



Technische Universität
Braunschweig



INSTITUT FÜR DATENTECHNIK
UND KOMMUNIKATIONSNETZE

TN-IDA-RAD-13/4

TID 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	Test facility	3
2.2	DUTs.....	3
2.3	Test device	4
2.3.1	Memory controller.....	5
2.3.2	Cooling	5
2.3.3	Shielding.....	6
2.4	Test sequence	7
3	Test procedures and test results	8
3.1	Mutual influence between DUTs	8
3.2	Error pattern	8
3.2.1	Precharge at test end.....	10
3.2.2	Precharge after periodic read.....	10
3.2.3	Test timing.....	11
3.2.4	Conclusion.....	11
3.3	Error count	11
3.3.1	DUT arrangement.....	11
3.3.1.1	Verification according to documentation/source code	12
3.3.1.2	Verification by measuring DM activity	13
3.3.1.3	Verification by error count	13
3.3.2	Influence of the error pattern.....	14
3.3.3	Influence of the DUT position.....	16
3.3.4	Influence of the column position.....	17
3.4	Idle current	19
3.5	Temperature dependence	19
3.5.1	Exposure to high temperature.....	20
3.6	Error annealing.....	20
4	Recommendations.....	23
5	Future work.....	23
6	References.....	24

1 Abstract

From October 22 to October 26, 2012, we performed a TID test campaign with DDR3 SDRAM at ESTEC, Noordwijk, Netherlands. This document reports on the findings.

2 Test setup

2.1 Test facility

The tests were performed at ESTEC's ^{60}Co source in Noordwijk, Netherlands.

2.2 DUTs

We tested three 4 Gbit devices from Samsung and Hynix, described in table 1. All of the devices were soldered to SODIMM modules, as shown in table 2. Other than that, none of the devices were prepared in any way, such as opening or thinning.

Table 1: tested parts (all 4 Gbit)

Designator	Manufacturer and part number		Lot code	Samples	Photo
A	Samsung	K4B4G0846B-HCH9	GMK3599Q	16	
B	Hynix	H5TQ4G83MFR-H9C	DTLB0241BH	8	
C	Hynix	H5TQ4G83MFR-H9C	DTLB0213HA	8	

Table 2: tested SODIMMs

Name	Manufacturer	Capacity	DUTs	Irradiation
Sam4SO1	Samsung	4 Gbit	8·A	Operated
Sam4SO2	Samsung	4 Gbit	8·A	Unbiased
Hyn4SO1	Hynix	4 Gbit	16 (8·B, 8·C)	Unbiased

2.3 Test device

The test bed, RTMC6 (figure 1), is capable of operating the first rank of one SODIMM with 8 DUTs in $\times 8$ or 4 DUTs in $\times 16$ 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.

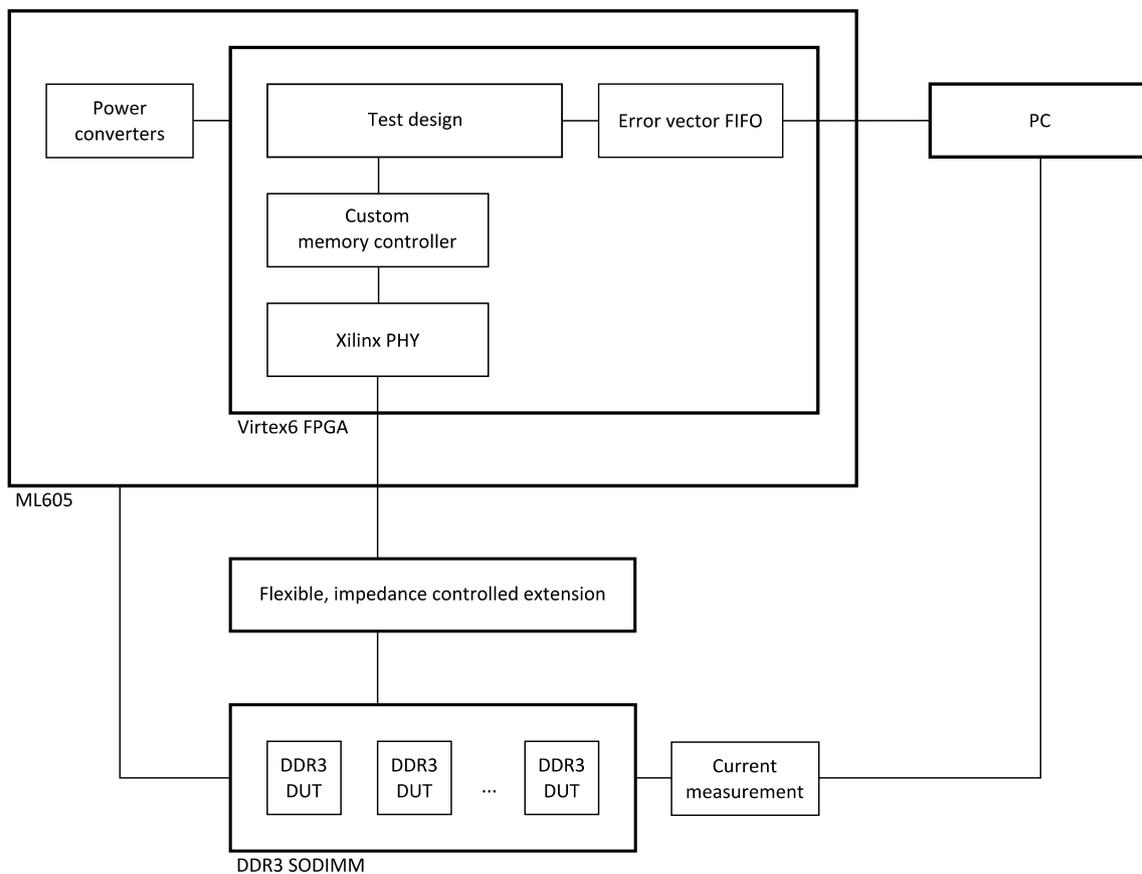


Figure 1: an overview of the RTMC6 test bench (simplified)

The FPGA contains a custom test design which writes a constant, counting or pseudo-random pattern to one of the DUTs, reads the data from the DUTs 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 DUTs have a higher data transfer rate than the USB connection, error vectors are buffered in an on-FIFO in order to be able to handle large runs of consecutive errors without slowing down the test. If the error record FIFO runs full due to too many errors, error vectors are either discarded or the test is slowed down, at the user's choice. In the latter case, when the FIFO runs full, the test is stalled until enough error vectors have been transmitted to make sufficient space in the buffer, and then resumed.

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

Due to the common command bus for all devices on an SODIMM, all commands are issued to all DUTs simultaneously and the DUTs generally operate in unison. Writing to one DUT selectively is achieved by using the *data mask* (DM) signal individual to each DUT. The associated activate and precharge commands are still performed by all DUTs simultaneously. Reading from one DUT selectively is not possible, apart from discarding the data from the other DUTs.

The total supply current for the whole SODIMM is measured at a sampling rate of 1 Hz and is logged by the PC. The supply current for individual DUTs cannot easily be measured as this would require a modification of the SODIMM.

2.3.1 Memory controller

An SDRAM controller consists of two parts: the memory controller proper (MC), which interfaces with the user logic, and the physical interface (PHY), which interfaces with the device. The MC is responsible for tracking the state of the DDR3 device and high-level timing. The PHY is responsible, among others, for low level (sub-clock-period) timing, data capturing and DDR translation. MC and PHY are connected via the DDR PHY interface (DFI).

Our test device uses a custom memory controller that has been developed specifically for memory tests. It provides fine-grained control over 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. Our controller implements an open page policy, which means that a row of the DRAM is kept open as long as possible after an access. In addition to simplifying the design, this policy has been shown to significantly increase DRAM performance [5]. It is particularly efficient for highly localized access patterns, as is the case with our memory test.

Our memory controller interfaces with the PHY developed by Xilinx and included in their Memory Interface Generator (MIG) package [3]. A peculiarity of the Xilinx DDR3 PHY is that it requires a read access (to no particular location) at least every microsecond in order to maintain internal timing parameters. If the user logic does not perform enough read operations, a read operation called “periodic read”, or PRD, is initiated by the MC. The controller uses the address that currently happens to be applied to its inputs, which is typically the last address accessed by the user logic.

2.3.2 Cooling

The FPGA and the power regulators dissipate several watts of heat. We devised a water cooler consisting of copper block placed above the board (see figure 2). The block is equipped with several threaded bolts to match the different heights of the various components to be cooled.

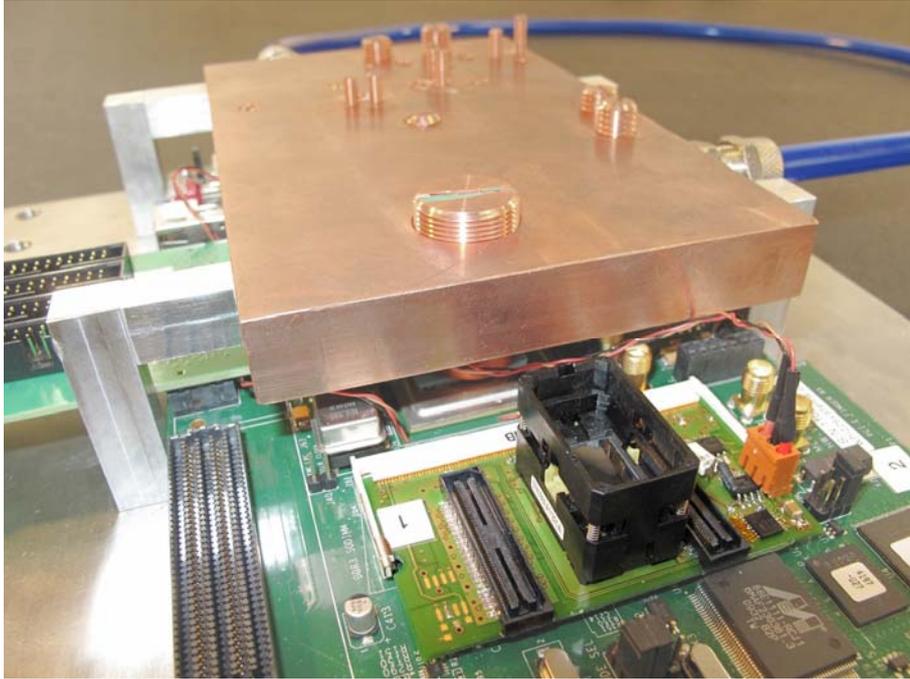


Figure 2: the water cooler is the same as in this SEE test setup

2.3.3 Shielding

In order to perform irradiation tests with a high total dose, all sensitive parts except the DUTs must be shielded from the radiation. For this purpose, we developed a shielding box (shown in figures 3 and 4) made of lead (where space is critical) and steel. The box weighs about 500 kg and is assembled from individual parts of about 10 kg each. It has several curved channels for feeding electrical wires and water tubes for the cooling system into the box.

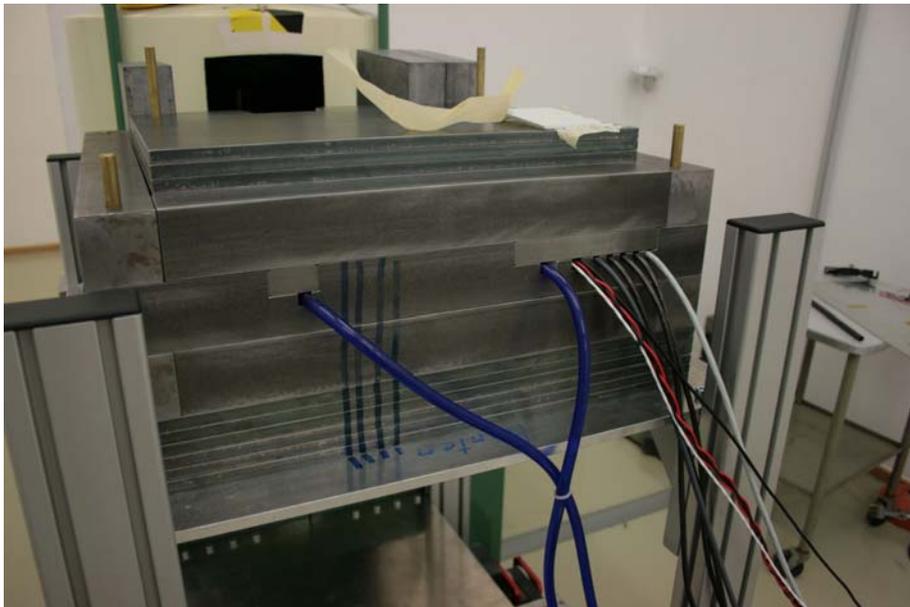


Figure 3: the shielding box with water tubes (blue) and electrical wires



Figure 4: our test device in the shielding box, with the top part of the shielding box removed

On the ML605, the SODIMM socket is located within centimeters of the FPGA and other sensitive devices. In order not to expose these devices to irradiation, we developed an impedance controlled extension consisting of a rigid PCB and a flexible part, shown in figure 5. The extension protrudes from the shielding box through a slit that is no wider than the thickness of the PCB. It also contains a shunt for current measurement.



Figure 5: the flexible extension with an SODIMM in the ML605 (with a fan on the FPGA instead of the water cooler)

A dosimeter placed inside the shielding box indicated a dose that was lower than the dose outside the box by a factor of $5 \cdot 10^3$. None of the devices inside the box failed during the entire irradiation with a total dose (outside the box) of more than 420 krad (silicon).

2.4 Test sequence

At the beginning of the test, a pseudo-random pattern was written to all DUTs. After that, the DUTs were tested in a round-robin fashion. Each DUT was first read. After a pause of 15

minutes, the DUT was written with the original pattern and the sequence repeated with the next DUT. This results in a total period of 120 minutes (2 hours) for each DUT, with 105 minutes between a write operation and the following read operation.

DDR3 devices experience significant self-heating during operation. Since all DUTs are operated in unison (as described in section 2.3), accessing any device has an influence on the temperature of other devices. The staggered operation of the DUTs serves to minimize the influence of this effect: were all DUTs to be read one after another, the last DUT would be read immediately following 7 other read operations, and therefore at a higher temperature than the first.

3 Test procedures and test results

3.1 Mutual influence between DUTs

During irradiation, we performed some manual tests. We found that, with our test device, reading from one device can cause errors in another device.

In particular, the test device was configured to not discard any errors vectors, but to slow down the test instead if the error vector FIFO runs full. We wrote a pattern to two devices and read the first one in order to verify that it did not have an unusually high number of errors. We then read another device, comparing the data against a different pattern than was written in order to simulate a very high number of errors in the device. When subsequently reading the first device again, it also contained a high number of errors. When reading only a part of the second device, the first device afterwards contained a high number of errors only in the same part of the address space. All of the error bits were changed from 1 to 0. A detailed analysis was not possible at this point due to shortcomings of the test device control software.

Slowing down the test causes the row that is currently open to remain open for a longer time than for regular test operation, up to the maximum row open time (t_{RAS} , [1]). As explained in section 2.3, this happens on all devices simultaneously. It has been hypothesized that active rows, or the sense amplifiers themselves, are more sensitive to radiation damage than the array, or that the process of activating a row involves other sensitive components within the device.

3.2 Error pattern

All of the Samsung DUTs (both on the operated and the unbiased DUT) exhibited a very peculiar set of error patterns, depending strongly on the mode of operation. The technical details of these devices (device A) are summarized in table 3.

Table 3: Technical details of the Samsung 4 Gbit device

Part number	K4B4G0846B-HCH9
Capacity	4 Gbit
Word size	8 bits
Number of banks	8
Number of rows	$2^{16} = 65536 = 0x10000$
Number of columns	$2^{10} = 1024 = 0x400$
Page size	1 kByte
Number of pages	$2^{19} = 524288 = 0x80000$

Detailed analysis of the error patterns was performed with the unbiased SODIMM (Sam4SO2) after the test campaign. The DUTs were heated to 95 °C, the maximum allowable temperature for the device, because the error patterns are more pronounced at higher temperature.

A detailed inspection of the error vectors reveals three distinct *error classes*, which can be recognized in figure 6:

1. randomly distributed errors
2. four regions, each 512 row wide, of errors
3. two single rows with a high number of errors close to the end of the device

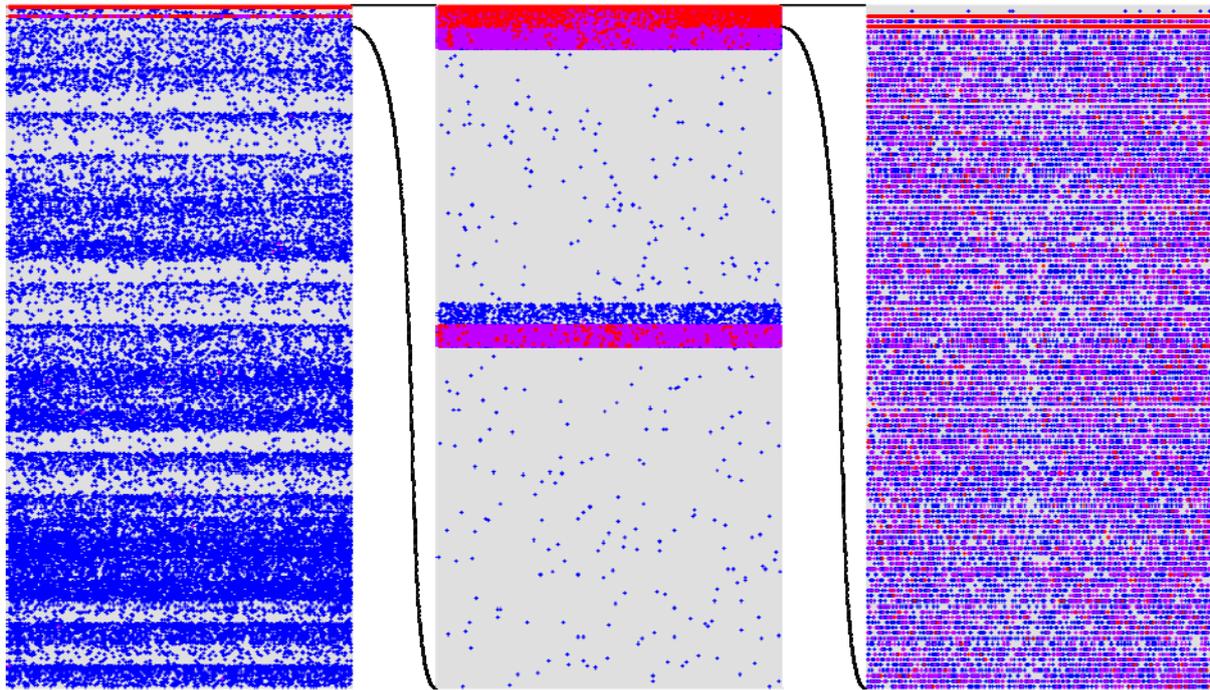


Figure 6: error map at different zoom levels (Samsung 4 Gbit, unbiased). The horizontal axis shows the column address, the vertical axis shows the concatenated bank and row address. The last error map shows individual pages; the white space in between does not represent pages without errors.

The first class consists of randomly distributed bit errors, the majority of which are single-bit errors. In figure 6, these errors are visible in the first and second error map. They exhibit some banding with a period of $2^{14} = 16384$ (0x4000) pages.

The second class consists of four *error regions* in the last bank, all 512 consecutive rows long:

- Row address 0xE000 to 0xE1FF
- Row address 0xE200 to 0xE3FF (the error density in this region is about 5% of the error density in the first region, making this region only visible for some of the DUTs and at high temperatures)
- Row address 0xFC00 to 0xFDFF (with the same error density as in the first region)
- Row address 0xFE00 to 0xFFFF (with an error density of about twice the error density in the first region, with the exceptions noted below)

For the “worst” DUT, that is, the one with the highest error count, about one percent of all bits were corrupted in the last error region. In figure 6, these errors can be recognized in the second error map.

The third class consists of the second and the fourth row from the end of the bank, row addresses 0xFFFC and 0xFFFE. In these rows, as many as 20% of all bits were corrupted. In figure 6, these rows can be recognized in the third error map.

Surprisingly, the very last row of the bank, row address 0xFFFF, contained very few errors, albeit more than would be expected from the random errors (class 1).

These error patterns (error classes 2 and 3) occur close to the end of the test, regardless of how many banks are tested: if only the first 4 banks are tested, the errors occur close to the end of bank 3. If all 8 banks are tested, the errors occur close to the end of bank 7 and the respective regions of bank 3 (and all other banks) show no particular error pattern (this is the case shown in figure 6). It has therefore been hypothesized that these errors are “caused by ending a test”. Note that except where otherwise noted, and to the best of our knowledge, the DUTs are always operated according to the specification [1]. The mode of operation is therefore not the root cause of the errors but rather triggers an error caused by radiation damage.

Writing all 8 banks, then writing the first 4 banks, and then reading all 8 banks caused the error pattern to appear both near the end of banks 3 and 7. This suggests that errors can be triggered by the end of the write operation.

Writing all 8 banks, then reading the first 4 banks, and then reading the first 4 banks a second time shows no errors during the first read, but shows the error pattern described above during the second read. This suggests that errors can also be triggered by the end of a read operation.

If the test is ended 128 rows before the end of the bank (testing only 0xFF80 rows of the bank), the row address of the four error regions (error class 2) remain the same as when testing the whole bank. This indicates that the errors do not only depend on the mode of operation but also on some property of the device. The last region, of course, is truncated to 384 rows in this case because the test is stopped before the end of the region is reached.

The second and fourth row from the end (error class 3), on the other hand, appear at row addresses 0xFF7C and 0xFF7E in this case, indicating that they depend solely on the mode of operation.

3.2.1 Precharge at test end

According to the open page policy implemented by the memory controller, the last row is not closed immediately by the memory controller after the test. The row is kept open until the next refresh operation, which requires all banks to be precharged. The mean time until that operation is half of the refresh interval, or 3.9 μ s.

In order to test the hypothesis that this is the cause for the errors near the end of the test, the test design was modified to explicitly precharge all banks at the end of the test. This does not seem to change the error patterns significantly (this has not been confirmed by detailed analysis).

The hypothesis that this is the cause for the errors is further made implausible by the fact that the same pattern should appear at the end of each bank, where the row is also kept open until the next refresh operation.

3.2.2 Precharge after periodic read

The periodic read operation required by the Xilinx PHY is performed every microsecond. The address currently applied to the memory controller is used, opening the row if required. Dur-

ing idle operation, the row is kept open until the next refresh operation, and then opened again when the next periodic read is scheduled. This leaves the row open for approximately $6.8 \mu\text{s}$ / $7.8 \mu\text{s}$, or about 78% of the time.

In order to test the hypothesis that this triggers the errors near the end of the test, the controller was modified to precharge the current bank immediately after the periodic read operation. This implements a closed page policy for periodic reads only. Note that this significantly slows down write operation because the row has to be re-opened in order to continue writing. Read operation is not affected because no periodic read operations are performed by the controller when enough read operations are initiated by the user logic.

This change makes the error regions (error class 2) disappear, or at least reduces the error density to a point where the regions can no longer be recognized. The second and fourth rows from the end (error class 3) remain, but the error density is greatly reduced.

The banding of the randomly distributed errors remains unchanged whether or not the device is precharged after periodic read.

3.2.3 Test timing

It has been hypothesized that the test logic slows down towards the end of the test, triggering the errors by keeping rows open longer than necessary. This hypothesis has been disproven by analyzing the DDR3 command bus with a logic analyzer.

3.2.4 Conclusion

The observed error patterns seem to depend in part on the mode of operation and in part on a property of the device.

It has been hypothesized that, due to bit line capacity issues, each bank of the DDR3 device does not contain one but several sense amplifiers, each associated with a subset of the rows in this bank, but only one of which is normally used at a time. This could have an influence on the error regions (error class 2), seeing that their extents seem to be related to some property of the device.

3.3 Error count

The DUTs on the Hynix SODIMM (Hyn4SO1) that were irradiated while unbiased showed no errors at all at room temperature. Additionally, this SODIMM was tested more thoroughly in a PC with the Memtest86+ software, also at room temperature. No errors were detected by this test either.

3.3.1 DUT arrangement

The physical arrangement of the DUTs on the SODIMM can be relevant. The used Samsung SODIMMs follow the “raw card version B2” layout from [2]. The mapping from device position to DUT number is shown in figure 7, along with a device label. The terms “front side” and “back” side are in accordance with the standard [2].

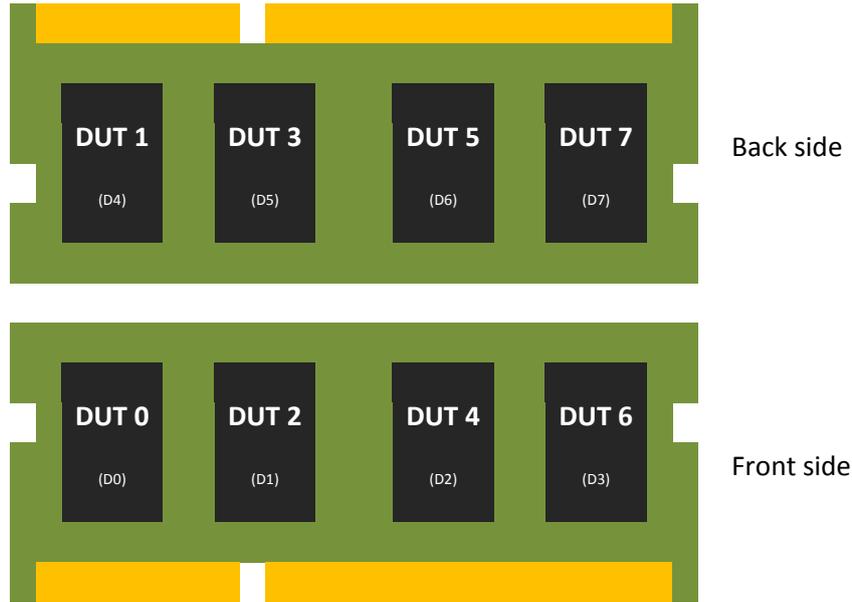


Figure 7: SODIMM DUT positions and device labels (in parentheses)

3.3.1.1 Verification according to documentation/source code

The device position can be mapped to the DUT number according to table 4. The individual mappings follow from the following sources:

- Device position to device label: from [2], figure 2 (page 4.20.18-13), right
- Device label to DM signal SODIMM net name: from [2], figure 2 (page 4.20.18-13), left
- DM signal SODIMM net name to DM signal SODIMM pin number: from [2], table 6 (page 4.20.18-11)
- DM signal SODIMM pin number to DM signal ML605 net name: from [4]
- DM signal ML605 net name to DM signal FPGA pin: from [4]
- DM signal FPGA pin to DM signal VHDL signal name: from the RTMC6 FPGA source code, `sodimm_standard.ucf`
- DM signal VHDL signal name to DUT number: RTMC6 FPGA and UI source code

Table 4: device label to DUT number mapping

Device position	Device label	DM signal SODIMM net name	DM signal SODIMM pin number	DM signal ML605 net name	DM signal FPGA pin	DM signal VHDL signal name	DUT number
See figure 7	D0	DM0	11	DDR3_DM0	E11	ddr3_dm[0]	0
	D1	DM2	46	DDR3_DM2	E14	ddr3_dm[2]	2
	D2	DM4	136	DDR3_DM4	B22	ddr3_dm[4]	4
	D3	DM6	170	DDR3_DM6	A29	ddr3_dm[6]	6
	D4	DM1	28	DDR3_DM1	B11	ddr3_dm[1]	1
	D5	DM3	63	DDR3_DM3	D19	ddr3_dm[3]	3
	D6	DM5	153	DDR3_DM5	A29	ddr3_dm[5]	5
	D7	DM7	187	DDR3_DM7	A31	ddr3_dm[7]	7

3.3.1.2 Verification by measuring DM activity

The mapping outlined in section 3.3.1.1 was verified in two steps.

The mapping from the device position to the DM signal SODIMM pin number was verified by measuring the connection between the SODIMM pins and the pads for the DM signals of the individual DUTs on an unpopulated SODIMM board using a continuity tester.

The mapping from the DM signal SODIMM pin number to the DUT number was verified by writing to one DUT at a time and monitoring the DM signal SODIMM pins using an oscilloscope.

3.3.1.3 Verification by error count

The mapping was further verified directly in one single step.

It is well known that DRAM data retention with disabled refresh depends on the temperature of the device [6]. The whole SODIMM was therefore heated to an estimated temperature of 50 °C using a 50 W halogen lamp. One of the DUTs was cooled to about 10 °C using coolant spray. A pattern was written to all DUTs, the refresh disabled for 16 seconds, and all DUTs read.

The result for cooling DUT 1 is shown in figure 8. DUT 1 has the lowest number of errors by far. The next higher error count is found in DUT 0 (opposite from DUT 1), DUT3 (adjacent to DUT 1) and DUT 2 (opposite from DUT3). The other DUTs are sufficiently far away from the cooled DUT to be unaffected.

The process was repeated for all DUTs, verifying the mapping between the device position on the SODIMM and the DUT number.

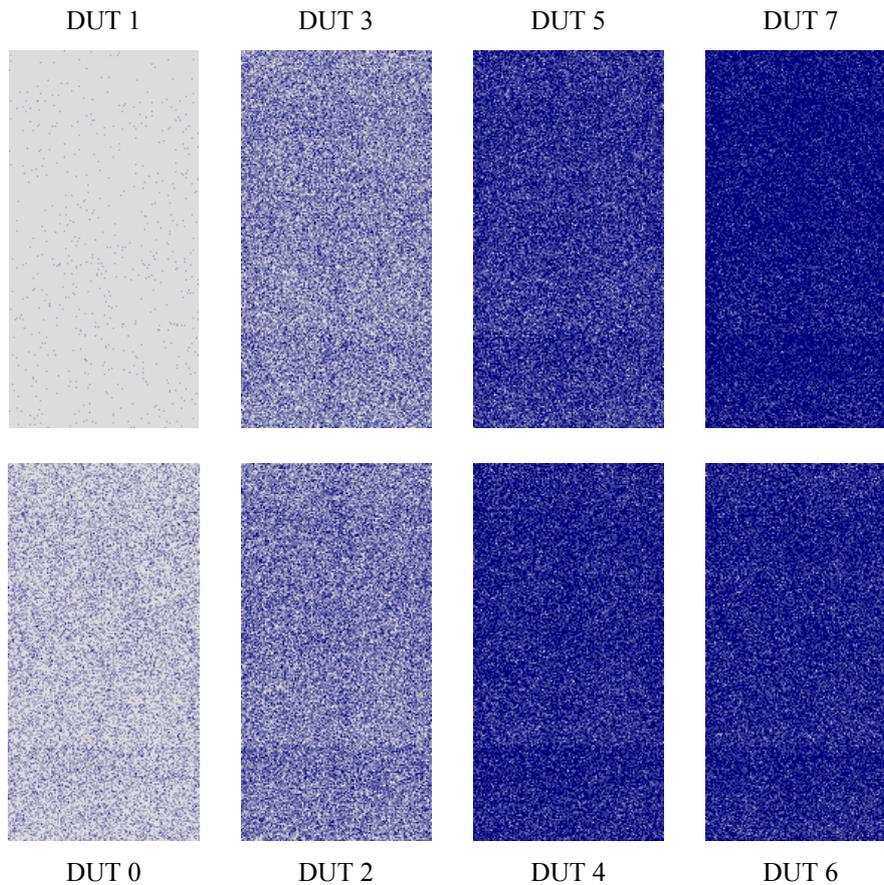


Figure 8: error maps of 8 DUTs after 16 seconds without refresh. DUT 1 was cooled.

3.3.2 Influence of the error pattern

Figure 9 shows the cell error ratio (also called *error density*) versus the dose for all DUTs of the SODIMM that was operated during irradiation (Sam4SO1). The cell error ratio is the number of bit errors in a region of the device, divided by the total number of bits in that region. The gap at 350 krad (silicon) represents the time when the manual tests mentioned in section 3.1 were performed.

The values in figure 9 include the error regions described in section 3.2. Since these errors are believed to be, at least in part, an artifact of the mode of operation, we are also interested in the cell error ratio excluding these errors. Furthermore, the error regions are large enough to cause overflows of the error vector FIFO in some cases, making the total bit error count unreliable.

Since the aforementioned error regions always occur near the end of the device (for the relevant test runs), we can filter them out by only considering the first half of the device. The cell error ratio for this case is shown in figure 10.

Figure 11 compares the unfiltered (from figure 9) with the filtered cell error ratio (from figure 10) for some selected DUTs. The influence of the error pattern is large low dose. As the dose increases, the influence of the error pattern decreases as the number of errors in the rest of the device increases.

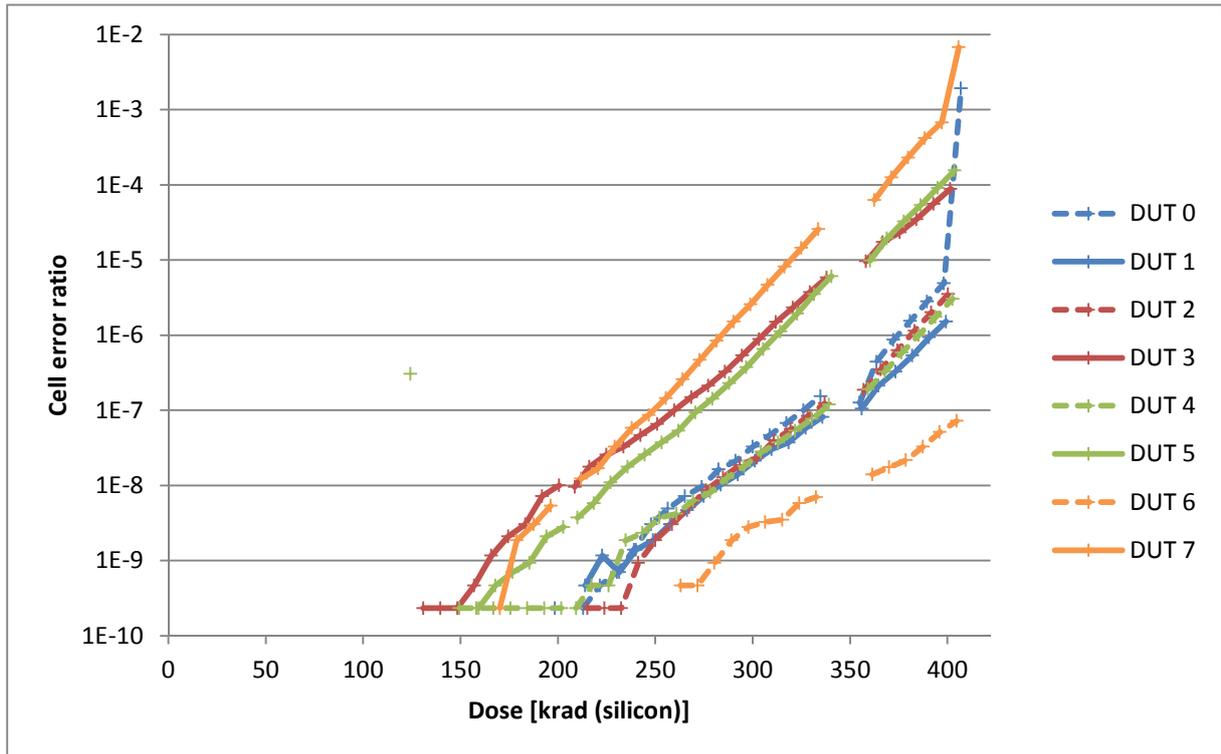


Figure 9: errors vs. dose, Samsung 4 Gbit, whole device

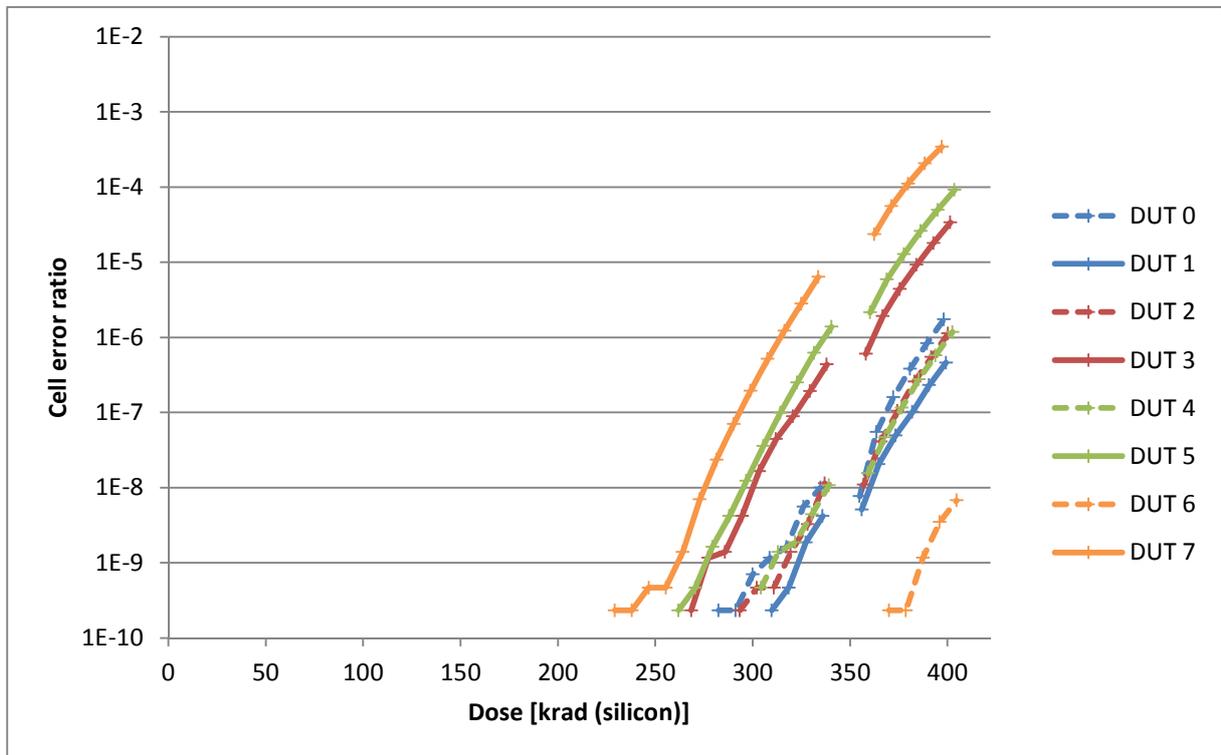


Figure 10: errors vs. dose, Samsung 4 Gbit, first half of device

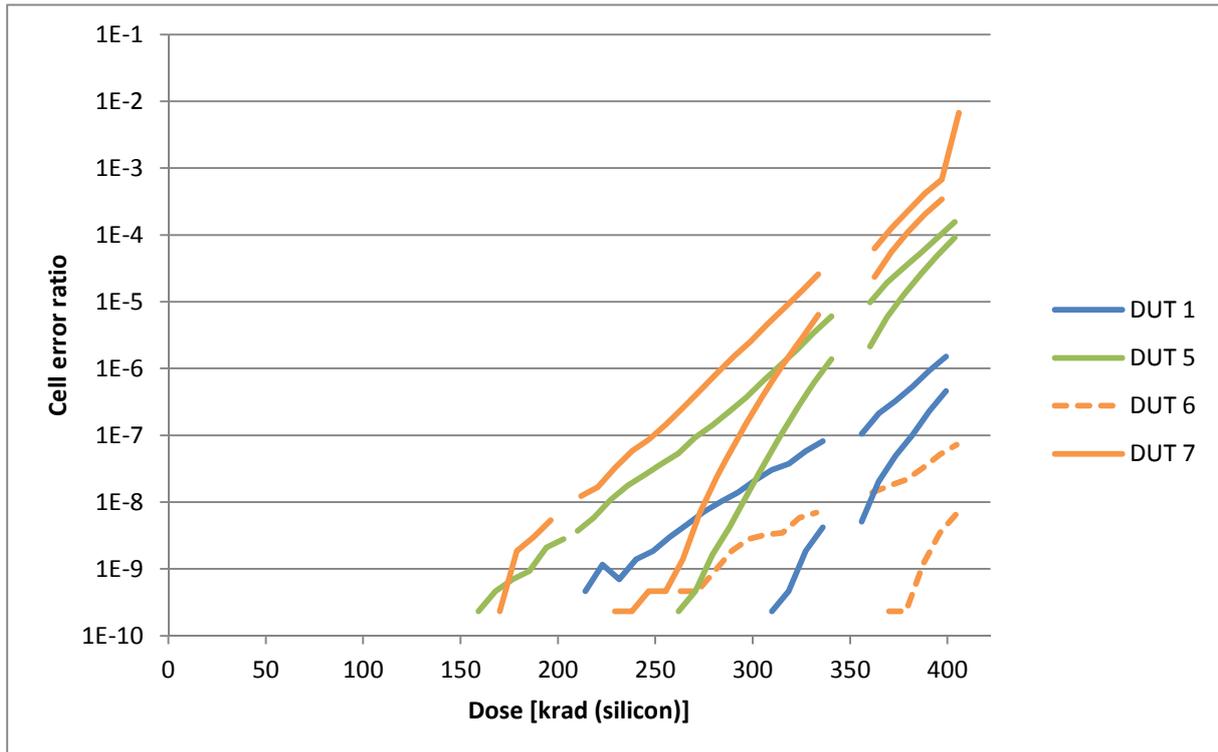


Figure 11: errors vs. dose, Samsung 4 Gbit, comparison of whole device (higher values, from figure 9) with first half of device (lower values, from figure 10) for selected DUTs

3.3.3 Influence of the DUT position

The cell error ratio shown in figure 10 varies strongly between the DUTs. DUTs on the back side of the SODIMM (facing the source) are shown with a solid line. DUTs on the front side of the SODIMM (facing away from the line) are shown with a dashed line. DUTs opposite from each other are shown in the same color. Figures 12 and 13 show the same data just for devices on the side facing the source and facing away from the source, respectively.

DUTs on the side facing the source have much more errors than DUTs on the side facing away from the source. This cannot be explained by γ radiation, which should not significantly be shielded by the SODIMM board or the DUTs. It is suspected that the difference might be caused by electrons, either from the β decay of ^{60}Co or from photon scattering in the test chamber.

It has also been hypothesized that the DUTs on one side of the SODIMM are, in fact, different from the DUTs on the other side. For example, this could be the case because during production of the SODIMM, the devices on the side that was soldered first were heated a second time when the devices on the other side were soldered.

Even on the same side of the SODIMM, the sensitivity of the four DUTs shows significant variation. This cannot be explained by non-uniform irradiation since the DUT with the highest error number on the front side (DUT 7) is opposite from the DUT with the lowest error number on the back side (DUT 6), and vice versa (DUTs 1 and 0).

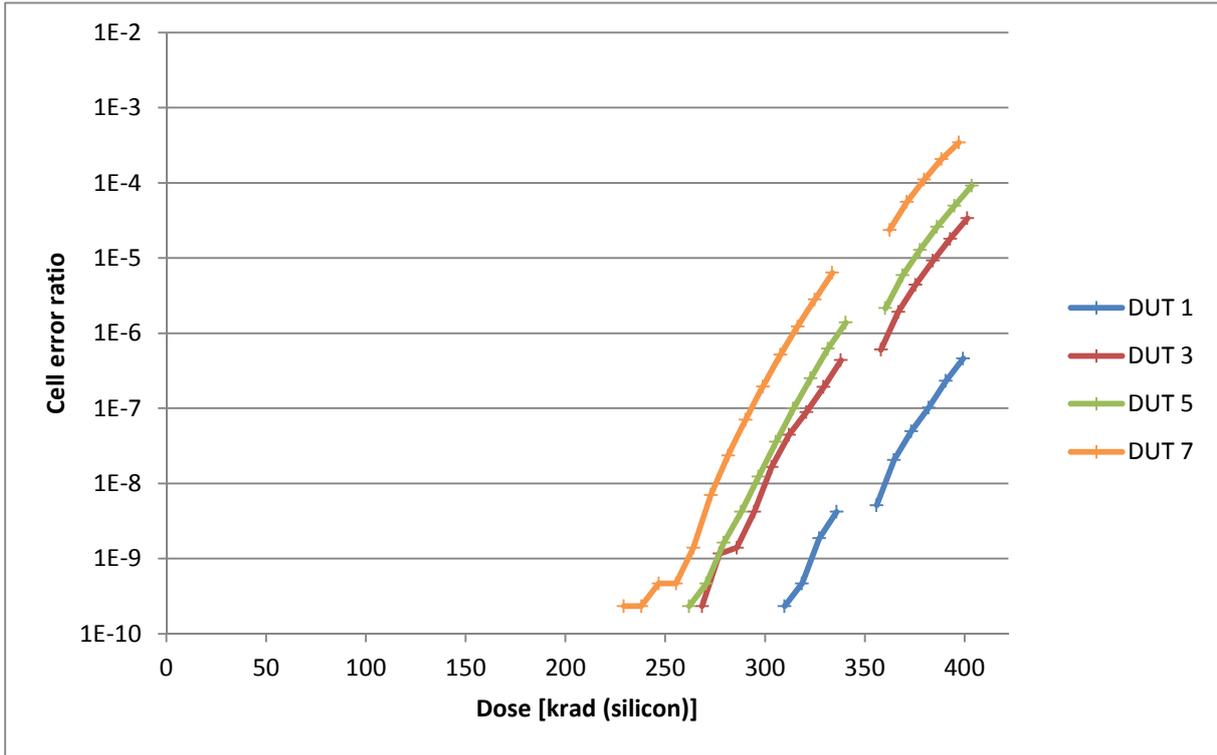


Figure 12: errors vs. dose, Samsung 4 Gbit, DUTs facing the source

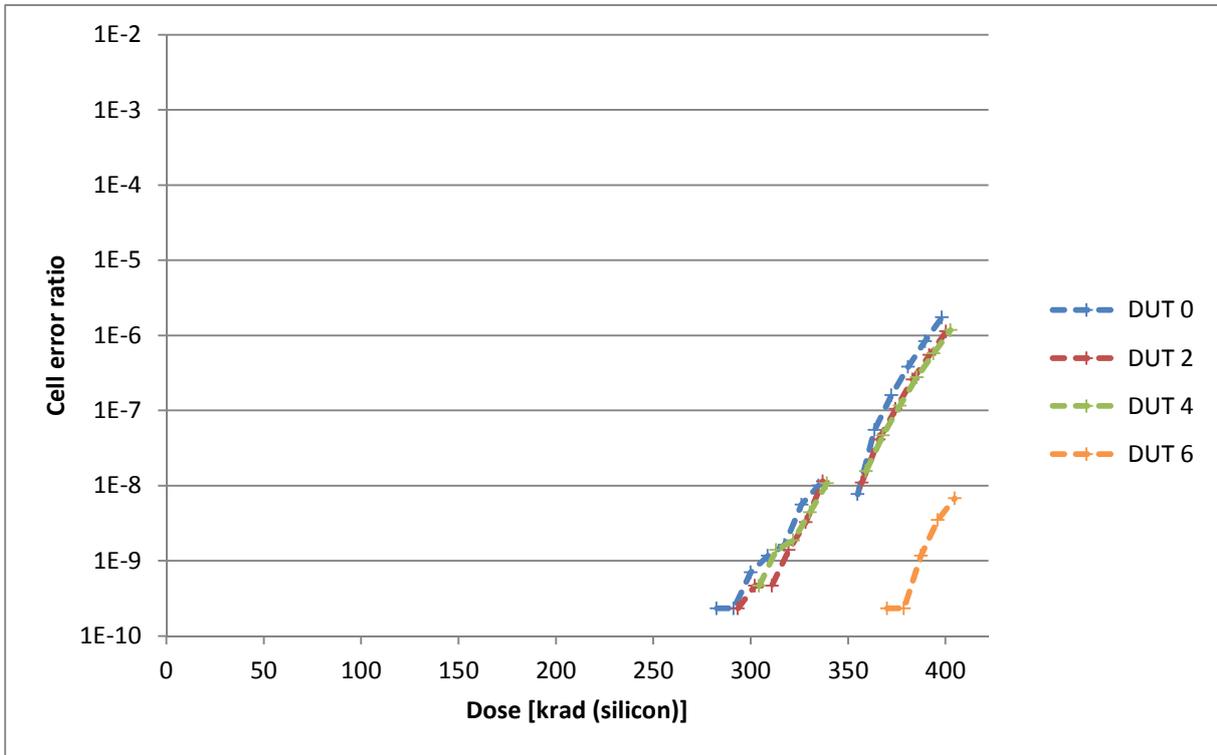


Figure 13: errors vs. dose, Samsung 4 Gbit, DUTs facing away from the source

3.3.4 Influence of the column position

The number of errors in the middle of a device is slightly lower than in the columns towards the edges of the device. An example for three devices is shown in figure 14. More detail can be recognized when plotting the data on a linear scale, shown for one DUT in figure 15. Note in particular the steep rise of the error count towards column 0.

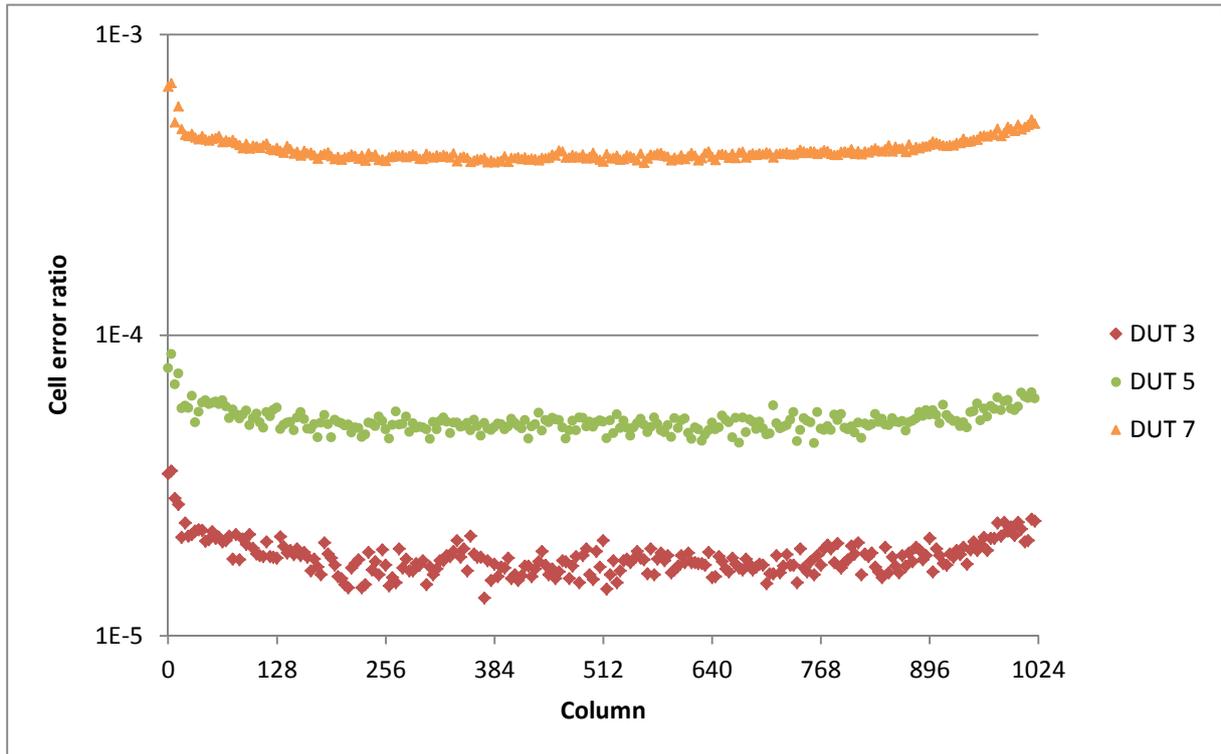


Figure 14: error count vs. column position, Samsung 4 Gbit, operated, ≈ 390 krad (silicon), logarithmic scale

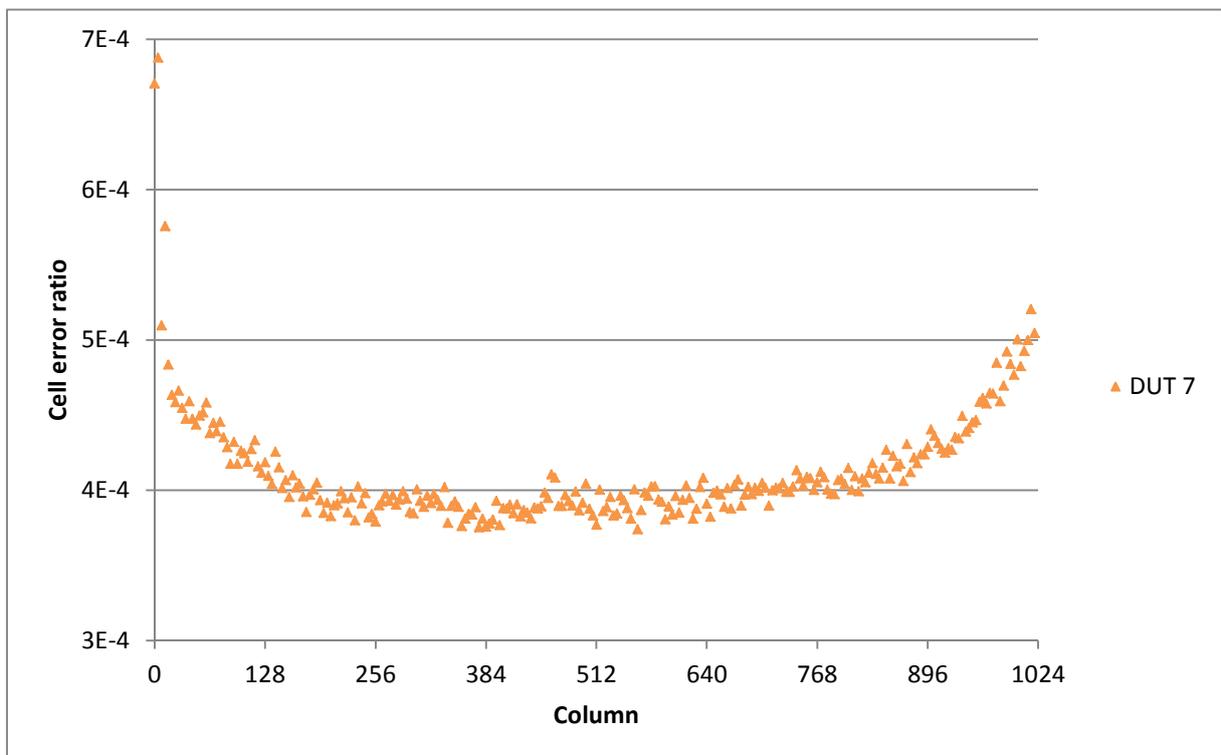


Figure 15: error count vs. column position, Samsung 4 Gbit, operated, ≈ 390 krad (silicon), linear scale

3.4 Idle current

During each loop, the idle current was measured and logged automatically. This measurement was performed as the first action after the 15 minute wait time, in order to avoid the warming caused by the write operation to affect the idle current.

Additionally, the current was logged manually about every hour during the time where the test was attended, even at times where the regular testing procedure was not followed.

As described in section 2.3, the individual current for each DUT could not be measured. The measured current is the total current for the whole SODIMM. Dividing the current by the number of DUTs on the SODIMM yields the average current per DUT.

The current vs. dose is shown in figure 16. Over the course of ≈ 420 krad (silicon), the average idle current increased by less than 25% from 15.7 mA to 19.3 mA.

Some of the automatically measured values are slightly increased over the nearby values. This may be due to the measurement coinciding with the refresh or periodic read operation.

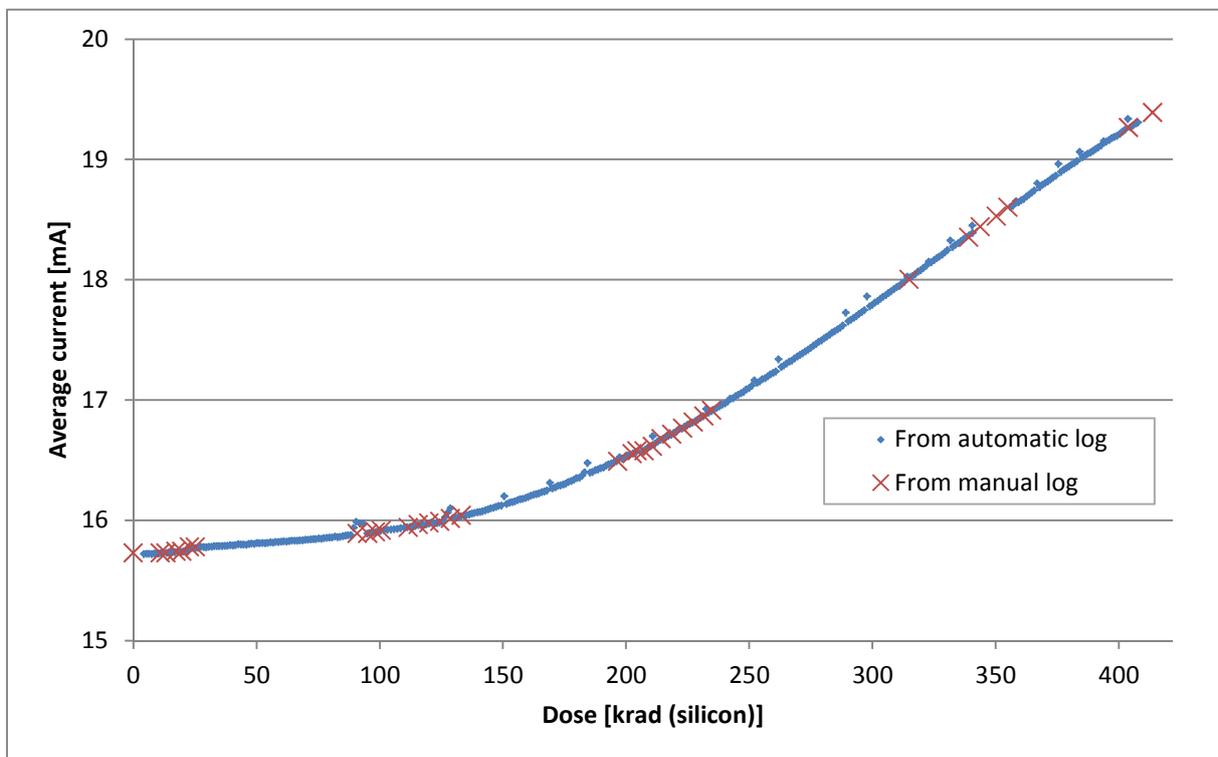


Figure 16: idle current vs. dose, Samsung 4 Gbit

3.5 Temperature dependence

After the test, the unbiased SODIMM, Sam4SO2, was tested for the influence of the temperature on the number of errors by writing a pseudo-random pattern to the device and reading it back about one second later.

Figure 17 shows the total number of bit errors n in the whole device vs. the temperature T . This relationship is described approximately by $n \propto e^{T/12K}$. The number of errors thus doubles every 8 K or increases by a factor of 10 every 28 K in temperature change.

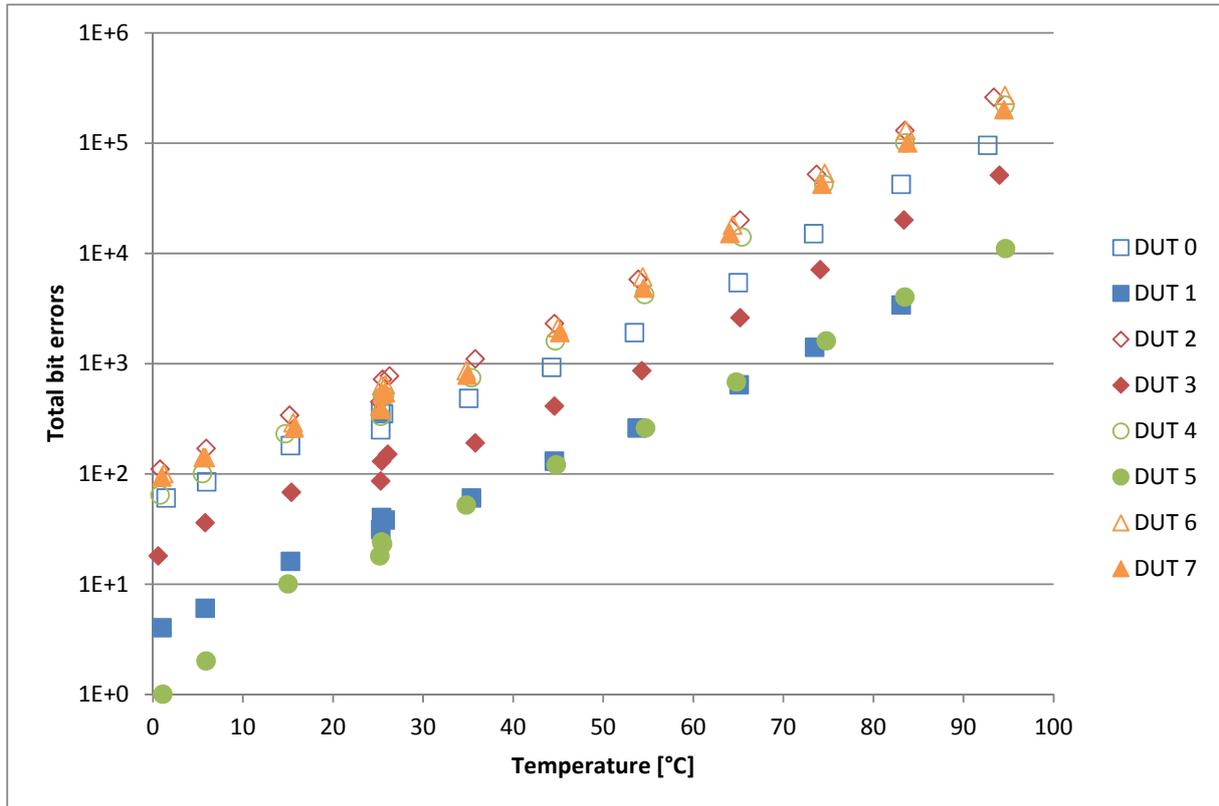


Figure 17: bit errors vs. temperature, Samsung 4 Gbit, unbiased, ≈ 420 krad (silicon), refresh rate $7.8 \mu\text{s}$

The first measurement was performed at 25°C . After the device had been heated to 95°C and repeatedly written and read over the course of approximately 2 hours, another measurement was performed at 25°C . Compared to the first measurement at 25°C , the error count was increased by a factor of approximately 1.5 (a factor that, were it caused by temperature, would correspond to a change in temperature of about 5 K). After cooling the device to 0°C and heating it back up to 25°C , the increased number of errors remained, confirming that the increased error count was not caused by temperature hysteresis.

The refresh interval was kept at the default of $7.8 \mu\text{s}$ for all tests, even at $T > 85^\circ\text{C}$, where the standard requires the refresh rate to be doubled [1].

3.5.1 Exposure to high temperature

In an attempt to examine individual devices, 4 of the 8 DUTs were desoldered from the SODIMM board. In the process, they were exposed to a temperature of approximately 250°C for about 10 seconds. After reballing, 3 of the devices were tested at 80°C . None of the errors described in section 3.2 could be observed for any of the devices. This is assumed to be due to annealing at high temperature. One of the devices showed a new type of error with column errors, which was not examined in detail.

3.6 Error annealing

In the days after the test, the number of errors remaining in all tested Samsung DUTs (Sam4SO1 and Sam4SO2) was repeatedly measured. The results are shown in figures 18 to 22. The values in figures 20 and 21 are normalized to the first measurement in the days after the test. The values in figure Figure 22 are normalized to the values immediately after irradiation. For the unbiased SODIMM (Sam4SO2), this data is not available.

These measurements were performed at room temperature which was neither controlled nor measured. Additionally, self-heating of the devices was not considered. Since the temperature has been shown to have a significant influence on the number of errors (see section 3.5), the data must be considered unreliable.

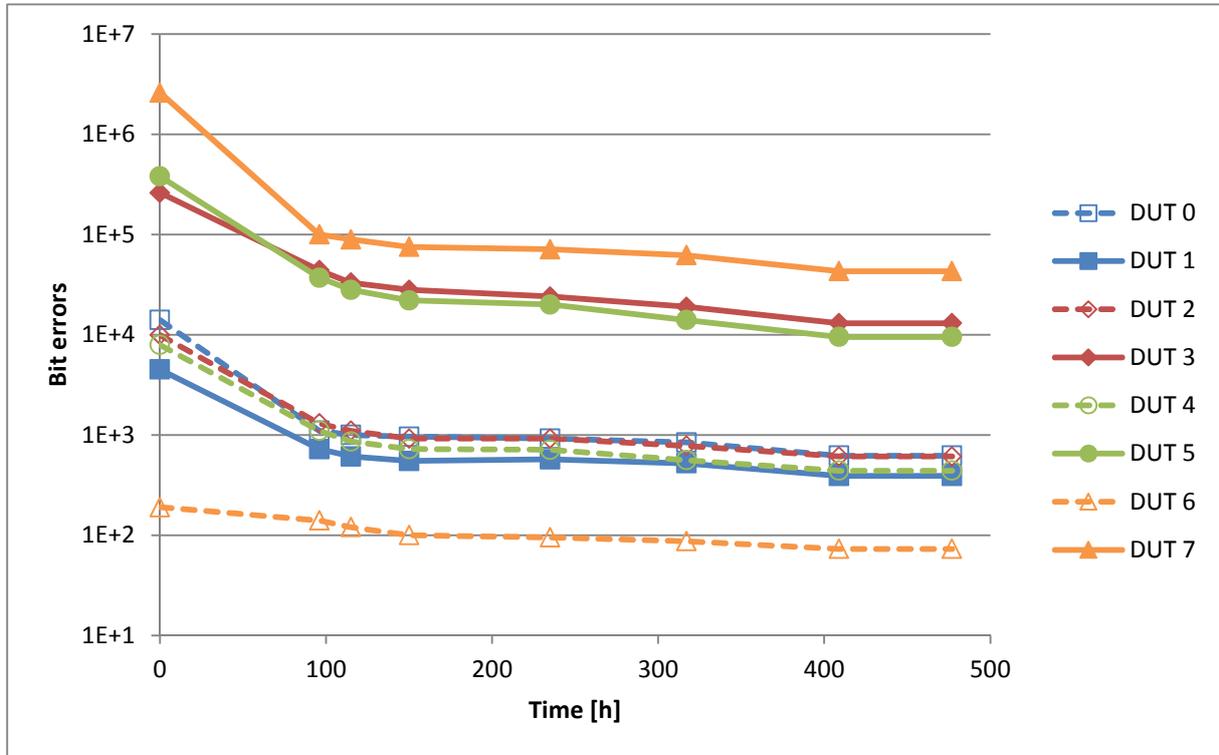


Figure 18: error annealing, Samsung 4 Gbit, operated, raw

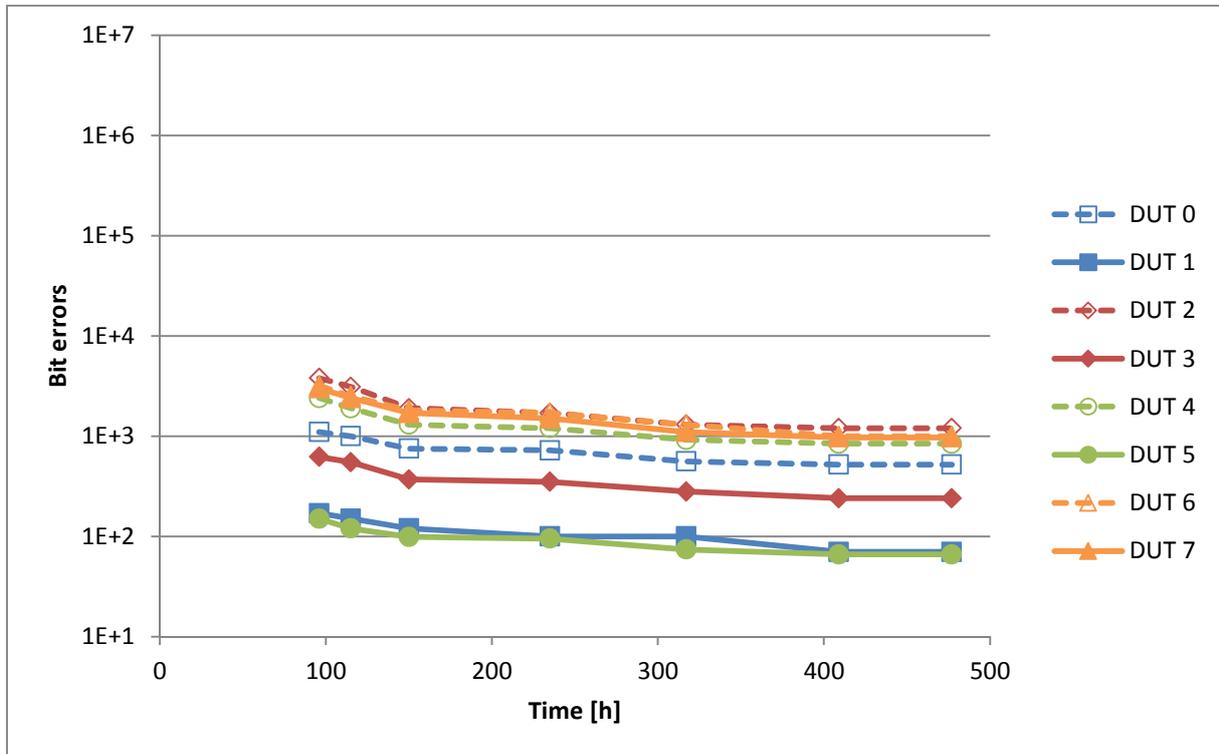


Figure 19: error annealing, Samsung 4 Gbit, unbiased, raw

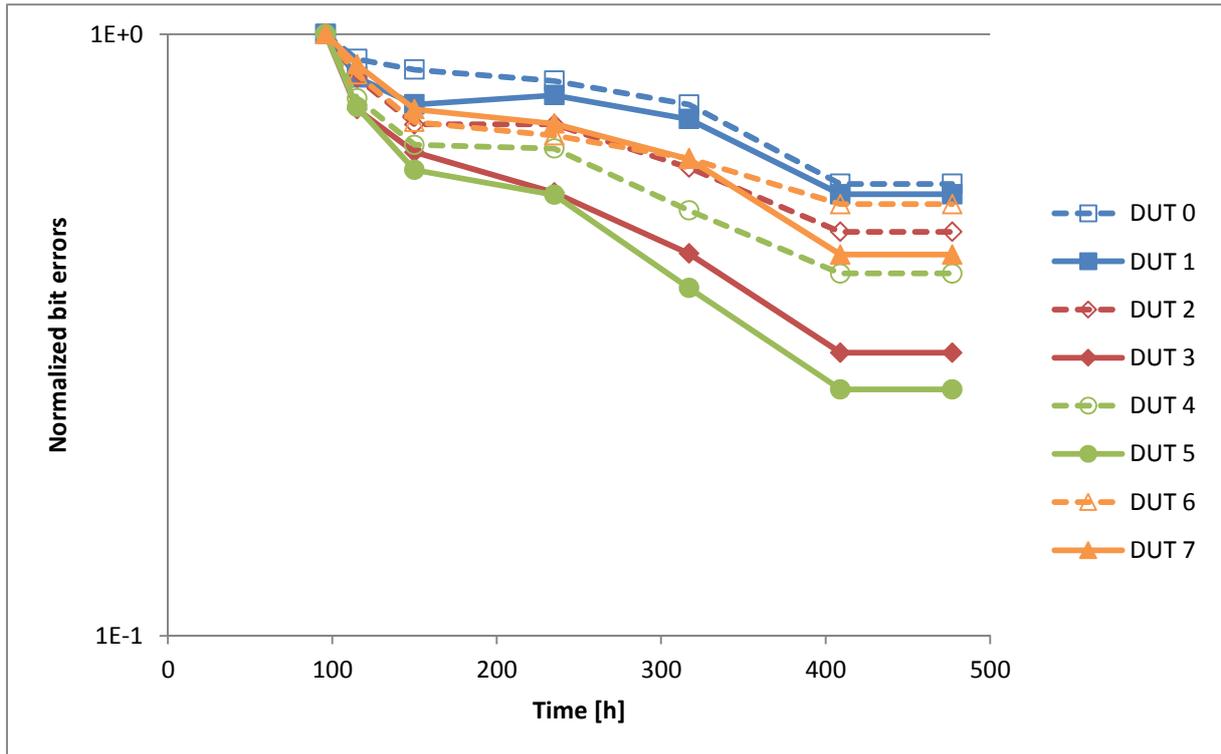


Figure 20: error annealing, Samsung 4 Gbit, operated, normalized to value after ≈100 hours

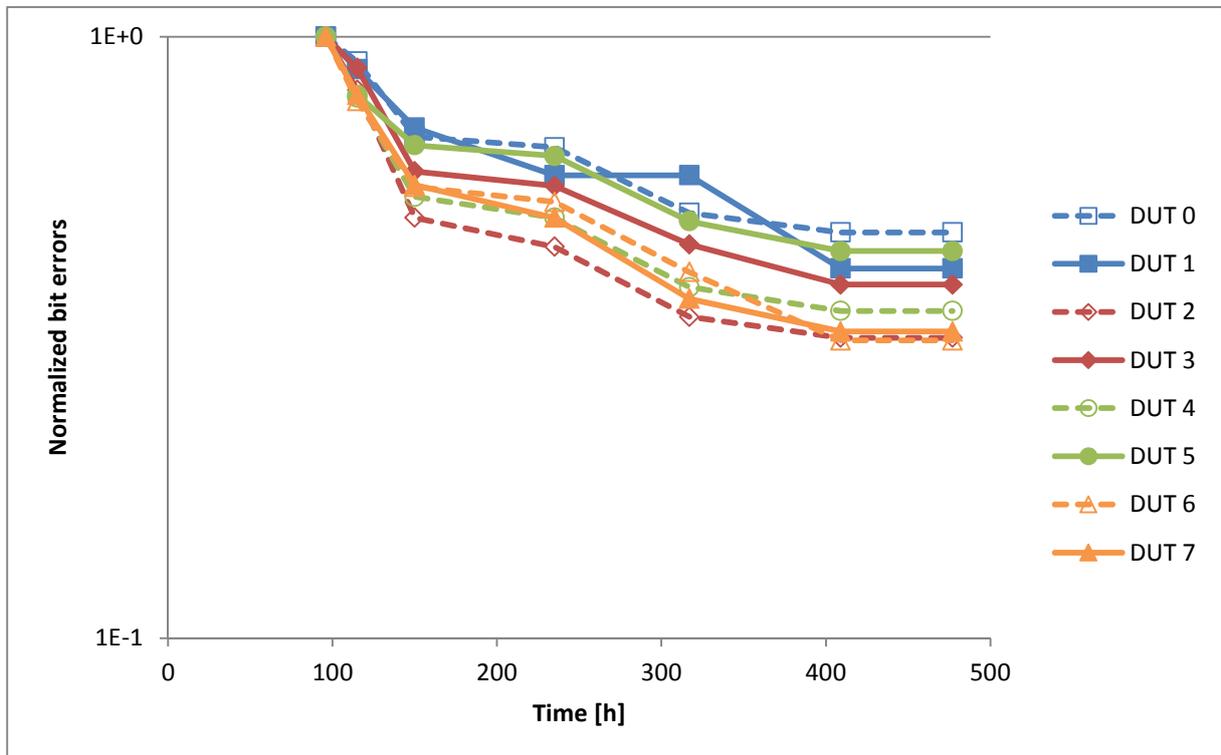


Figure 21: error annealing, Samsung 4 Gbit, unbiased, normalized to value after ≈100 hours

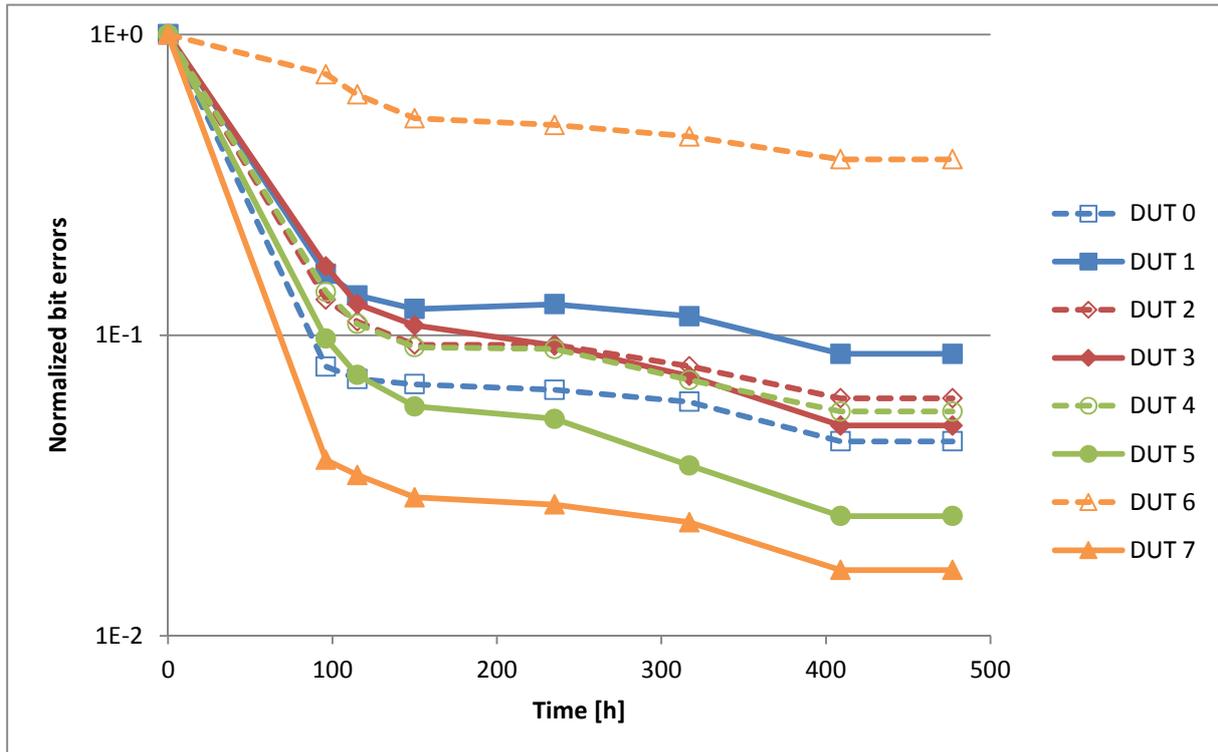


Figure 22: error annealing, Samsung 4 Gbit, operated, normalized to values right after irradiation

4 Recommendations

Pending further investigation of the error pattern described in section 3.2, when operating a DDR3 memory device in a radiation environment, it is recommended not to keep rows open for an extended period of time. If the performance impact is acceptable, a closed page policy should be implemented by the controller, or banks should be explicitly precharged by the user logic as soon as they are no longer required active. It may be desirable to use a PHY which does not require periodic read operations.

5 Future work

As described in section 3.2, the observed error patterns depend both on the mode of operation and on some property of the device. There are several tests that can be performed in order to further investigate the issue. These tests include modifying the mode of operation, which is relatively easy, and verifying the test device itself. Not all tests may be necessary or useful, depending on the results of earlier tests.

- End a test at different points within a bank to see if the error regions (error class 2) also appear at these positions
- End a test after an even row address (that is, test an odd number of rows) to see if the error rows (error class 3) still appears at the second and the fourth row from the end or if the locations are also influenced by a property of the device.
- Write the whole device, then read only one or a few rows close to the end of bank 3, and then read the whole of bank 3 in order to find out the conditions under which the error regions (error class 2) appear

- Instead of reading and/or writing, just activate and precharge the rows in order to determine whether this already causes the error. Vary the time between opening and closing the row, from the minimum to the maximum possible time.
- Explicitly specify an address to use for the periodic read operations
- Perform only periodic read operations to determine the influence of the periodic read on the error pattern
- Use a different test data pattern to find patterns
- Begin a test at a different row address
- Disable periodic read completely, provided that the Xilinx PHY still works correctly in this case
- Verify by detailed analysis that precharging all banks at the end of the test does not change the error patterns at all
- Examine the supply current with an oscilloscope to find unusual current spikes, for example, if several sense amplifiers of a bank are activated simultaneously (contrary to regular operation)
- Examine the DDR3 data bus using a logic analyzer to find out if the errors are caused by wrong timing of the device
- Examine the complete DDR3 interface using a protocol analyzer to verify that the device is operated according to the specification. No such protocol analyzer is available to us at this time.
- Check the Hynix DUTs for similar errors at high temperature

Concerning the mutual influence (section 3.1), the test sequence described in section 2.4 should be modified to not access any other DUT between a write operation and the following read operation of the same DUT.

Concerning the discrepancy between DUTs (section 3.3.3), the DUTs should be shielded, for example with a thin (e. g. 2 mm) layer of aluminum, to decrease the effect of electrons. If the difference between the front and back sides of the SODIMM persists, it might be desirable to rotate the SODIMM by 180 degrees at some time during the test in order to determine whether the difference between the sides is caused by the test setup or is related to some property of the SODIMM, as suggested in section 3.3.3.

Concerning the influence of temperature (section 3.5), additional temperature cycles, heating the device to 95 °C and cooling it back down to 25 °C, could be performed in order to see if the error count increases further.

6 References

- [1] JEDEC standard 79-3E – *DDR3 SDRAM Specification*
- [2] JEDEC standard 21C – 204-Pin DDR3 SDRAM Unbuffered SO-DIMM Design Specification
- [3] Xilinx, UG406 – Virtex-6 FPGA Memory Interface Solutions User Guide
- [4] Xilinx, XTP052 – ML605 schematics
- [5] J. A. a. J. Niittylahti, *A Comparison of Precharge Policies with Modern DRAM*, in *9th International Conference on Electronics, Circuits and Systems*, 2002

- [6] J. Alex Halderman et al., *Lest We Remember: Cold Boot Attacks on Encryption Keys*