Single Event Response of the Samsung 16G NAND Flash

Timothy R. Oldham[1], Mark R. Friendlich[2], and Edward T. Wilcox[2]

Abstract— **This study was undertaken to determine and compare the susceptibilities of 16 Gbit NAND Flash memories from Samsung to destructive and nondestructive single-event effects (SEE).**

I. Introduction

This study was undertaken to determine and compare the susceptibilities of 16 Gbit NAND Flash memories from Samsung (part number K9FAG08U0M) to destructive and nondestructive single-event effects (SEE) for use in possible future NASA missions. The devices were monitored for SEUs (Single Event Upset), errors from individual cells, for SEFIs (Sinele Event Functional Interrupts), errors arising in the control logic, and for destructive events, including latchup, induced by exposing them to a heavy ion beam at the LBNL (Lawrence Berkeley National Laboratory) Cyclotron.

II. Description of Devices under test

In this test, two parts were used (part number K9FAG08U0M-HCB0, Lot Date Code (LDC) 1026). The parts have 2Gx8 organization with large blocks. That is, the blocks are (512K+40K)x8, with 64 pages/block. Each page is nominally 8Kx8, but they also have 640 redundant columns, which makes the total page size (8K+640)x8. NAND flash normally has some bad blocks which can be screened off. The specification is that 4056 of the 4152 blocks will be good. In our experience, the parts almost always have a few bad blocks, but it is usually a single digit number. The device technology is 32 nm minimum feature size CMOS (Complementary Metal Oxide Semiconductor) NAND flash memory. All the parts are single die, SLC (single level cells). The chips came in a 132-pin BGA (ball grid array) package, but the plastic had been dissolved on the topside to expose the chips, allowing the beam to reach the chip surface. The parts have a single power supply which is 2.7-3.6 V, full range, with 3.3 V nominal. This means the parts have a high voltage charge pump to produce the high voltages necessary for Program (Write) and Erase operations.

1. Ball Aerospace, Boulder, CO 80301
2. MEI Technologies, Inc., Landover, MD 20706

III. Experimental Procedure

Because Flash technology uses different voltages and circuitry depending on the operation being performed, testing was performed for a variety of test patterns and bias and operating conditions. Test patterns included all 0's, all 1's, checkerboard and inverse checkerboard. In general, all zeroes is the worst-case condition for single bit errors. For a zero, the floating gate is fully charged with electrons. An ion can have the effect of introducing positive charge, which may be enough to cause a zero-to-one error. However, a checkerboard pattern (AA) was used in many tests because errors in the control circuitry can cause errors of both polarities. One-to-zero errors are an indication that the errors are coming from the control circuits. Between exposures, all patterns can used to exercise the DUT (device under test), to verify that it was still fully functional. However, all patterns are not used routinely, just because it is too time consuming to do so. The maximum clock frequency for these devices was 40 MHz, which is also the frequency used in the dynamic testing.

Bias and operating conditions included:

1) Static/Unbiased irradiation, in which a pattern was written and verified, and then bias was removed from the part and the part was irradiated. Once the irradiation reached the desired fluence, it was stopped, bias was restored, and the memory contents were read and errors tallied.

2) Static irradiation, which was similar to unbiased irradiation, except that bias was maintained throughout irradiation of the part.

Note that these conditions provide no opportunity to monitor functional or hard failures that may occur during the irradiation. It was also not possible to monitor the power supply current during the unbiased tests, but this was done in all the other test modes.

3) Dynamic Read, in which a pattern was written to memory and verified, then subsequently read continuously during irradiation, and the errors counted. This condition allows determination of functional, configuration and hard errors, as well as bit errors. In this mode, the number of static bit errors is determined by reading the memory again, after the beam is turned off. In reporting for this test mode, we normally distinguish between transient

Presented by Timothy R. Oldham at the International Electrical and Electronics Engineering (IEEE) Nuclear and Space Radiation Effects Conference (NSREC), July 11, 2013, San Francisco, CA, published in the IEEE Radiation Effects Data Workshop and published on nepp.nasa.gov and radhome.gsfc.nasa.gov.

errors, due to transient noise in the Read circuit, and static errors, where the bits are actually corrupted.

4) Dynamic Read/Write, which was similar to the Dynamic Read, except that a word in error was rewritten. In this mode, the words that are read are compared to the stored pattern. When errors are detected, the addresses are rewritten. Because the Write operation uses the charge pump, it is expected that the Flash could be more vulnerable to destructive conditions in this test mode. In this mode, a block is Read, and rewritten, if necessary, followed by the next block, and so on. Note that this mode only corrects one polarity of error, which is zero-to-one.

5) Dynamic Read/Erase/Write, which is similar to the R/W mode, except that it corrects both polarities of errors. A block in error is Erased, and then rewritten. Since the high voltage charge pump is used for both the Erase and Write operations, it is expected that the parts will be most vulnerable to destructive effects in this mode.

6) In this set of experiments, we have attempted to look at angular effects, which may include multiple bits grazed by the same ion, and other effects due to charge sharing by multiple nodes in the control logic. This test was done with the DUT tilted 45 degrees. In previous testing, with TSOP (thin small outline package) packaged parts, the socket used to hold the DUT prevented tests at any higher angle. This restriction does not really apply with the BGA package, but there was not enough beam time to test at a variety of angles. Since we only had time for one angle, we used the same angle we had used previously, which allows us to compare these results with previous tests.

The ions used in testing were obtained using the LBNL 10 MeV/nucleon cocktail, and are summarized I Table I.

Table I. Ions used in testing, where LET (Linear Energy Transfer) is given in MeV/(mg/cm^2).

Ion	E(MeV)	Surface LET	Angles	Eff. LET
Ne	216	3.49	0, 45	3.49, 4.9
Ar	400	9.74	0, 45	9.74, 11.8
Kr	906	30.2	0, 45	29.3, 42
Xe	1232	58.8	0, 45	58.8, 82

IV. Results

In this test, we observed only one functional failure, and we believe it was the result of cumulative damage over many beam runs, rather than the result of a single ion interaction. The results shown in Fig.1illustrate a pattern we have observed numerous times. In this case, a fresh part has Write current of 18.5 mA, but as the part is exposed to one beam run after another, the Write current increases to 20 mA to 22 mA to 25 mA, and then the part fails, as it does here. After failure, the Write Current jumps up to 30 mA, in this case. In a case like this, we report the failure by plotting it for the shot where it was observed, as we will do here, later. But the failure is clearly the result of cumulative damage, and not something that happened on just the last shot. We have observed similar results on previous technology generations of Samsung flash going back many years, although we have not plotted the results in this manner, previously. In this case, the failure occurred on a shot with Kr ions at high angle, but a fresh part survived a similar, later, shot with Xe ions. One would expect the shot with Xe to be a more severe test, because of the higher LET. We have to report failures if they occur, and if we did not plot them where they occurred, where would we plot them? But it is important to recognize that they are not usually due to single ion interactions, but rather to cumulative damage from many shots. The total number of particles that this part was exposed to corresponds to millions of years in space. The fact that parametric degradation was evident almost from the beginning indicates that cumulative damage was also present from the beginning.

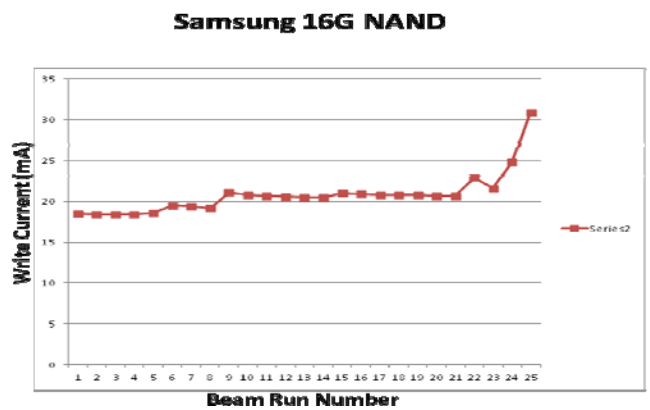

Fig. 1. Parametric degradation, increasing Write current, over many beam runs, followed by eventual Write mode failure.

Presented by Timothy R. Oldham at the International Electrical and Electronics Engineering (IEEE) Nuclear and Space Radiation Effects Conference (NSREC), July 11, 2013, San Francisco, CA, published in the IEEE Radiation Effects Data Workshop and published on nepp.nasa.gov and radhome.gsfc.nasa.gov.

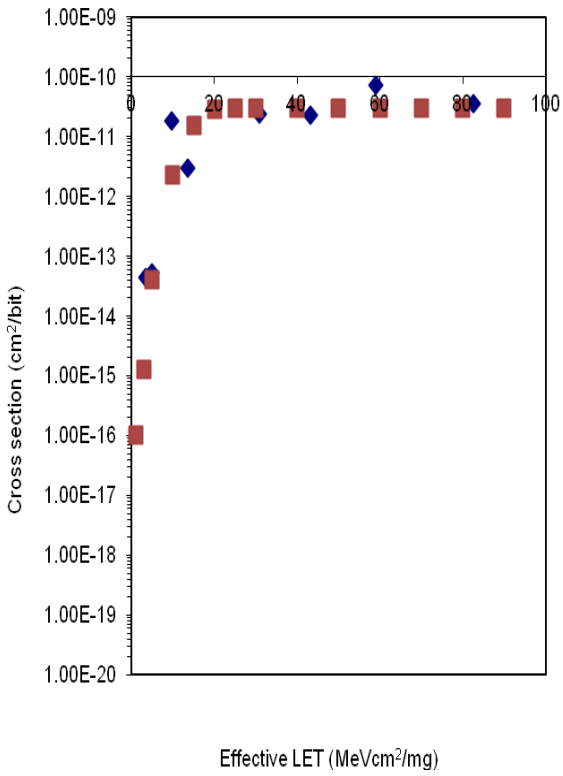

Fig. 2. Measured single bit errors (blue diamonds), compared to Weibull fit (red squares), in static mode with no bias applied.

During testing, the DUTs were irradiated with the ions indicated in Table I. The DUT was oriented normal to the incident beam, or at 45 degrees. The errors observed in static SEU testing are shown in Fig. 2, with no bias applied. The 45 degree data is plotted at the effective LET (LET/cos θ). This is done so that one can distinguish between the normal incidence shots and the 45 degree shots. It is <u>not</u> done because effective LET is necessarily expected to be a useful concept for other reasons. In Fig. 2, the cross section for bit errors is plotted in units of cm^2/bit. In this mode, there were no SEFIs of functional failures (destructive effects). The same results are plotted in Fig. 3, in units of cm^2/device, which we normally do so that SEFIs and functional failures can be plotted on the same scale as bit errors. In this case, there are not any such events, so there is no real difference between the plots. In Fig. 2, the red squares indicate a Weibull curve [1] fitted to the measured cross section, where the threshold LET is taken to be 1 MeV/mg/cm^2, the width, W, is taken to be 15 MeV/mg/cm^2, the exponent or shape factor is 5.0, and the saturated cross section is taken to be $3x10^{-11}$ cm^2. These parameters were used to calculate the Petersen Figure of Merit (FOM) [2-3], which was used, in turn, to estimate the bit error rate for the Adams worst case environment (geosynchronous orbit at solar minimum). In this case, the FOM is approximately $1.86x10^{-13}$, which corresponds to an error rate of about 5 x 10^{-11} errors/ bit-day. This is slightly higher than previous generations of Samsung NAND flash technology [4], but not unexpected, however.

We note here that NAND flash manufacturers typically specify the use of error correction for their parts. This is because without error correction, the parts would not meet any of their performance or reliability specifications. Since error correction is critical to the performance of the parts, we included it in this test. We did not actively error correct during the test, but in post processing, we examined the bit error map to determine which errors would have been corrected. Previous generations of NAND flash, by Samsung and other manufacturers, used SEC/DED (single error correction/ double error detection) Hamming codes [5], capable of correcting one bit of every 64 bytes (512 bits). For example, in the 4G technology generation, the page size was 2Kx8, which meant there were 32 segments of 512 data bits per page to be error corrected. With Hamming codes, a segment of 2^N bits requires N+1 bits for single error correction (SEC), and one more bit, N+2, for double error detection (DED). That is, 11 check bits are required for SEC/DED for a 512 bit segment, or 352 check bits for a 2Kx8 page. The number of bits available for that purpose was actually 512, or more than enough. The 16G parts reported here have more extra bits set aside for error correction, so other, more intensive, error correction schemes could be considered. In this test, the 512 bit error correction scheme described above would have corrected the vast majority of single bit errors, but not all of them. When the higher LET ions, Kr and Xe, were incident at 45º and oriented along the columns, a single ion sometimes caused two bit errors. The same bit would be in error in two adjacent 32 bit strings. We did not test with the beam oriented along the rows, but we would expect adjacent bits to be in error in that case, as reported in [6]. The ions that cause these double errors are rare in space, so the double error rate in space would be low, but not zero.

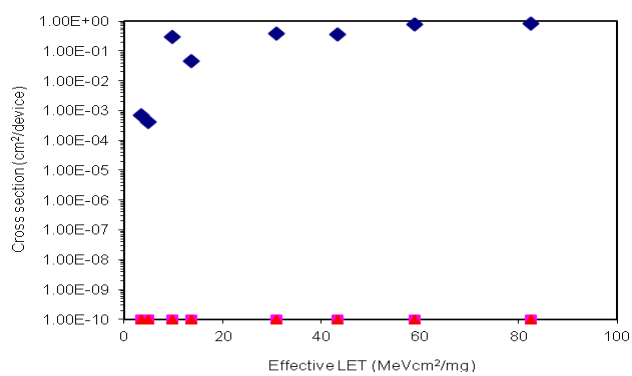

Fig.3. Single bit errors (blue diamonds) in static mode without bias, along with SEFIs (red squares), where SEFI rate was zero.

Presented by Timothy R. Oldham at the International Electrical and Electronics Engineering (IEEE) Nuclear and Space Radiation Effects Conference (NSREC), July 11, 2013, San Francisco, CA, published in the IEEE Radiation Effects Data Workshop and published on nepp.nasa.gov and radhome.gsfc.nasa.gov.

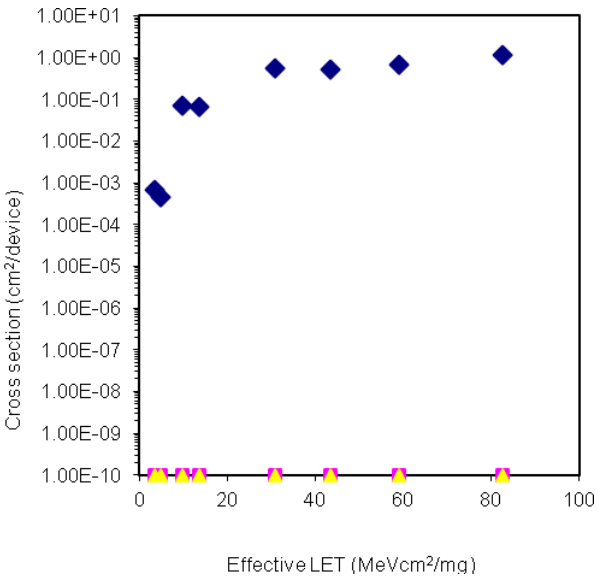

Fig. 4. SEU, SEFI, and Destructive results for static mode, with bias applied.

The results for static mode with bias applied are shown in Fig. 4. The main effect is single bit upsets (blue diamonds), along with no SEFIs (pink squares), and no destructive events (yellow triangles), again. The bit error rate is essentially the same as for the unbiased condition, as shown in Fig. 3.

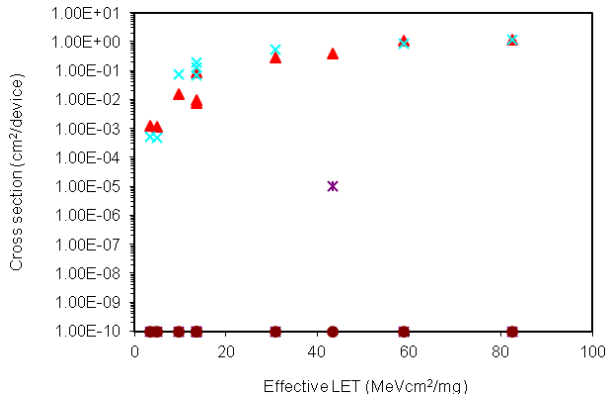

Fig. 5. Results for the Samsung 16G NAND flash in Dynamic Read mode. Transient (teal X) and static errors (red triangles) are shown, along with one SEFI (purple symbol), from Kr ions at 45° incidence. There were no functional failures (brown circles).

For the Dynamic Read condition, the parts showed exhibited transient read errors in addition to the bit errors, which are plotted in Fig. 5. In this mode, the DUT reads continuously with the beam on. The significance of the transient errors is not always completely clear, because the entire memory can be read multiple times, which means static errors will be read multiple times. On the other hand, sometimes the beam is turned off before even one loop through the memory is completed, in which case some static errors will not be counted at all, until after the beam is off. In addition there are errors due to transient noise in the read circuit or the control logic. The static errors shown in Fig 5 are bit errors read after the beam is turned off. Generally the static error count and the transient error count are similar, which suggests that the transient errors are really just static bit errors. For the one shot with a SEFI, it was not possible to determine the static error count. The reason was that there were so many control logic errors that they swamped the handful of real bit errors.

Results of the dynamic R/W tests are shown in Fig. 6. The main reason for including this test was the expectation that the high voltage write operation would contribute to more errors in the control circuits, which does not appear to have happened. There were no functional failures, and only one SEFI, which was, again, with Kr at 45° incidence. The only puzzling result here is that the static errors and transient errors appear to be roughly equal. The problem is that the static errors are supposed to be corrected in this mode. Probably, we did not wait long enough after the beam was off to stop the rewrite operation, which prevented some of the errors from being corrected.

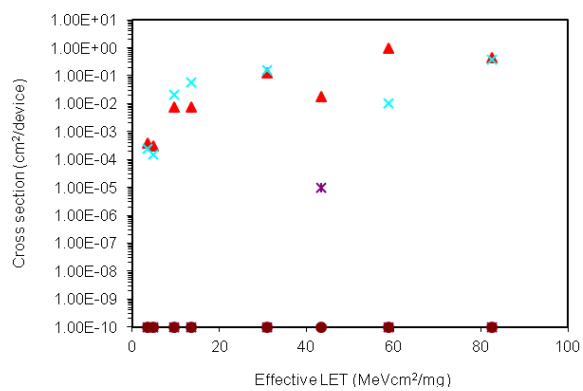

Fig. 6. SEU, SEFI, and Destructive results in the dynamic R/W mode. Symbols have the same meanings as in Fig. 5.

Presented by Timothy R. Oldham at the International Electrical and Electronics Engineering (IEEE) Nuclear and Space Radiation Effects Conference (NSREC), July 11, 2013, San Francisco, CA, published in the IEEE Radiation Effects Data Workshop and published on nepp.nasa.gov and radhome.gsfc.nasa.gov.

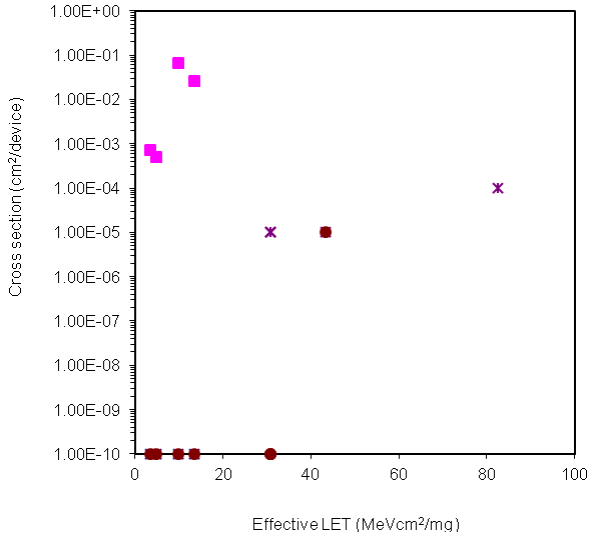

Fig. 7. SEU, SEFI, and Destructive results in Dynamic Read/Erase/Write mode. In this mode errors are corrected as they occur (indicated by pink squares), except that the number cannot be determined when SEFIs (purple symbols) or functional failures (brown circles) occur.

Results for the dynamic R/E/W tests are shown in Fig 7. In this mode, errors are counted as they are read, but then they are erased and rewritten. Therefore, there are no static errors read after the beam is turned off, and bits in error for a time are counted as transient errors. Because the high voltage erase and program operations are performed constantly, it is expected there will be errors in the control logic in this mode. In fact, there were multiple SEFIs and one destructive failure in this mode. The point indicating a SEFI at LET~30 MeV/mg/cm^2 is actually three points, superimposed. The failure was with Kr ions at 45° incidence. After the first part failed, a second part was used for another shot with Kr, and for all the shots with Xe. There were SEFIs on both Xe shots in this mode, but no functional failure, even though Xe might be expected to be more damaging than Kr. On the shots where there were SEFIs, it was not possible to count the single bit errors because there were so many control logic errors. There was no Xe shot at normal incidence because we ran out of time, but there was a Xe shot at high angle.

V. Conclusions

The Samsung 16G NAND flash memory appears to be extremely well suited for space applications. The only functional failure in this test appears to be the result of cumulative damage over many beam runs, and not the result of a single ion effect. There were no SEFIs below LET=30 MeV/(mg/cm^2) in any test mode, and no SEFIs in any static test mode at any LET. There were SEFIs in dynamic test modes starting at LET=30 MeV/(mg/cm^2), but not on every shot. This indicates that the rate for such SEFIs would be very low in space, if not quite zero. The rate for single bit errors, without error correction, is slightly higher than for previous generations of NAND flash, but still many orders of magnitude better than for standard volatile memories. Simple error correction schemes correct the vast majority of these errors, but not all of them. In this 32 nm technology, the individual bits are close enough together that a single ion at the right angle can apparently upset two physically adjacent cells. We had not noticed this in previous generations of NAND flash, with larger feature sizes. With continued scaling, more elaborate error correcting schemes will almost certainly be necessary, at least for space applications. We also note that the Samsung 16G has not yet been tested for total dose, but the 8G predecessor was still fully functional after exposure to 400 krad (SiO_2) [7]. For this reason, we expect the total dose response of the 16G to be excellent.

Acknowledgment

This work was supported in part by the NASA Electronic Parts and Packaging Program (NEPP), and the Defense Threat Reduction Agency (DTRA) under IACRO# 11-4395I. The authors wish to thank Martha O'Bryan for technical assistance in preparing both the poster and the final manuscript for this paper.

References

[1] W. Weibull, J. Applied Mechanics, pp. 293-297, Sept. 1961.
[2] E.L. Petersen, J.B. Langworthy, and S.E. Diehl, Suggested single event upset figure of merit, IEEE Trans. Nucl. Sci., vol. 30, no. 6,pp. 4533-4539, Dec. 1983.
[3] E.L. Petersen, The SEU figure of merit and proton upset rate calculations, IEEE Trans. Nucl. Sci., vol. 45, no. 6, pp. 2550-2562, Dec. 1998.
[4] M. Bagatin, S. Gerardin, G. Cellere, A. Paccagnella, A. Visconti, S. Beltrami, R. Harboe-Sorensen, and A. Virtanen, Key contributions to the cross section of NAND flash irradiated with heavy ions, IEEE Trans. Nucl. Sci., vol. 55, no. 6, pp. 3302-3308, Dec. 2008.
[5] R.W. Hamming, Error detecting and error correcting codes, Bell System Technical Journal, vol. 29, no. 2, pp. 147-161, April 1950.
[6] G. Cellere, A. Paccagnella, A. Visconti, M. Bonanomi, R. Harboe-Sorensen, and A. Virtanen, Angular dependence of heavy ion effects on floating gate memory arrays, IEEE Trans. Nucl. Sci., vol. 54, no. 6, pp. 2371-2378, Dec. 2007.
[7] T.R. Oldham, D. Chen, M. Friendlich, M.A. Carts, C.M. Seidleck, and K.A. LaBel, Effect of radiation exposure on the retention of commercial NAND flash memory, IEEE Trans. Nucl. Sci., vol. 58, no. 6, pp. 2904-2910, Dec. 2011.