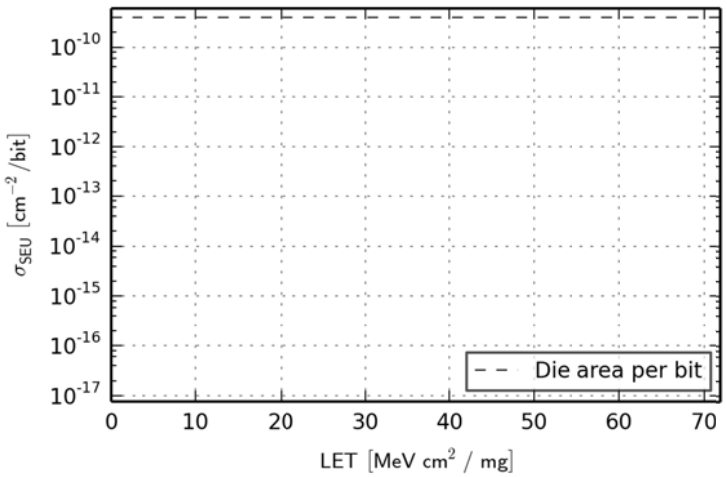

Figure 34 – SEU cross sections, 4-Gbit Elpida, without software conditioning, 4/2012

6.1.2 SEU – 4-Gbit Samsung

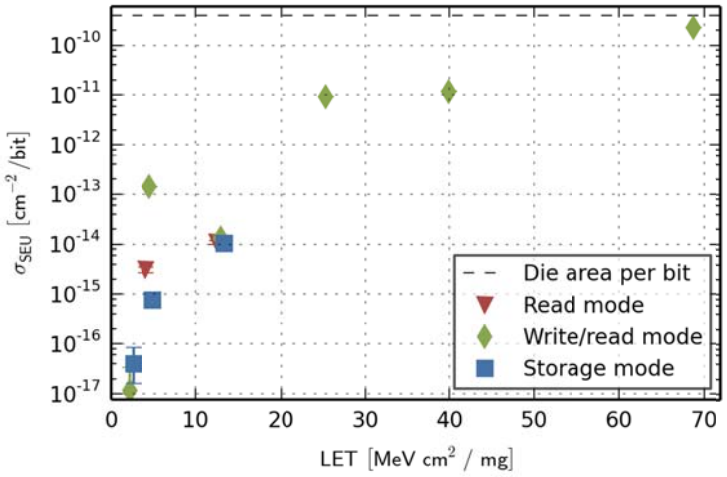

Figure 35 – SEU cross sections, 4-Gbit Samsung, without software conditioning, 1/2012

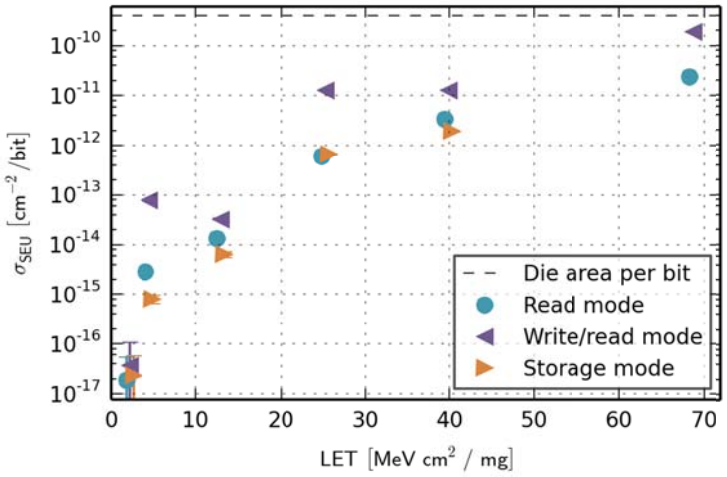

Figure 36 – SEU cross sections, 4-Gbit Samsung, with software conditioning, 1/2012

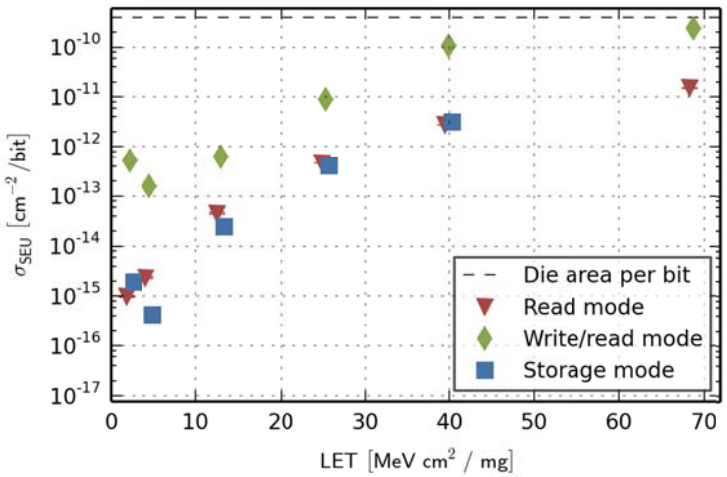

Figure 37 – SEU cross sections, 4-Gbit Samsung, without software conditioning, 4/2012

6.1.3 SEU – 2-Gbit Samsung (revision B)

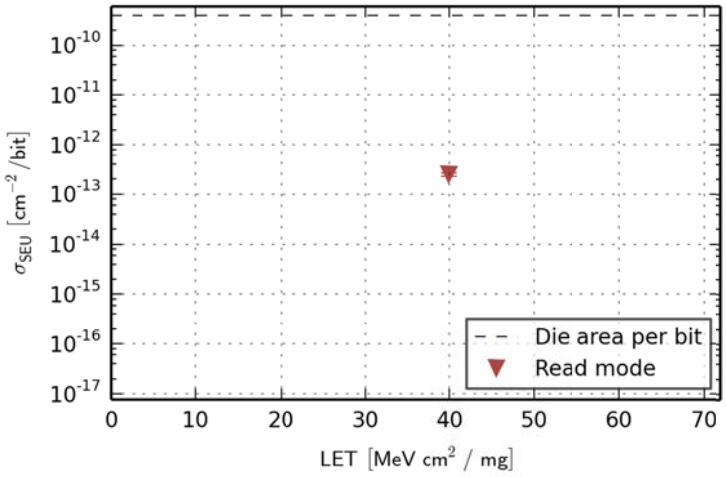

Figure 38 – SEU cross sections, 2-Gbit Samsung (revision B), without software conditioning, 1/2012

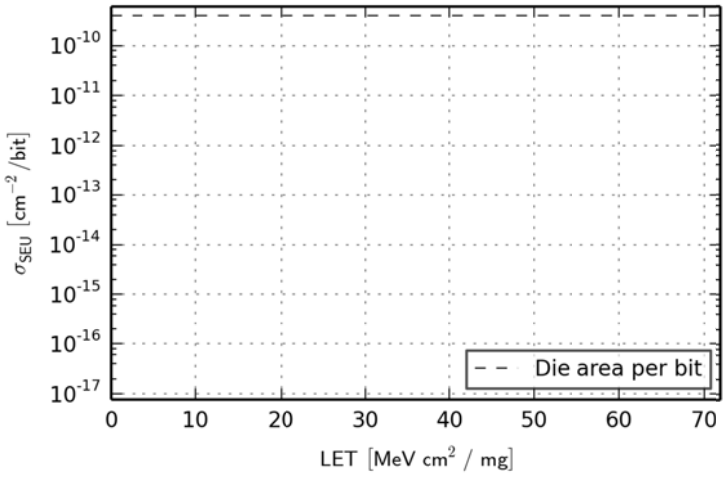

Figure 39 – SEU cross sections, 2-Gbit Samsung (revision B), with software conditioning, 1/2012

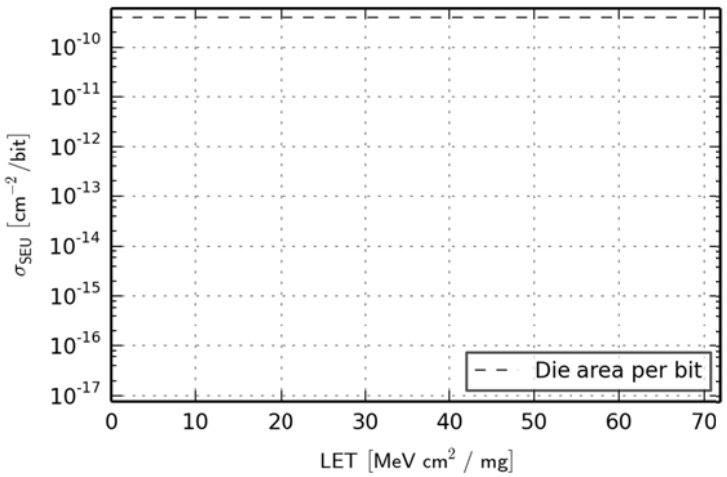

Figure 40 – SEU cross sections, 2-Gbit Samsung (revision B), without software conditioning, 4/2012

6.1.4 SEU – 2-Gbit Samsung (revision D)

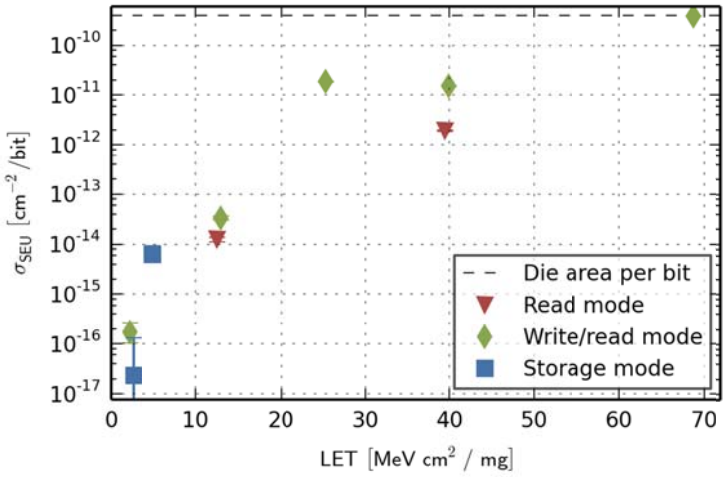

Figure 41 – SEU cross sections, 2-Gbit Samsung (revision D), without software conditioning, 1/2012

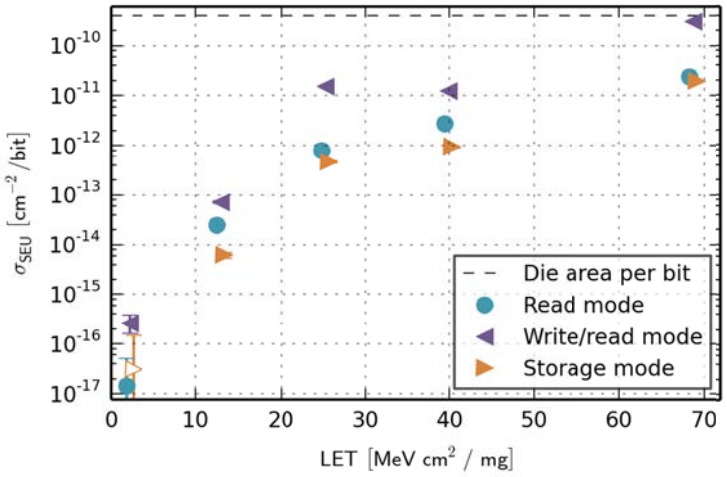

Figure 42 – SEU cross sections, 2-Gbit Samsung (revision D), with software conditioning, 1/2012

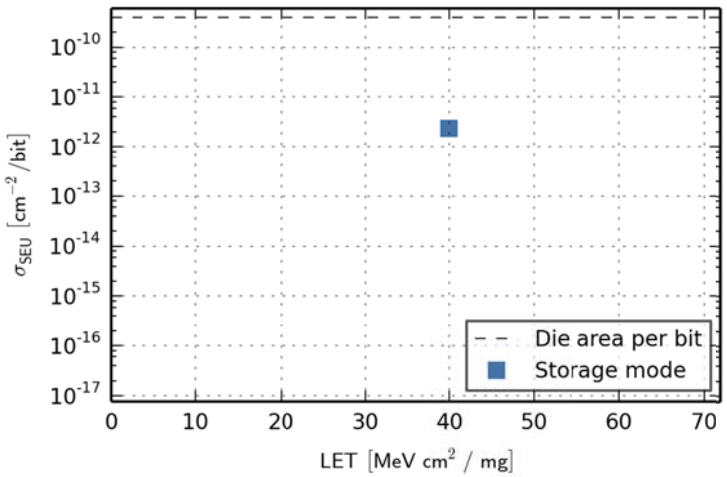

Figure 43 – SEU cross sections, 2-Gbit Samsung (revision D), without software conditioning, 4/2012

6.1.5 SEU – 2-Gbit Hynix

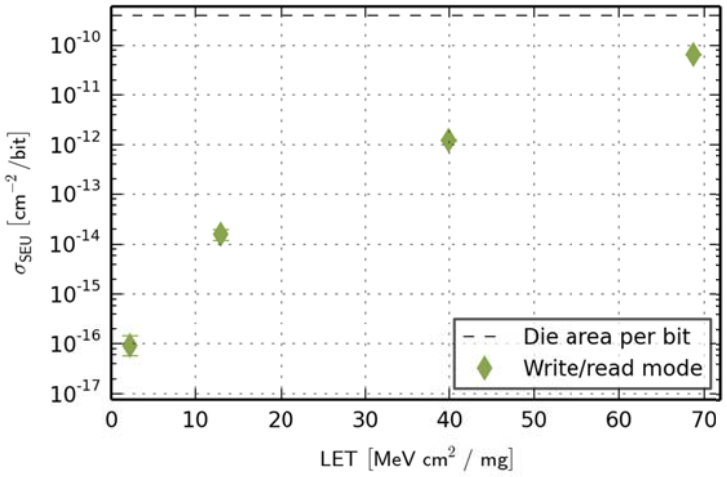

Figure 44 – SEU cross sections, 2-Gbit Hynix, without software conditioning, 1/2012

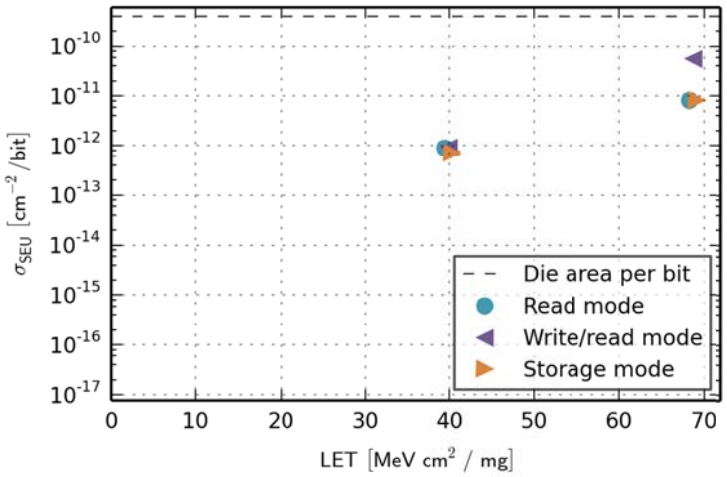

Figure 45 – SEU cross sections, 2-Gbit Hynix, with software conditioning, 1/2012

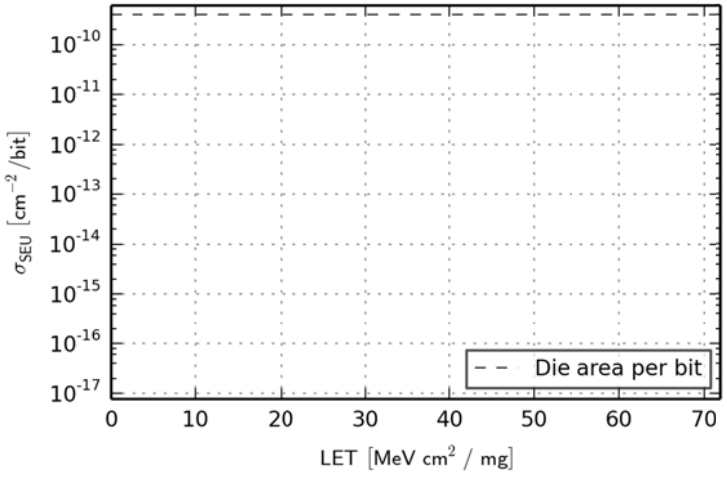

Figure 46 – SEU cross sections, 2-Gbit Hynix, without software conditioning, 4/2012

6.1.6 SEU – 2-Gbit Micron

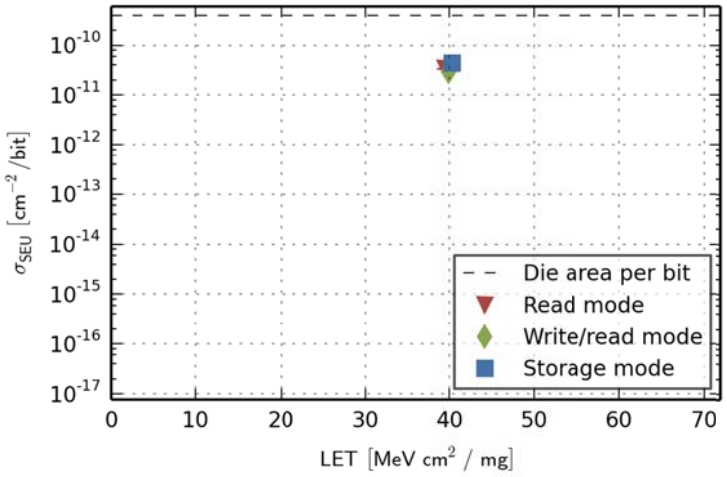

Figure 47 – SEU cross sections, 2-Gbit Micron, without software conditioning, 1/2012

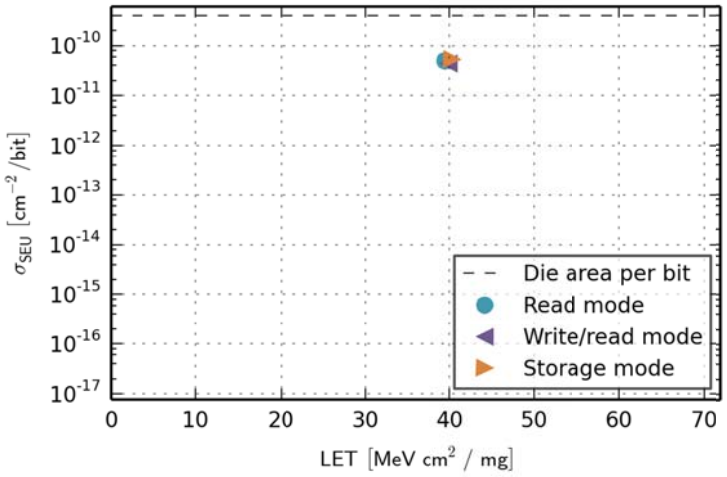

Figure 48 – SEU cross sections, 2-Gbit Micron, with software conditioning, 1/2012

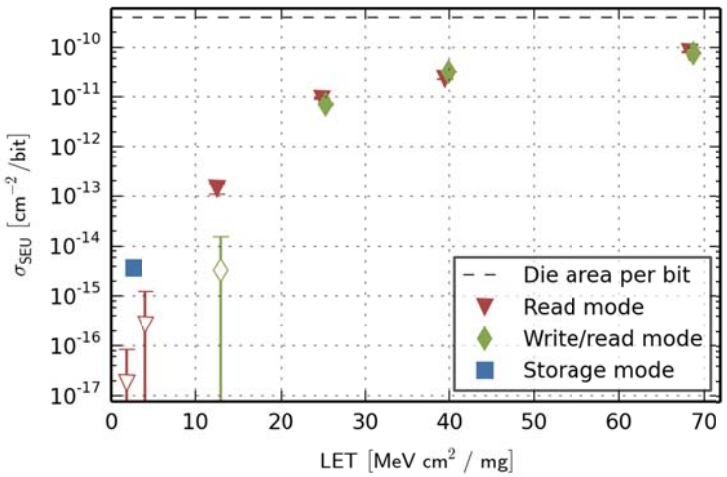

Figure 49 – SEU cross sections, 2-Gbit Micron, without software conditioning, 4/2012

6.1.7 SEU – 2-Gbit Nanya

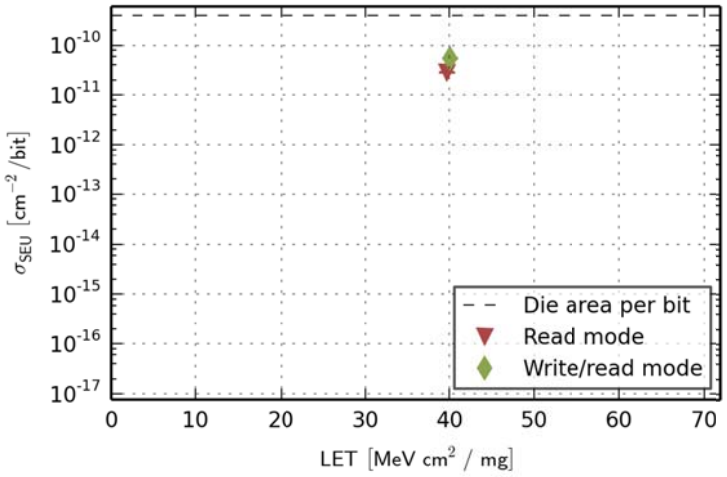

Figure 50 – SEU cross sections, 2-Gbit Nanya, without software conditioning, 1/2012

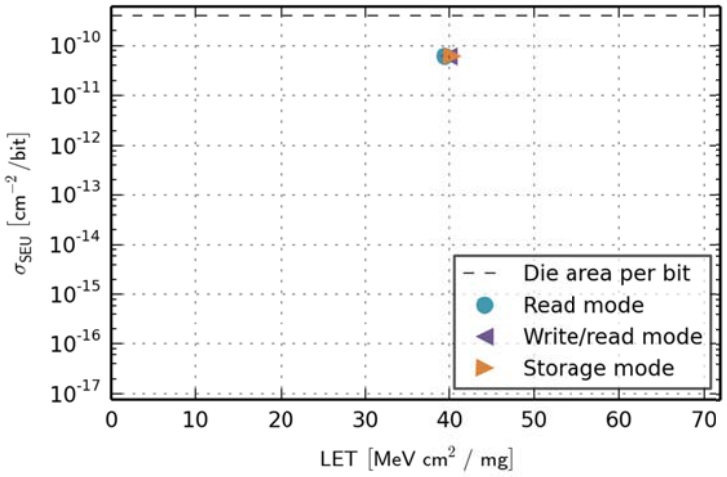

Figure 51 – SEU cross sections, 2-Gbit Nanya, with software conditioning, 1/2012

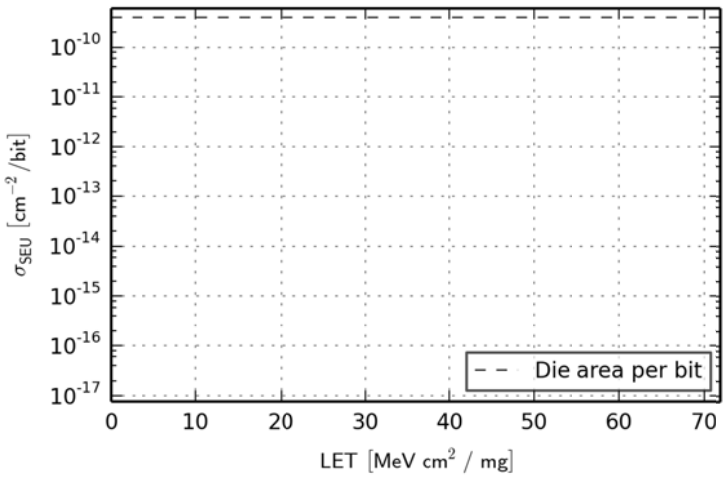

Figure 52 – SEU cross sections, 2-Gbit Nanya, without software conditioning, 4/2012

6.1.8 Row SEFI – 4-Gbit Elpida

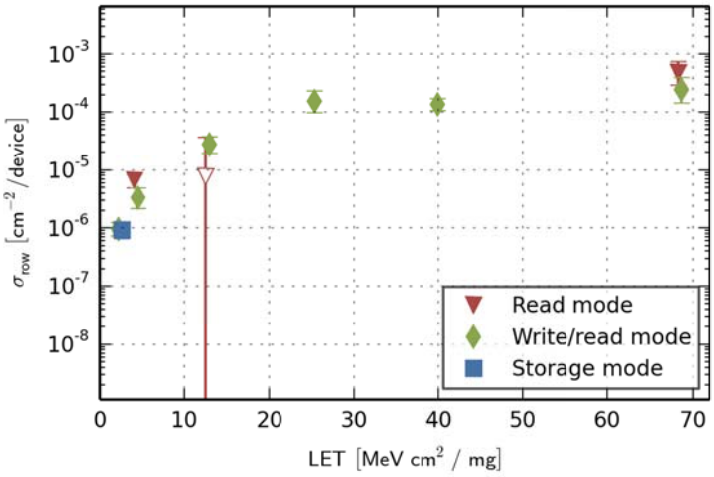

Figure 53 – Row SEFI cross sections, 4-Gbit Elpida, without software conditioning, 1/2012

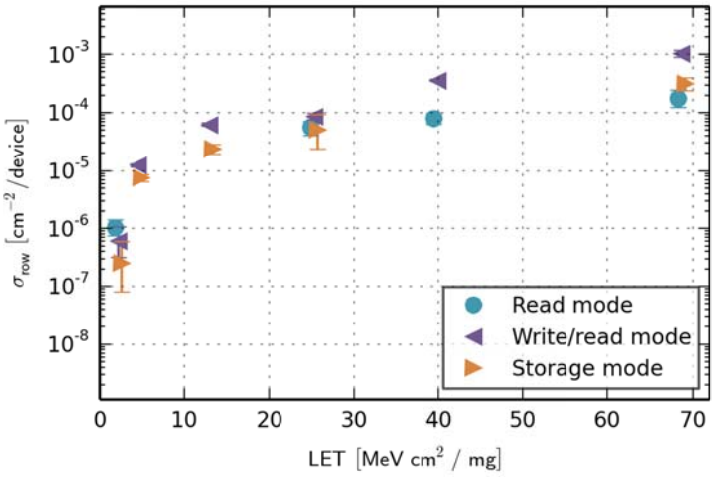

Figure 54 – Row SEFI cross sections, 4-Gbit Elpida, with software conditioning, 1/2012

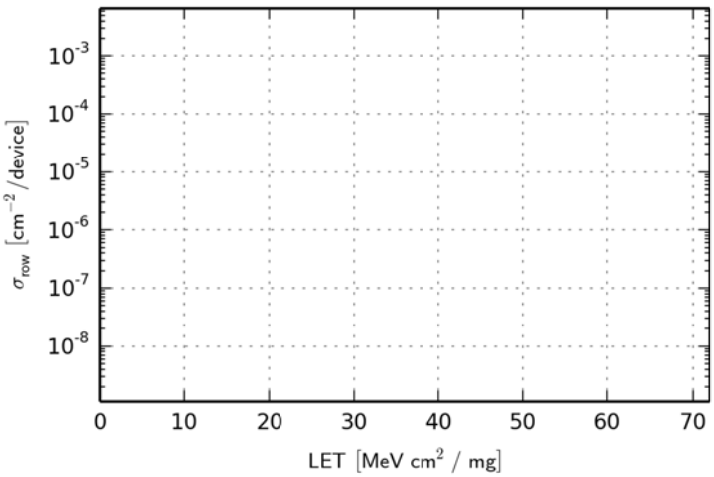

Figure 55 – Row SEFI cross sections, 4-Gbit Elpida, without software conditioning, 4/2012

6.1.9 Row SEFI – 4-Gbit Samsung

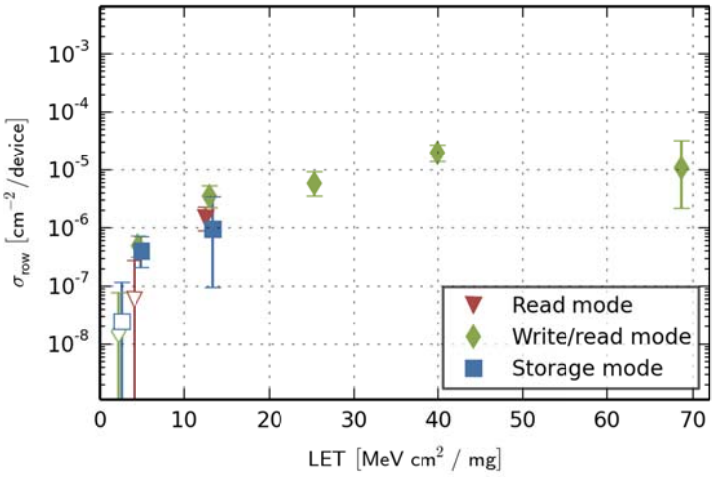

Figure 56 – Row SEFI cross sections, 4-Gbit Samsung, without software conditioning, 1/2012

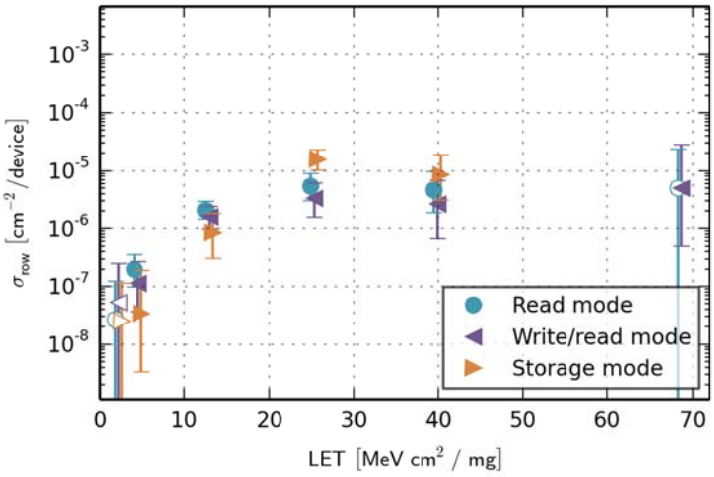

Figure 57 – Row SEFI cross sections, 4-Gbit Samsung, with software conditioning, 1/2012

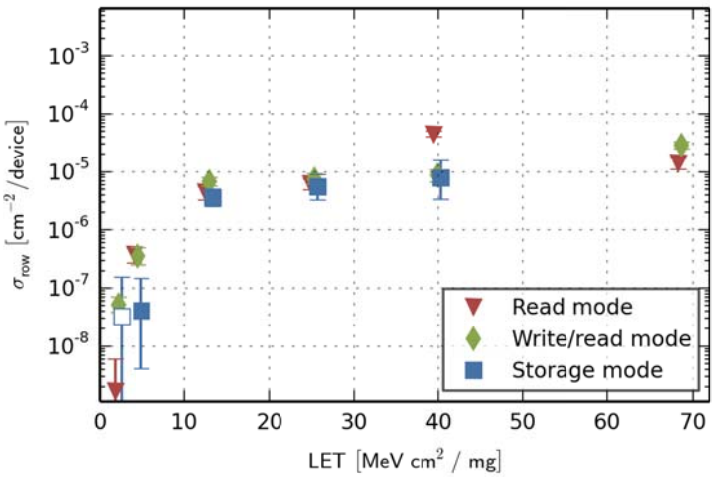

Figure 58 – Row SEFI cross sections, 4-Gbit Samsung, without software conditioning, 4/2012

6.1.10 Row SEFI – 2-Gbit Samsung (revision B)

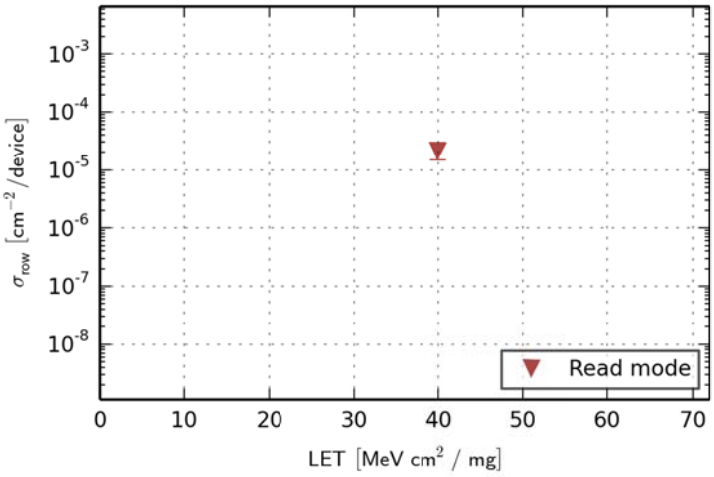

Figure 59 – Row SEFI cross sections, 2-Gbit Samsung (rev. B), without software conditioning, 1/2012

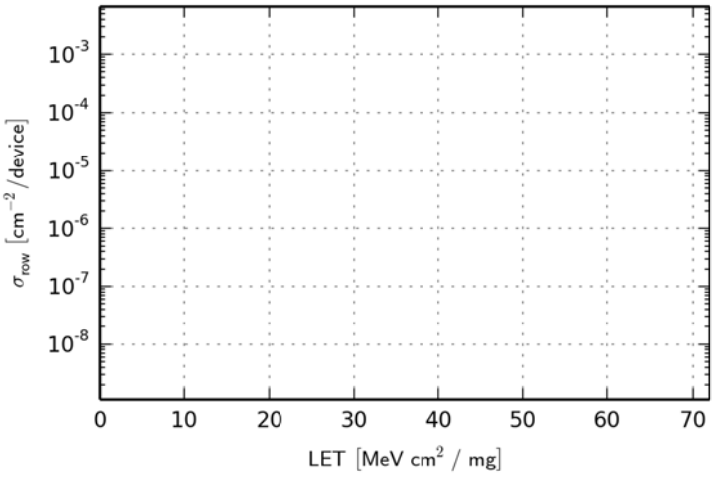

Figure 60 – Row SEFI cross sections, 2-Gbit Samsung (rev. B), with software conditioning, 1/2012

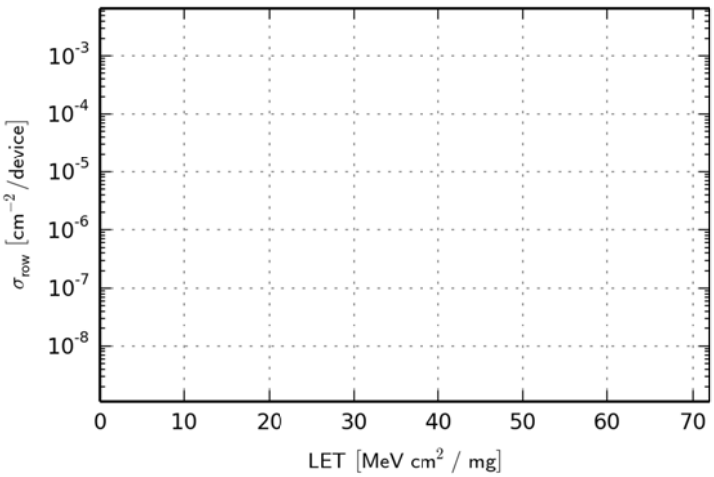

Figure 61 – Row SEFI cross sections, 2-Gbit Samsung (rev. B), without software conditioning, 4/2012