



Technische Universität
Braunschweig



INSTITUT FÜR DATENTECHNIK
UND KOMMUNIKATIONSNETZE

TN-IDA-RAD-13/1A

Heavy Ion SEE Test of 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.....	3
2	Test setup.....	3
2.1	DUTs.....	3
2.1.1	DUT preparation.....	4
2.2	Test facility	5
2.3	Test bench	9
2.4	Test sequence	11
2.5	Error classification	11
3	Test results.....	12
3.1	Cross sections.....	13
3.1.1	SEUs.....	13
3.1.2	Row SEFIs.....	15
3.1.3	Column SEFIs	16
3.1.4	Device SEFIs.....	18
3.1.5	Approximation.....	19
3.1.6	Error rates	19
3.1.7	Comparison with previous results	19
3.2	Single-event latch-up	21
3.3	Current	21
3.3.1	Current before and after irradiation.....	22
3.3.2	Current during irradiation.....	23
4	Future work.....	27

1 Abstract

From December 4 to December 7, 2012, we performed a heavy ion test campaign with DDR3 SDRAM and NAND Flash devices at Texas A&M University (TAMU), College Station, Texas, USA. This document reports on the findings for DDR3 SDRAM devices.

2 Test setup


2.1 DUTs

We tested five devices from Samsung and Hynix, as described in tables 1 and 2. All tested devices were thinned from approximately 300 μm to either 60 μm or 200 μm , with the exception of the Samsung 2 Gbit parts, which have an original thickness of 200 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 parts

Manufacturer and part number		Capacity	Lot code	Samples	Notes
Samsung	K4B4G0846B-HCH9	4 Gbit	GKM35990	5	Thinned to 200 μm
Samsung	K4B4G0846B-HCH9	4 Gbit	GMD02390	1	Thinned to 60 μm
Hynix	H5TQ4G83NFR-H9CR	4 Gbit	DTLB2032HM	5	Thinned to 200 μm
Samsung	K4B2G0846B-HCH9	2 Gbit	GLJ423AC	2	Not thinned, 200 μm
Samsung	K4B2G0846D-HCH9	2 Gbit	GEDN60GKU	2	Not thinned, 200 μm

Table 2: part photos

Part	Lot code	Photo
4-Gbit Samsung	GKM3599Q	
4-Gbit Samsung	GMD02390	Not available ¹
4-Gbit Hynix	DTLB2032HM	Not available ¹


2-Gbit Samsung (rev. B)

GLJ423AC



2-Gbit Samsung (rev. D)

GEDN60GKU



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

2.1.1 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 has been performed by Fraunhofer-Institut für Angewandte Optik und Feinmechanik (IOF). The thinning process is described by Fraunhofer 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 DDR3 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)

2.2 Test facility

The tests were performed at the Cyclotron Institute of TAMU. We used the 25 MeV/amu ion cocktail, described in table 3. This particular cocktail was chosen because it provides the necessary range (in particular for xenon) and has reasonable change times between different ions (see figure 3).

The LET vs. range curves for the 15, 25 and 40 MeV/amu cocktails are shown in figures 4 to 8.

Table 3: TAMU 25 MeV/amu ion cocktail

Ion	Energy	Range in Silicon [μm]	LET [$\text{MeV cm}^2 \text{mg}^{-1}$]		
			Bragg peak	200 μm	60 μm
^4He	99	3449	1.5	Not used	Not used
^{14}N	347	1009	6.7	1.0	Not used
^{22}Ne	545	799	9.7	1.8	Not used
^{40}Ar	991	493	20.1	7.3	Not used
^{84}Kr	2081	332	41.4	30.5	22.1
^{129}Xe	3197	286	63.4	61.1	Not used

Estimated Beam Change Times for 15 A MeV Beams

Final 15 A MeV Beam Configuration	⁴ He		45	55	55	55	55	55	55	55	70	70	70	60	60
	¹⁴ N	40		40	40	40	40	40	40	40	50	50	50	50	50
	²⁰ Ne	50	40		25*	35	30	30	30	30	45	45	45	50	50
	⁴⁰ Ar	50	40	25*		35	30	30	30	30	45	45	45	50	50
	⁶³ Cu	50	40	35	35		35	35	35	35	45	45	45	50	50
	⁸⁴ Kr	50	40	30	30	35		25*	25	30	45	45	45	50	50
	¹⁰⁹ Ag	50	40	30	30*	35	15		15	30	45	45	45	50	50
	¹²⁹ Xe	50	40	30	30	35	25	25*		30	45	45	45	50	50
	¹⁴¹ Pr	50	40	25	20*	35	25	25	25		45	45	45	50	50
	¹⁶⁵ Ho	60	50	45	45	45	45	45	45	45		30	30	65	65
	¹⁸¹ Ta	60	50	45	45	45	45	45	45	45	30		20	65	65
	¹⁹⁷ Au	60	55	50	50	50	50	50	50	50	30	20		70	70
	He	N	Ne	Ar	Cu	Kr	Ag	Xe	Pr	Ho	Ta	Au	25 MeV	40 MeV	
	Initial 15 A MeV Beam Configuration											w/Energy Changes			

Estimated Beam Change Times for 25 A MeV Beams

Final Beam Config.	⁴ He		50	40	40	40	20	60	60
	¹⁴ N	50		35	35	35	45	50	50
	²² Ne	40	35		30	30	45	50	50
	⁴⁰ Ar	40	35	30		30	45	50	50
	⁸⁴ Kr	40	35	30	30		45	50	50
	¹²⁹ Xe	40	60	50	50	50		70	70
	He	N	Ne	Ar	Kr	Xe	15 MeV	40 MeV	
	Initial Beam Config.						w/Energy Changes		

Estimated Beam Change Times for 40 A MeV Beams

Final Beam Config.	¹⁴ N		35	35	35	80	80
	²⁰ Ne	35		20	30	80	80
	⁴⁰ Ar	35	20		30	80	80
	⁷⁸ Kr	50	50	50		120	120
		N	Ne	Ar	Kr	15 MeV	25 MeV
	Initial Beam Config.				w/Energy Changes		

- R** Extra tuning and conditioning times are usually required
- O** Extra tuning and conditioning times are often required
- Y** Extra tuning and conditioning times are sometimes required
- G** Changes are typical, but some fine tuning is required
- B** Changes are typical and usually go smoothly
- V** Very few changes are made, rarely takes longer than specified.

Possible Delays

- several hours
- >30 minutes
- <30 minutes
- a few minutes
- a few minutes
- a few minutes

*These changes can be very fast (<15 minutes) if you give your SEE Manager advanced notice

Figure 3: TAMU beam change times (TAMU)




Figure 4: LET versus range, 15 MeV/amu cocktail, low LET ions (TAMU)

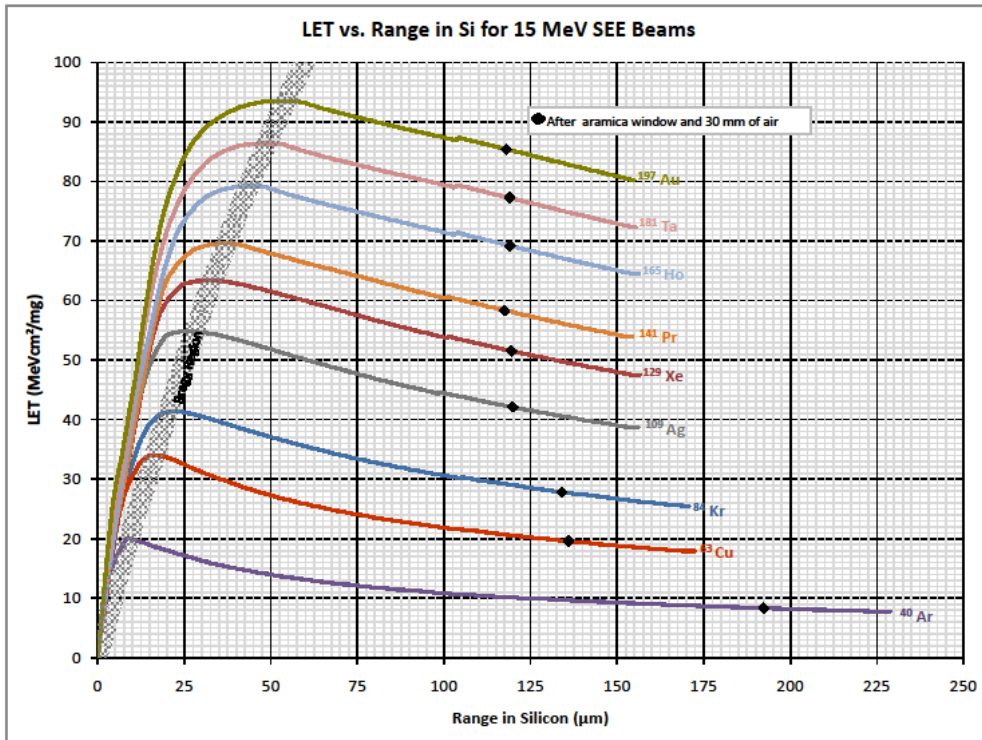


Figure 5: LET versus range, 15 MeV/amu cocktail, high LET ions (TAMU)




Figure 6: LET versus range, 25 MeV/amu cocktail, low LET ions (TAMU)




Figure 7: LET versus range, 25 MeV/amu cocktail, high LET ions (TAMU)




Figure 8: LET versus range, 40 MeV/amu cocktail (TAMU)

2.3 Test bench

The test bench, RTMC6 (figures 9 and 10), 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.

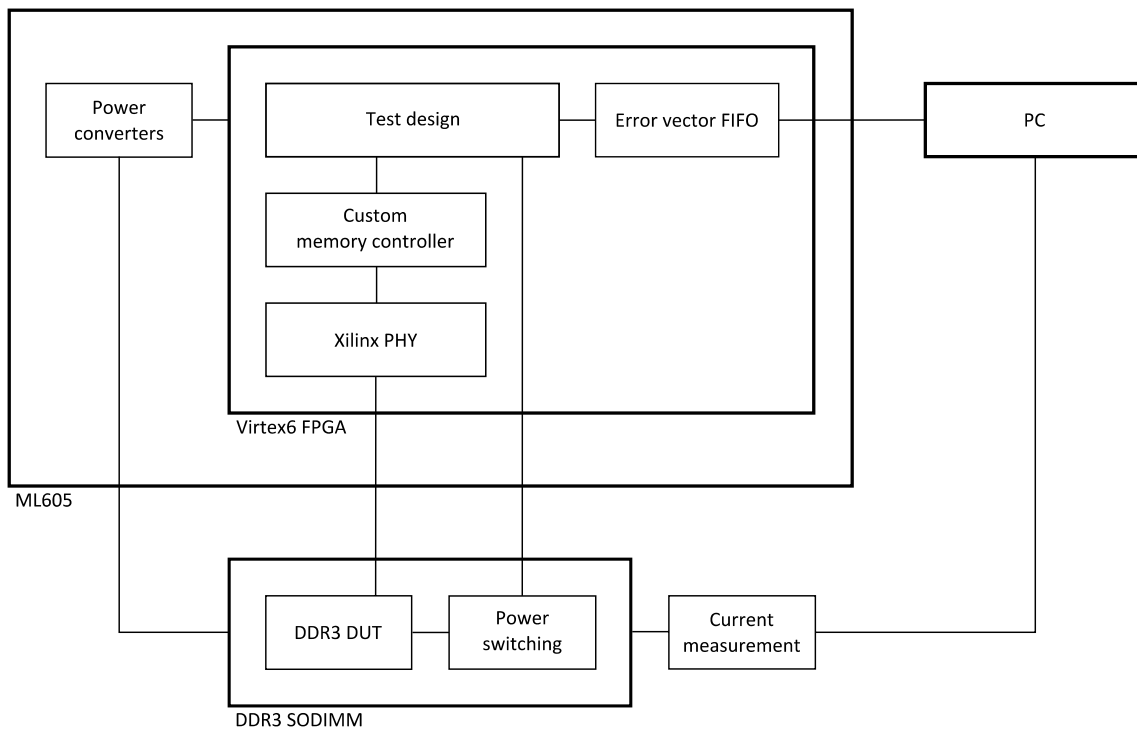


Figure 9: an overview of the RTMC6 test bench (simplified)




Figure 10: 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 DUT current is measured at a sampling rate of 1 Hz and is logged by the PC.

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 either discarded or the test is slowed down, at the user's choice.

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.

2.4 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.

All test modes can be performed with and without *software conditioning*. 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)

2.5 Error classification

Several error classes are distinguished, according to the overview shown in figure 11:

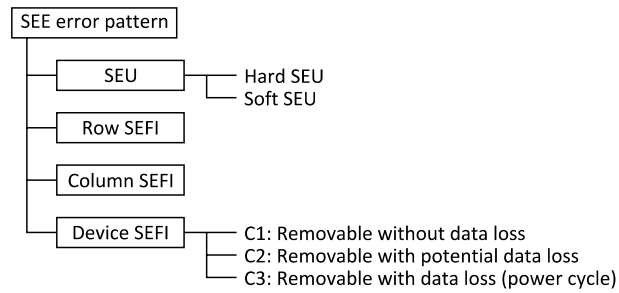


Figure 11: error classification overview

- 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 12.




Figure 12: an error map showing row SEFIs, column SEFIs and SEUs

3 Test results

Due to the limited available test time, all tests were performed *without* software conditioning. A pseudo-random pattern was used for all tests.

3.1 Cross sections

The following diagrams show the cross section versus LET for different devices, for a given error class and test mode. Open symbols represent runs where no errors of the respective class were observed. In this case, the value represents an upper limit.

The results for read mode are almost identical 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 read operation.

3.1.1 SEUs

In addition to the SEU cross sections, the diagrams 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 (small) amount of the die area used for the periphery (control electronics, row buffers, I/O etc.) The values for the die area per bit are shown in table 4.

Table 4: die area per bit for different DDR3 devices

Manufacturer and part number		Capacity	Die area per bit
Samsung	K4B4G0846B-HCH9	4 Gbit	$1.6 \cdot 10^{-10} \text{ cm}^2$
Hynix	H5TQ4G83NFR-H9CR	4 Gbit	$1.9 \cdot 10^{-10} \text{ cm}^2$
Samsung	K4B2G0846B-HCH9	2 Gbit	$4.0 \cdot 10^{-10} \text{ cm}^2$
Samsung	K4B2G0846D-HCH9	2 Gbit	$1.4 \cdot 10^{-10} \text{ cm}^2$




Figure 13: SEU cross section in read mode

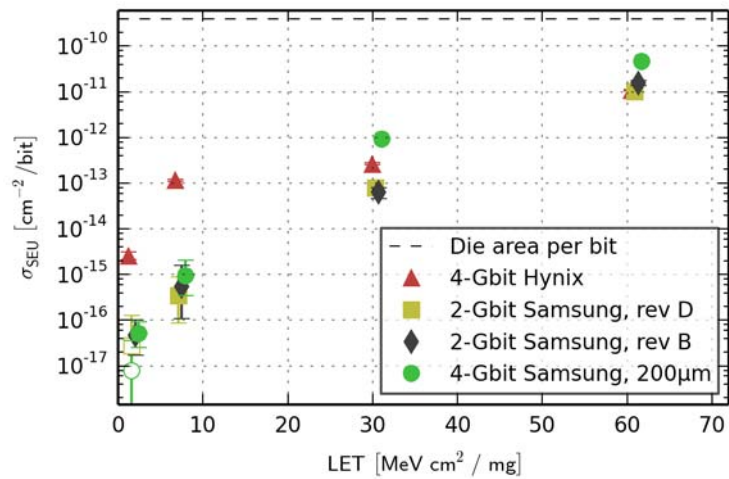


Figure 14: SEU cross section in write/read mode

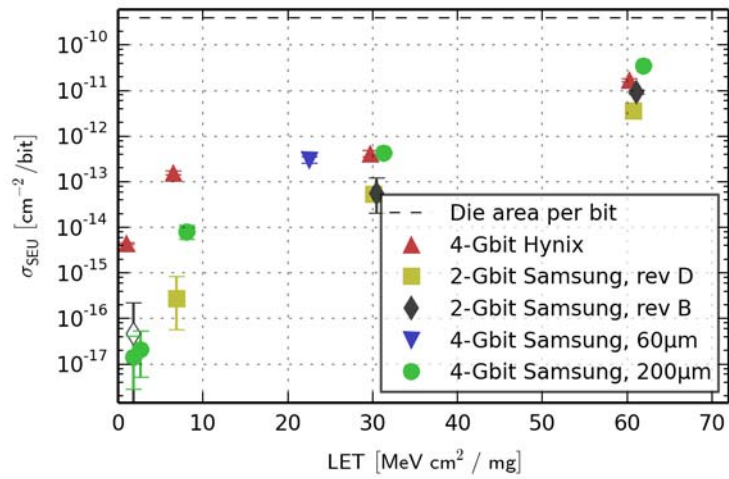


Figure 15: SEU cross section in storage mode




Figure 16: SEU cross section in any mode

3.1.2 Row SEFIs




Figure 17: row SEFI cross section in read mode




Figure 18: row SEFI cross section in write/read mode




Figure 19: row SEFI cross section in storage mode




Figure 20: row SEFI cross section in any mode

3.1.3 Column SEFIs

Column SEFIs are very rare. Out of 118 runs that could be analyzed, only 2 showed any column SEFIs at all.

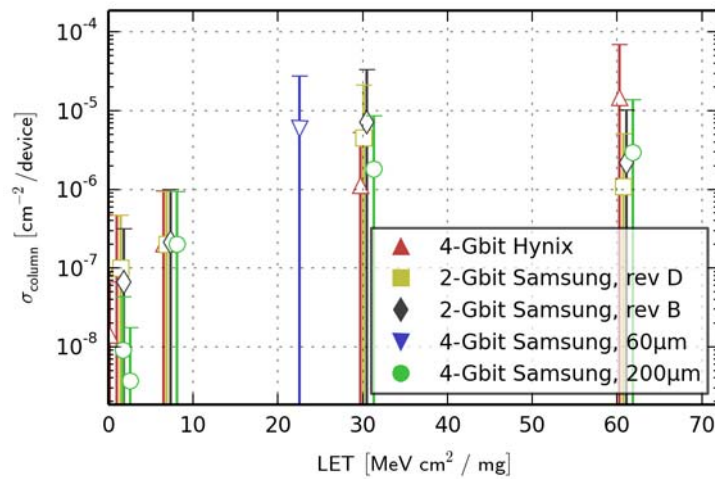


Figure 21: column SEFI cross section in read mode




Figure 22: column SEFI cross section in write/read mode




Figure 23: column SEFI cross section in storage mode




Figure 24: column SEFI cross section in any mode

3.1.4 Device SEFIs




Figure 25: device SEFI cross section in read mode




Figure 26: device SEFI cross section in write/read mode




Figure 27: device SEFI cross section in storage mode

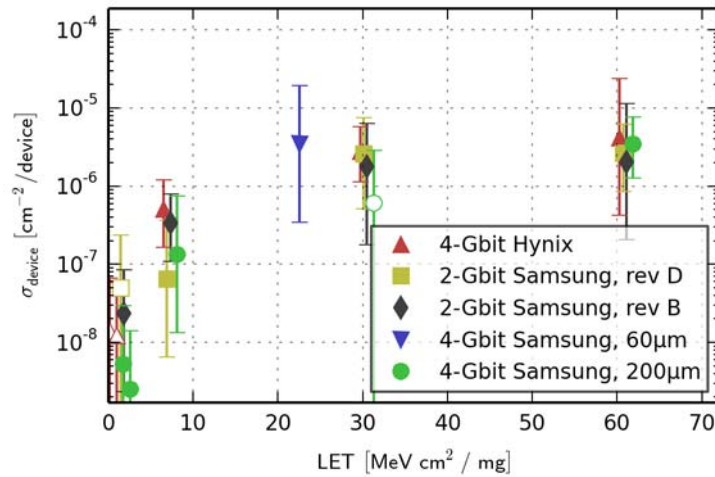


Figure 28: device SEFI cross section in any mode

3.1.5 Approximation

For the 200 μm, Samsung, 4 Gbit device, the parameters for a Weibull approximation of the cross sections for SEUs, row SEFIs and device SEFIs are given in table 5. These parameters were calculated using OMERE 3.4.0.0 (2009).

Table 5: Weibull approximation parameters for Samsung, 4 Gbit, 200 μm

Error class	W	S	Limit cross section [cm ²]	LET threshold [MeV cm ² / mg]
SEU cross section	57.56	3.63	$3.38 \cdot 10^{-11}$ / bit	0.47768
Row SEFI cross section	32.87	2.20	$4.78 \cdot 10^{-05}$ / device	0.47768
Device SEFI cross section	25.10	2.10	$3.49 \cdot 10^{-06}$ / device	0.47768

3.1.6 Error rates

Calculated error rates for errors caused by galactic cosmic rays (GCR) in a 1 AU solar orbit are shown in table 6. For each error class, three values are given, assuming a cell depth of 0.01 μm, 0.1 μm and 1 μm, respectively. The reasons for assuming a cell depth of 0.01 μm are detailed in TN-IDA-RAD-12/4. These error rates were calculated using OMERE 3.4.0.0.

Table 6: error rates for 1 AU orbit, GCR

Error class	Error rate (0.01 μm)	Error rate (0.1 μm)	Error rate (1 μm)
SEUs [errors/(bit-day)]	$3.21 \cdot 10^{-12}$	$1.33 \cdot 10^{-13}$	$1.17 \cdot 10^{-15}$
Row SEFIs [errors/(device-day)]	$8.55 \cdot 10^{-5}$	$8.55 \cdot 10^{-5}$	$8.19 \cdot 10^{-5}$
Device SEFIs [errors/(device-day)]	$1.05 \cdot 10^{-5}$	$1.05 \cdot 10^{-5}$	$9.19 \cdot 10^{-6}$

3.1.7 Comparison with previous results

Figures 29, 30 and 31 show the results of the 12/2012 TAMU campaign alongside the results from former test campaigns at RADEF, Jyväskylä, Finland. The figures show the SEU, row SEFI and device SEFI cross section, respectively, for Samsung 4 Gbit in all test modes without software conditioning.

The SEU cross section from the TAMU test campaign is about one order magnitude lower than the cross section from the RADEF campaigns (figure 29). This may have been caused by a difference in DUT temperature: the tests at TAMU were performed in air at a temperature of less than 20°C. The tests at RADEF, on the other hand, were performed in vacuum. The temperature of the DUT in vacuum is not known, but based on the power dissipation of approximately 150 mW, it is estimated to be about 60°C.

For the device SEFI cross sections, the TAMU values and the RADEF values are equal within the error margin resulting from the low number of observed SEFIs.




Figure 29: Samsung 4 Gbit, SEUs, any mode, no software conditioning

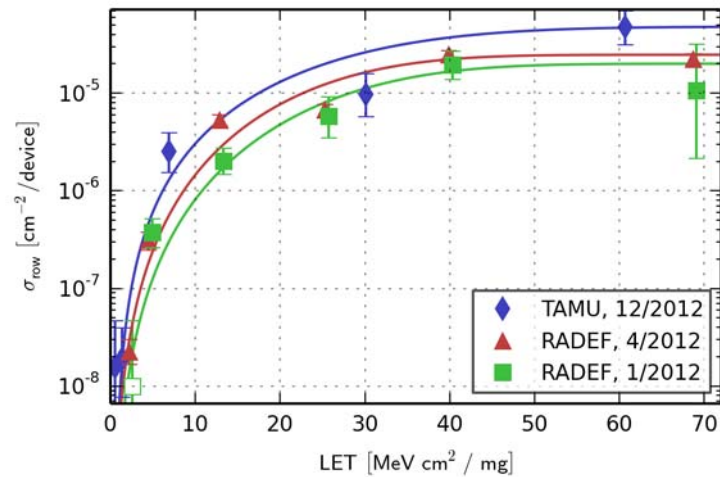


Figure 30: Samsung 4 Gbit, row SEFIs, any mode, no software conditioning




Figure 31: Samsung 4 Gbit, device SEFIs, any mode, no software conditioning

3.2 Single-event latch-up

In order to test for single-event latch-up (SEL) events, five 4 Gbit DUTs (two from Samsung and three from Hynix) were heated to 80 °C and irradiated with 10^7 cm^{-2} 61.1 MeV $\text{cm}^{-2} \text{ mg}^{-1}$ xenon ions while operating them in write/read mode. The heating was performed using the air flow from a heat gun and monitored with four temperature sensors arranged around the DUT.

No latch-up events were observed for any of the devices at the given LET.

3.3 Current

For the single-event latch-up test (section 3.2), the current was logged with a sampling rate of 1 Hz before, during, and after irradiation.

Figure 32 shows, for a Samsung 4 Gbit device, the current for 4 loops of a test run in write/read mode *without* irradiation at room temperature. The green line represents the current measured with a sample rate of 16 Hz. At 0 s, the device is idle with an idle current of approximately 20 mA. At 4 s, the device is power cycled and the test run is started. A current of approximately 100 mA indicates a write operation, while a current of approximately 65 mA indicates a read operation.

The blue crosses in figure 32 represent the result of the 1 Hz measurement. These values are the average current over a 1 second interval. It is clearly visible that a sample rate of 1 Hz is not sufficient to resolve the individual operations. Unfortunately, the 16 Hz measurement was not available for logging at the time of the test campaign.




Figure 32: high and low sample rate current measurement in write/read mode

The idle current is the current while only performing the required idle operations – refresh (every 7.8 μ s), periodic read (required by the Xilinx PHY, every 1 μ s) and short ZQ calibration (every 128 ms). This value does not represent any of the IDD_x values defined in the DDR3 specification.

The approximate *baseline* current values (i. e. the current values without irradiation) for idle, write and read operation are shown in table 7.

Table 7: approximate baseline current values

Device	Idle current	Write current	Read current
Samsung, 4 Gbit	20 mA	100 mA	60 mA
Hynix, 4 Gbit	20 mA	110 mA	75 mA
Samsung, 2 Gbit (revision B)	35 mA	120 mA	90 mA
Samsung, 2 Gbit (revision D)	15 mA	90 mA	55 mA

3.3.1 Current before and after irradiation

Table 8 shows the idle current before and after irradiation. This value represents the baseline idle current. No significant permanent increase in idle current was observed for any of the devices at the given LET. After the end of the irradiation, the idle current quickly drops to its pre-irradiation value (within seconds).

Table 8: current before and after irradiation at 80 °C

DUT	Current before irradiation	Current after irradiation
Samsung, 4 Gbit, sample 1	17.6 mA	17.6 mA

Samsung, 4 Gbit, sample 2	17.9 mA	17.9 mA
Hynix, 4 Gbit, sample 1	22.0 mA	21.4 mA
Hynix, 4 Gbit, sample 2	22.8 mA	22.9 mA
Hynix, 4 Gbit, sample 3	21.3 mA	21.3 mA

3.3.2 Current during irradiation

The current during irradiation is shown in figures 33 (Samsung devices) and 34 (Hynix devices). The dashed lines at the bottom represent the baseline idle current, as shown in table 8. Note that the flux – and therefore the irradiation time – was not identical for the five test runs.

For comparison, the baseline idle current for a Samsung device in the same test mode (write/read) is shown in figure 35 (including a twenty-second idle period at the beginning and at the end). This is the same data as shown in figure 32, but on a time scale that makes it easier to compare to the current during irradiation. The current is very similar for the Samsung and Hynix devices. Heating the device (while it is not being irradiated) does not significantly change the baseline current values.




Figure 33: current during irradiation 80 °C (4 Gbit Samsung devices, write/read mode, flux $2.5 \cdot 10^4 \text{ cm}^{-2}\text{s}^{-1}$)




Figure 34: current during irradiation at 80 °C (4 Gbit Hynix devices, write/read mode, flux $2.5 \cdot 10^4 \text{ cm}^{-2}\text{s}^{-1}$)




Figure 35: current without irradiation at room temperature (Samsung, 4 Gbit, write/read mode)

This was our first test campaign where we performed current measurement. While the sample rate was too low to analyze the current in detail, there are still some interesting findings: