Total Dose Test Report for Samsung 8G NAND Flash Nonvolatile Memory

Timothy R. Oldham[1], Mark Friendlich[2] , and Steve Cox[3]

1. QSS Group, Inc., Seabrook, MD 20706
2. Muniz Engineering, Inc., Seabrook, MD 20706
3. NASA Goddard Space Flight Center, Greenbelt, MD 20771

Test Date: 6-13 March 2008 Report Date: 5 May 2008

I. Introduction

The purpose of this test was to determine the susceptibility to total ionizing radiation dose (TID) of the Samsung 8G NAND flash nonvolatile memory (part number K9F8G08U0A, Lot Date Code 716). This test was supported by the NASA Electronics Parts and Packaging (NEPP) Program.

II. Devices Tested

The Samsung NAND Flash Memory is a non-volatile memory that uses a floating gate NAND cell, implemented in 60 nm technology. It also provides a standard interface for pin and functional drop-in compatibility. We believe these parts were burned-in before leaving the factory, so it is not possible to do a controlled experiment to look at burn-in effects. In any case, there is no plan to do our own burn-in. Detailed device information is provided in Table I. The parts have 8K blocks, a few of which can be "bad," as identified by the manufacturer. The blocks are 128Kx8, with 64 pages, each 2Kx8. In this case, five samples were irradiated, all of which had some bad blocks. There was also one unirradiated control device. The parts have a nominal 3.3 V power supply, plus an internal charge pump to generate higher voltages for writing and erasing. These parts come with the capability for a 4-bit ECC (error correcting code), which was not used during the test. But it could possibly have corrected some of the single bit errors.

Generic Part Number:	
Full Part Number	K9F8G08U0A
Manufacturer:	Samsung
Lot Date Code (LDC):	716
Quantity Tested:	5
Serial Numbers of Control Sample:	6
Serial Numbers of Radiation Samples:	1,2, 3, 4, 5
Part Function:	NAND Flash Memory
Part Technology:	CMOS
Case Markings:	Samsung 716 K9F8G08U0A PCB0 FHB025BX
Package Style:	48 pin TSOP
Test Equipment:	Power Supply (+3.3V) Digital test board. Multimeters
Test Engineer:	M. Friendlich
Dose Levels (krad (Si)):	10, 20, 30, 50, 75, and 100krads(Si) continuing in 50krads (Si) steps until functional failure.
Target dose rate (rad (Si)/min):	1200-1800

Table I. Device information

III. Test Facility

Testing was at the Co-60 facility at GSFC, which is a room air source, where the pencils are raised up out of the floor, during exposures. Active dosimetry is performed, using air ionization probes. Testing is done in a step/stress manner, using a standard Pb/Al filter box. Dose rate typically varies slightly from one exposure to the next, up to 30 rads/s. Most exposures are near the maximum dose rate, as required by MIL-STD Test Method 1019.6. Time intervals for testing between exposures are also within the limits stated in 1019.6 (one hour after exposure to start electrical characterization, two hours to begin the next exposure). Parts were under DC bias during exposures, but not actively exercised.

IV. Test Procedure

The test devices were programmed with a checkerboard pattern (AA) during exposures, and biased at 3.6 V (3.3 V nominal power supply, plus 10%), but the devices were not actively exercised during exposures. Two parts were read (only) between exposures, to look for problems related to the integrity of the individual bits. The other three parts were exercised between exposures—read, erased, and written into four different patterns. The patterns were checkerboard (AA), checkerboard complement (55), all ones, and all zeroes. In each of these tests, the entire memory is read, or erased, or programmed in one operation, with the commands entered manually. There is also a dynamic test mode, where each block is read, erased, and programmed, then the next block, and so on until the entire memory is completed. A block diagram of the test apparatus is shown in Fig. 1.

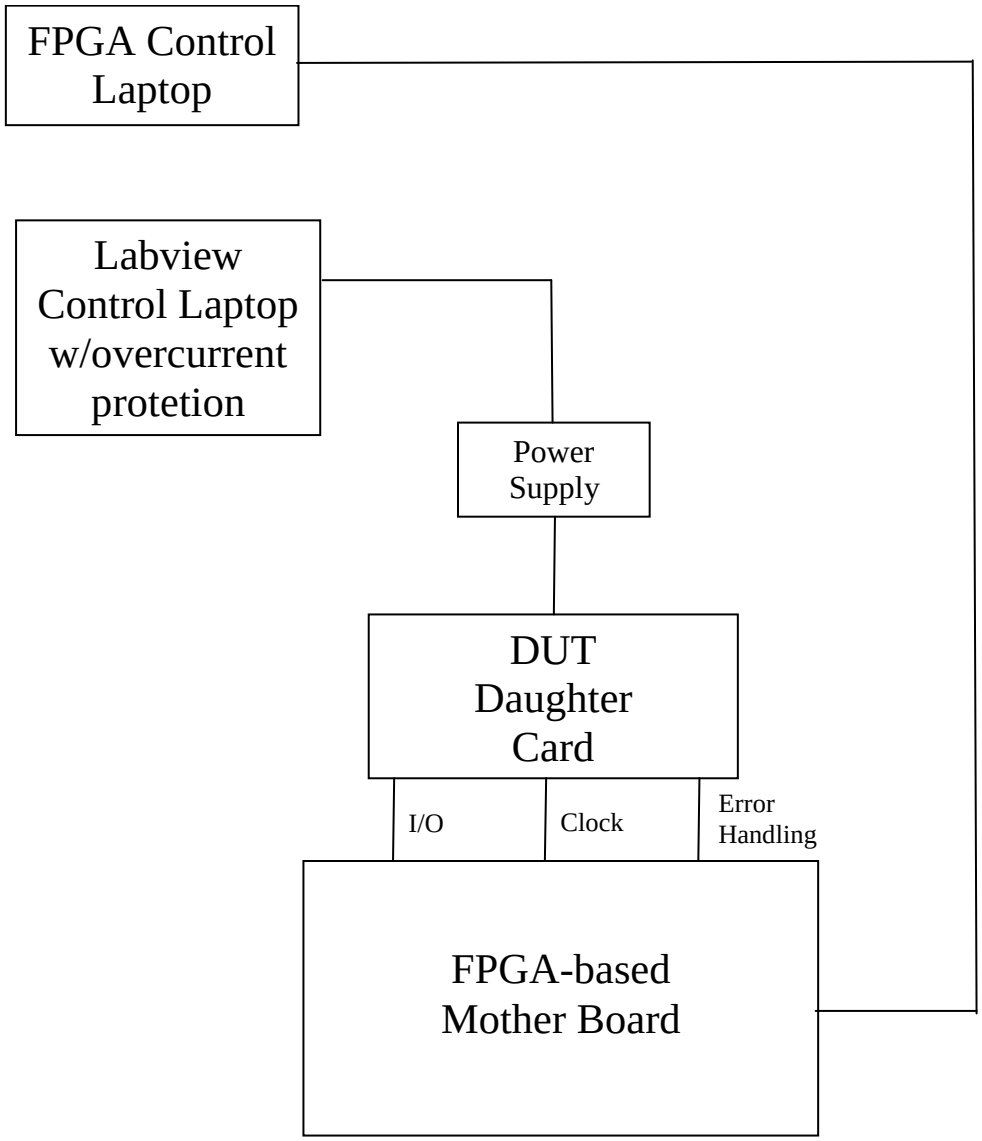

Fig. 1. Block diagram of the flash memory test apparatus.

V. Results

DUTs 1-2 were tested in read-only mode, while DUTs 3-5 were exercised in all the test patterns and the dynamic mode, as described above. All the DUTs had some bad blocks. At the 10, 20, and 30, 50, and 75 krad (SiO_2) exposure levels, there were no errors in any device, in any test mode, except those identified in the bad blocks prior to irradiation, with three exceptions. DUTs 1 and 2 (both read only) had 21 and 366 errors, respectively, in the initial read. For DUT 1, these all reset correctly, and all but eight reset for DUT 2. Unfortunately, the bit map for these eight errors was not preserved, but they were later revealed to be one to zero errors, in the next exposure. For both DUTs, the initial errors were all zero-to-one, as would be expected from the buildup of radiation-induced positive charge. DUTs 3 and 4 had no errors initially, or while being exercised. DUT 5, after 75 krads(SiO_2), had no errors in the initial read, but when the pattern was

reversed (to 55), there were 70 errors. These were all single bit errors, where a zero was written instead of a one. Then at the end, when the pattern was reset to AA for the next exposure, there were 55 errors, also ones written incorrectly as zeroes. This error mode is something new, and unique, in our experience with TID testing of flash memories. In all our previous tests, the high voltage charge pump has been the first thing to fail, with the result that the erase function, or the write function, is lost, or sometimes both. In this case, the erase function is still working perfectly, and the write function is working, too—actually too well. Bits are being written, even when they are not supposed to be. This result indicates that the problem is in the control circuitry that turns the write function on and off, and not with the charge pump. A further indication of this point is that the errors are distributed in a highly non-random manner. For example, in the AA pattern, there were 22 consecutive errors in blocks 3884 and 3885, where column 275 contained the error. Page numbers differed by four or a multiple of four from one error to the next, and the same bit was in error each time. Chances of such a regular error pattern occurring from a random variation of the individual bits are vanishingly small. Even with these errors, all the DUTs were returned to the test cell for another exposure, because the ECC could have corrected them.

At 100 krads (SiO_2), DUT 1 had no errors, and DUT 2, which began with eight errors, had four of them corrected by the exposure. The four that were read were one-to-zero errors, and the radiation-induced positive charge was enough to turn the other four back into ones. (Later, at higher doses, and with larger error counts, this was a common effect—half or more of the incorrectly written zeroes were typically corrected by the exposure. Usually, none of the correctly written zeroes were changed, however.) DUTs 3 and 4 had no errors, either initially, or at any other point when being exercised. DUT 5 had 19 errors initially (that is, 35 corrected), but multiple “initial” reads were performed, and the error count went up, first to 1328, and finally to 1254. With the 55 pattern, the error count was 1547. When the AA pattern was restored for the next exposure, the error count was 1034. In all cases, these were all one-to-zero single bit errors. The pattern of the errors was highly non-random, as described previously. Another non-random feature of the error distribution was first noticed at this dose level, but proved to be a recurring thing at higher dose levels. That is, there were no errors, or very few at most, in the first thousand or so blocks, and also in the last thousand or so blocks. The errors were all concentrated in the middle of the logical address range. We don’t know the correlation of physical bit location with logical address, but it is possible that voltages on longer signal lines are not maintained, because of iR drops, when radiation-induced increases in parasitic leakage current come into play. This would be less of a problem near the edge of the chip. If this is happening, however, it has to be very localized, because the power supply current for the whole chip has changed hardly at all. Pre-radiation, the standby current was 1 mA, compared to 8 mA during erase operations, 10-14 mA during read operations, depending on the pattern, and 14-17 mA during writes, again depending on the pattern. At this dose level, standby and erase currents had not changed at all. There was in increase of at most 1-2 mA during read and write operations, depending on which DUT one looks at.

At 125 krads (SiO_2), DUT 1 had no errors, and DUT 2 (which had had eight and four at the two previous measurements) was down to two errors. DUT3 had no errors, either initially, or while being exercised. DUT 4 had two errors in the initial read, 16 errors in the 55 pattern, and 11 when reset to AA. These were all single bit one-to-zero errors. DUT 5, which had 1034 errors when reset after the previous exposure, had 550 initially. But when the part was rewritten, the number increased to 8319 in the 55 pattern, and 5739 in the AA pattern, for the next exposure.

At 150 krads (SiO_2), DUT 1 had no errors. DUT 2, which had two errors (one-to-zero) had 13 errors, but of the opposite polarity (zero-to-one). These had the same sort of clustered, non-random pattern we expect for one-to-zero errors. DUT 3 had no errors in the initial read, but when new patterns were written, in the course of exercising the part, errors started to occur. There were 50 one-to-zero errors in the 55 pattern, and 42 errors when the DUT was reset to AA for the next exposure. DUT 4 had seven errors in the initial read, but this number increased to 118 when 55 was written, and to 188 when AA was rewritten for the next exposure. These were all one-to-zero errors. DUT 5 had 2695 errors in the initial read, down from 5739 previously. But, as with other DUTs, the number increased when other patterns were written, 52433 for 55, 37982 when AA was rewritten for the next exposure.

At 200 krads (SiO_2), DUT 1, which had had no errors previously, suddenly had more than one million, most of which did not reset. Eight blocks (actually, four pairs of blocks) had gone completely bad. These were 2002 and 2003; 5272 and 5273; 6310 and 6311; and 6762 and 6763. Instead of a checkerboard pattern, they were reading all zeroes, and they did not respond to any write or erase commands. We have seen blocks fail in this manner in previous flash tests. However, in this test, only DUT 1 ever had any radiation-induced bad blocks, and DUT 1 did not have any more, after this first group of eight. DUT 1 also had 7675 single bit zero-to-one errors scattered throughout the rest of the memory, in the initial read. After an initial erase, there were no errors except for the bad blocks. When the DUT was rewritten for the next exposure, there were 3092 single bit one-to-zero errors (that is, opposite polarity from before the reset), in addition to the eight bad blocks. DUT 1 still had many fewer bad blocks than the manufacturer allows, so it remained in the test—the new ones could have been screened out, like those existing pre-radiation. DUT 2 had 40416 errors in the initial read, all of which were zero-to-one errors (compared to 13 errors of the opposite polarity after the previous exposure). After being reset, DUT 2 had 25433 one-to-zero errors (opposite polarity, compared to those before the reset). DUT 3 had 19 single bit one-to-zero errors in the initial read, compared to 42 errors before the exposure. When DUT 3 was reset, the result was 1974 errors, all one-to- zero. DUT 4 had 61 errors initially, compared to 188 before the exposure, all the same polarity. These errors erased, but when write operations were performed, the number increase to 3871 (for 55) and 3606, when AA was rewritten for the next exposure. DUT 5 had 11618 errors initially, compared to 37982 before the exposure. Similar to the other parts, the error count for DUT 5 increased after write operations, to 524,100 (for 55), and 421,273 (for AA).

At 250 krads (SiO_2), DUT 1 had 1300 single bit one-to-zero errors in the initial read, compared to 3092 before the exposure, in addition to the same eight bad blocks. DUT 2 had 10,065 errors, compared to 25,433 before the exposure. Neither of these parts was reset, because that would have increased the error count. DUT 3 had 486 errors initially, compared to 1974 before the exposure. When DUT 3 was reset, the error count was 9501 (for 55), and 11,187 for AA. When AA was set a second time, the result was 9928. DUT 4 had 1064 errors in the initial read (compared to 3606 before), which increased to 17,536 (for 55), and 17,094 for AA. When AA was written a second time, the result was 14,962. DUT 5 had 147,127 errors initially, compared to 414,967 after the previous exposure. When DUT 5 was rewritten, the error count was 1.5M (for 55), and 1.26M for AA. Early in this test, when the number of one-to-zero errors was small, one might have been tempted to argue that the errors were not due to the bits being written incorrectly, but that they were the result of erroneous read operations, instead. The idea would be that the cells are n-channel transistors, where the channel of a zero is "on", because of the negative charge on the floating gate, and the channel of a one is "off" because that charge is missing. The read operation consists of applying a voltage from source-to-drain, with the result that current flows in the zeroes, but not in the ones. Then the one-to-zero errors could be due to radiation-induced parasitic leakage current detected in the wrong place at the wrong time, rather than anything really amiss with the cells themselves. One problem with this idea has already been pointed out—power supply current has hardly changed at all, so far. But in this last exposure, the number of one-to-zero errors dropped by half to three quarters of the pre-radiation number on all five DUTs <u>during</u> irradiation, when leakage current would be expected to increase, if it changed at all. On the other hand, the error count goes up every time a write operation is performed, sometimes by an order of magnitude. So the parts are being written incorrectly, and not because the charge pump is failing.

At 300 krads (SiO_2), DUT 1 had 588 single bit one-to-zero errors, compared to 1300 after the previous exposure, and 3092 before that. The same eight bad blocks were also observed, again. DUT 2 had 3863 total errors, which included roughly equal numbers of both polarities. Before the exposure, there were 10,065 errors, all one-to-zero. That is, most of the original erroneous zeroes were repaired by the radiation, but some bits that were supposed to be zeroes were also corrupted. DUT 3 had 2878 errors in the initial read, which increased to 57,601 for the 55 pattern, and 61,420 when AA was reset. When AA was written a second time, there were 46,780 errors, all one-to-zero. DUT 4 had 4917 errors in the initial read, compared to 14,962 errors before the exposure. This number increased to 96,553 for the 55 pattern, and 85000 when AA was rewritten. When AA was written a second time, the error count was 74,102. In all cases, these were one-to-zero errors. DUT 5 had 441,837 errors in the initial read, which was down from 1.26M after the previous exposure. This number increased to 706,690 for the 55 pattern, and 3.24 M when the DUT was reset to AA for the next exposure.

At 400 krads (SiO_2), DUT 1 had 150 M errors. In many of the even numbered blocks, all addresses appeared bad, with a single bit one-to-zero error (1000 1010 vs. 1010 1010 expected). Many of the odd numbered blocks had a few zero-to-one errors. The eight previously bad blocks were still bad, and still all zeroes. After being erased, DUT 1 had

no errors except the same eight bad blocks. When the AA pattern was rewritten for the next exposure, there were about 500 K single bit one-to-zero errors, in addition to the bad blocks. DUT 1 could have been removed from the test at this point, but it wasn't because the test was going to continue, either way. DUT 2 had about 7.6 M errors in the initial read, and 2.6 M after being reset. DUT 3 had 6398 errors in the initial read, all one-to-zero. After resetting to the 55 pattern, there were 445,375 errors, and 431,874 after resetting to AA. Programming AA a second time resulted in 306,767 errors. DUT 4 had 13,621 errors in the initial read, and 477,694 when written in the 55 pattern. In the AA pattern, there were 464,405 errors, and 297,325 after a second write. DUT 5 had 471,734 errors in the initial read, 579,747 in the 55 pattern, and more than 10 M when the AA pattern was rewritten. As before, these were all single bit one-to-zero errors. Since this is one percent of the total word addresses in the DUT, it was removed from the test at this point.

At 500 krads (SiO_2), DUT 1 had about 150 K single bit errors (down from about 500 K), in addition to the bad blocks. When it was reset, this number increased about 8x, to 1.2 M. DUT 2 had about 413 K errors in the initial read, which increased to more than 5 M when the part was reset. DUT 3 had 38,975 errors in the initial read, and 870,247 in the 55 pattern. When the AA pattern was rewritten, there were about 1.5 M errors. DUT 4 had 61,189 errors in the initial read, and 470,511 in the 55 pattern. When the AA pattern was rewritten, there were about 1.5 M errors. At this point, the test could have been stopped, but the remaining 4 parts received one more dose increment, to 700 krads (SiO_2), because the charge pump still had not failed. It did not fail at 700 krads, either—the only thing that changed was that there were more errors. The error count increased by roughly an order of magnitude for all the remaining DUTs. After 700 krads, the test was stopped.

These results are summarized in Fig. 2, for DUTs 3 and 4, which shows the number of errors on the initial read after each dose increment, along with the (higher) number after the part is reset for the next exposure. One can see how many of the errors recover during each exposure, and how many more errors are introduced by the next write operation. We note that there is an erase operation between the read and the write, which is not shown, but the error count is zero after the erase. That is, a pattern of all ones is both expected, and actually observed. DUT 5 was qualitatively similar, in general, with the first errors attributable to a write step occurring at 75 K, but there were more errors, sooner. DUTs 1 and 2 were read-only, so they were not exercised in the same way, so they are not shown. They were not written as often, and so had no errors attributable to write operations until 200 krads (SiO_2).

One reason this test was continued as long as it was, is that most of the errors appeared to be single bit errors, which could possibly be corrected by the error correction code (ECC) that the manufacturer recommends. However, the post-processing of the raw data indicates, as we have pointed already, that the errors are not due to inherent properties of the memory cells. Rather, they are due to the fact that the programming is not being done correctly, after the peripheral logic accumulates enough dose. For this reason, it is not

clear that ECC will help with these errors. The ECC can identify single bit errors, and tell the DUT to correct the erroneous bit(s). But if the programming was not done correctly originally, the re-programming might not be correct either. Therefore, we consider the part to have failed, when programming starts to introduce errors. Errors may occur before this point, particularly in the read-only parts. But if they can be reset correctly, the part is not considered to have failed.

These parts were stored, unbiased, on a work-bench following the test. They were checked for annealing, a few weeks later—62 days, to be exact. Every DUT had significant numbers of errors, in every pattern, even after that much annealing time. They number of errors varied with the pattern and the DUT number, but the range was hundreds of thousands to millions. The number of errors was marginally lower than immediately after the exposure, but the recovery was far from complete. We plan to continue monitoring the annealing and recovery of these samples, but results so far do not indicate a significantly different response should be expected in a low dose rate irradiation.

VI. Recommendations

The Samsung 8G NAND, K9F8G08U0A, had no errors up to 50 krads (SiO_2), with the first errors occurring at 75 krads (SiO_2), which was the next data point. Only one of the five DUTs had errors at that level. Two more had errors at the next dose level, 100 krads, and the read-only parts had no uncorrectable errors until well past 100 krads. Most NASA systems require 50 krads (SiO_2) TID tolerance, and a factor of 2 margin (100 krads) is considered highly desirable. These parts meet the basic requirement, but not all of them have the desired margin. Therefore they can be considered for use in most NASA systems, but it is recommended that they be tested again. The radiation response of unhardened commercial microelectronics can vary widely, from date code to date code, for example, or even within a date code. This test should be repeated on parts more nearly representative of those that will actually be flown.

This part is promising for another reason. Every previous flash memory tested has failed because the high voltage charge pump circuit failed, sometimes at very low doses. Then the typical result is that the erase function, or sometimes the write function, could no longer be performed, and the chip experienced total functional failure. In this case, the charge pump was still working, even after 700 krads (SiO_2). However, other radiation-induced damage to the control logic became apparent at doses much lower. But the fact that the chronic, weakest link has been made very robust by some kind of design change (unknown, at least to us) is still a good sign. When the biggest problem is solved, one typically finds out what the second biggest problem is, and that is what has happened here.

The Samsung 8G NAND Flash, K9F8G08U0A, has not yet been characterized for SEE sensitivity. Previous flash memories have been very resistant to soft errors in the memory array itself, because they are nonvolatile. That is, the information storage does

not depend on voltages being maintained. However, the on-chip processor is sometimes very sensitive to single event functional interrupts (SEFI), and these interrupts are expected to be the main problem. This response has varied widely, in previous tests.