

# TID and SEE Response of Advanced 4G NAND Flash Memories

T. R. Oldham, *Fellow, IEEE*, M. Suhail, M. R. Friendlich, M. A. Carts, R.L. Ladbury, *Member, IEEE*, H. S. Kim, M. D. Berg, C. Poivey, *Member, IEEE*, S. P. Buchner, *Member, IEEE*, A. B. Sanders, C. M. Seidleck, and K. A. LaBel, *Member, IEEE*

**Abstract-- We present total dose and SEE responses for 4G NAND flash memories by three different manufacturers. The SEE response is scaled to predict the response to atmospheric neutrons at aircraft altitudes and at sea level using the Figure of Merit.**

## I. INTRODUCTION

We present total ionizing dose (TID) and single event effects (SEE) results obtained on 4G NAND flash memories obtained from three different manufacturers, Samsung, Hynix, and Micron. Measured heavy ion cross-sections have been used to calculate error rates expected for geosynchronous orbit at solar minimum (Adams ten percent worst case environment). The cross-sections have also been used to determine the Figure of Merit (FOM), as proposed by Petersen [1,2], which has then been used to calculate the expected error rate from atmospheric neutrons at aircraft altitudes, and at sea level. The parts have also been tested for SEL (Single Event Latchup), at their maximum rated operating temperature and voltage, with no latchup observed in any part, but other destructive events have been observed. Compared to standard volatile memories, flash memories are normally very robust against SEE, because they are designed to retain information with no voltage applied to the memory cells. The results reported here reflect this, for all environments. However, the peripheral circuits are sensitive to single event transients (SET) and single event functional interrupts (SEFI), much like any other CMOS. Rates for these effects will be discussed. Three different TID tests are reported: (1) high dose rate, nominal bias test; (2) a high dose rate, unbiased test; and (3) a low dose rate (LDR), nominal

bias test. We also present an initial TID test on a Samsung single chip, 8G NAND flash memory. There is a range of TID responses for the different manufacturers, and for the different test conditions, but the best results reported here are good enough for most NASA missions.

## II. DESCRIPTION OF SAMPLES

The samples used in this study are 4G NAND flash memories produced by Micron (part number MT29F4G08AAAWP, lot date code (LDC) 628), Hynix (HY27UF084G2M, LDC 0636), and Samsung (K9F4G08U0A, LDC 625). All are single die, with single level cells (SLC). We eventually ran out of the Samsung K9F4G08U0A parts, so, for the zero bias and LDR TID tests, we substituted K9K8G08U0A parts (LDC 622), which are two 4G chips, stacked in an 8G configuration. The chips are very similar to the single chip 4Gs, and the stacking arrangement was not expected to affect the TID response, nor did the results seem inconsistent. All are nominally 3.3 V parts, with an allowed range of 2.7-3.6 V. They have an on-chip charge pump to generate the higher voltages needed for erasing and programming. All have the same architecture: pages 2Kx8, blocks 64 pages or 128Kx8, and 4096 blocks. All are allowed to have up to 80 bad blocks, which are normally screened off. Our samples all had a few bad blocks, usually a single digit number, and never close to 80. The Samsung 8G single chip parts (K9F8G08U0A, LDC 719) were similar in organization, except that they had 8192 blocks.

## III. EXPERIMENTAL PROCEDURE

TID testing was done at the Co-60 Radiation Effects Facility at NASA/GSFC. This is a room air source, where the pencils are raised up out of the floor, during exposures. Active dosimetry was performed, using air ionization probes. Testing was done in a step/stress manner, using a standard Pb/Al filter box. The initial test was done in accordance with MIL-STD-TM 1019, although a zero-bias test has also been performed, and a low dose rate exposure is in progress. Parts were under DC bias during exposures, except as noted, but not actively exercised. Eight test devices, four each from Hynix and Micron, were programmed with a checkerboard pattern (AA) during exposures, and biased at 3.6 V (3.3 V

---

Manuscript received June 9, 2008. This work was supported in part by the NASA Electronic Parts and Packaging (NEPP) Program and the Defense Threat Reduction Agency (DTRA) under IACRO 08-43431.

Timothy R. Oldham is with Perot Systems Government Services, Inc., c/o NASA/GSFC, Code 561.4, Greenbelt, MD 20771 (USA), phone: 301-286-5489, email: timothy.r.oldham@nasa.gov.

M. Suhail is with Freescale Semiconductor, Austin, TX 78735 (USA).

M. Friendlich, M.A. Carts, H.S. Kim, M.D. Berg, and C.M. Seidleck are with MEI Technologies, Inc., Seabrook, MD 20706 (USA).

R.L. Ladbury and A.B. Sanders are with NASA Goddard Space Flight Center, Greenbelt, MD 20771 (USA).

C. Poivey is with European Space Agency, Noordwijk, Netherlands.

S.P. Buchner is with Perot Systems Government Services, Inc., c/o NASA/GSFC, Code 561.4, Greenbelt, MD 20771 (USA).

nominal power supply, plus 10%). (Samsung parts had been TID tested previously [3].) Four parts, two from each manufacturer, were read (only) between exposures, to look for problems related to the integrity of the individual bits. The other four 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 was also a dynamic test mode, where each block was read, erased, and programmed, then the next block, and so on until the entire memory was completed.

For SEE testing, bias and operating conditions included all the static and dynamic test modes previously reported [3], in addition to latchup testing at 70° C, and 3.6 V, which are the maximum rated operating conditions. Dynamic test modes with high voltage operations were expected to be more likely to produce SEL, so they were emphasized. SEE testing was done at the Texas A&M (TAMU) Cyclotron, using the 15 MeV/ nucleon tune. Ions are indicated in Table 1.

TABLE I  
IONS/ENERGIES AND LETS FOR THIS TEST

TAMU Ions	Energy (AMU)	Energy (MeV)	Angle	Effective LET (MeV•cm <sup>2</sup> /mg)
Ne	15	300	0	2.7
Ar	15	600	0	8.4
Kr	15	1260	0, 45	30.1, 42.1
Xe	15	1965	0	54.8

#### IV. TID RESULTS AND DISCUSSION

TID results are summarized in Fig. 1. For the Hynix parts, all passed at 30 krad (SiO<sub>2</sub>), the first part failed at 50, and the rest failed at 75 krad(SiO<sub>2</sub>). For Micron, all parts passed at 75 krad (SiO<sub>2</sub>), one failed at 100, two more failed at 150 krad (SiO<sub>2</sub>), which was the next exposure level. At this point, the test was terminated, with one part still working. The chart assumes it would have failed at the next level, 200 krad (SiO<sub>2</sub>). For Samsung, which is shown here for comparison, all parts passed at 100 krad (SiO<sub>2</sub>), one failed at 125, and the other seven failed between 150 and 200 krad (SiO<sub>2</sub>).

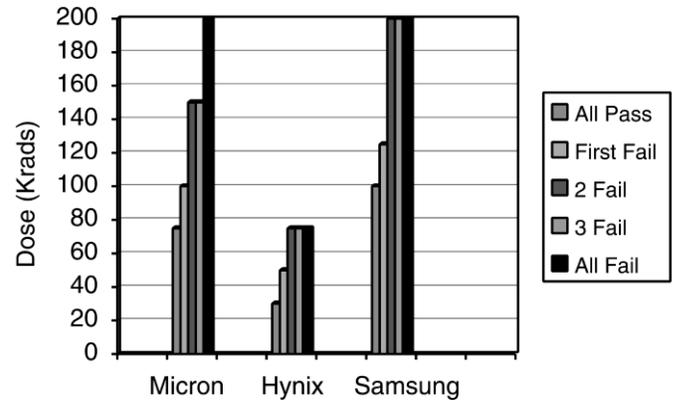


Fig 1. TID response for Micron, Hynix, and Samsung 4G NAND flash memories.

In the zero bias test, the irradiations were performed with the parts mounted on conducting foam. The Samsung parts had no failures through 400 krad (SiO<sub>2</sub>), although some errors had to be reset. First failure was at 500 krad (SiO<sub>2</sub>), with the last two parts failing at 700 krad (SiO<sub>2</sub>). The first Micron part failed at 200 krad (SiO<sub>2</sub>) (passed at 150), because errors could not be reset completely. Others failed at 300 krad(SiO<sub>2</sub>), because the write operation was not successful. Before failure, there were other errors that could be reset. The first Hynix part failed at 200 krad (SiO<sub>2</sub>), because errors could not be reset completely. Others did not fail until 500 krad (SiO<sub>2</sub>), when the write operation failed. However, all the parts had errors before total failure, which were reset successfully. This test was performed because the use of cold spares is sometimes considered as a way of extending system lifetime. Parts from all three manufacturers survived past the point where they failed in the biased test, which suggests that cold spares can be used successfully, in general.

We are also performing a low dose rate exposure (0.02 rads (SiO<sub>2</sub>)/sec) with bias. We have not studied the annealing response of any of these parts in detail, but some parts tested to failure have actually been fully functional after annealing for a few days, unbiased, at room temperature. For this reason, one might think the parts would survive to higher doses, possibly much higher, if there is more time for annealing during the exposure. Results have varied widely, however. The Micron parts all passed at 75 krad (SiO<sub>2</sub>), but failed at 100 krad (SiO<sub>2</sub>), because the parts stopped responding to any commands. This is roughly the same as the high dose rate result. The Hynix parts all passed at 75 krad (SiO<sub>2</sub>), but failed at 100 krad (SiO<sub>2</sub>), because the write and erase functions failed. This is only slightly better than results at high dose rate. The Samsung parts have had no failures at 300 krad (SiO<sub>2</sub>), (past the point at which all parts had failed in the high dose rate test), and the test is continuing.

For the single chip 8G parts from Samsung, we observed a failure mechanism different from most other flash memory we have previously tested. The high voltage charge pump never failed, even at 700 krad (SiO<sub>2</sub>). Consequently, the high voltage program (write) and erase operations never failed.

The failure mechanism was that the write circuit was actually working too well—the control logic that was supposed to turn off the write circuit eventually failed. There were no errors of any kind through 50 krad (SiO<sub>2</sub>). At 75 krad (SiO<sub>2</sub>), one read-only part, and one fully exercised part had errors in the initial read operation after the exposure. (Other parts had no errors of any kind). These errors were zero-to-one errors, where radiation-induced positive charge neutralized (or compensated) the negative charge stored on the floating gate. Often, errors of this kind can be erased and rewritten, so they are not considered fatal. In this case, the rewrite operation introduced errors of the opposite polarity. That is, the part was successfully erased, but, in the write operation, negative charge was injected in locations that were not supposed to be injected. The errors are distributed in a highly non-random manner, which also points to the control logic. For the fully exercised part, 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 in each word. Chances of such a regular error pattern occurring from a random variation of the individual bits are vanishingly small. These two parts could have been considered to have failed at this point, because they had errors that could not be corrected. However, they continued in the test, to see what would happen next. At the 100 krad (SiO<sub>2</sub>) level, there were no errors in the parts that had not had errors before. In the two parts that had errors from the reset operation after the previous exposure, the radiation actually corrected most of the errors. In subsequent exposures, this turned out to be a typical result—the radiation-induced positive charge was enough to correct half to three quarters of the incorrectly written zeroes. That is, it was enough to turn them back into ones. Normally, however, it was never enough to corrupt any of the correctly written zeroes. Apparently, the errors are only partially written. When these two parts were rewritten after this exposure, the number of errors increased again. At 125 krad (SiO<sub>2</sub>), one more fully exercised part started to have errors when rewritten. The parts with errors before had fewer after irradiation, but more when reset. At 150 krad (SiO<sub>2</sub>), the last fully exercised part started to have errors when rewritten. The last read-only part had large numbers of errors at 200 krad (SiO<sub>2</sub>), which could be erased but not rewritten. All the parts received several more dose increments, with qualitatively similar results—number of errors decreased with irradiation, but increased with rewriting. These results are summarized for two of the exercised parts in Fig. 2.

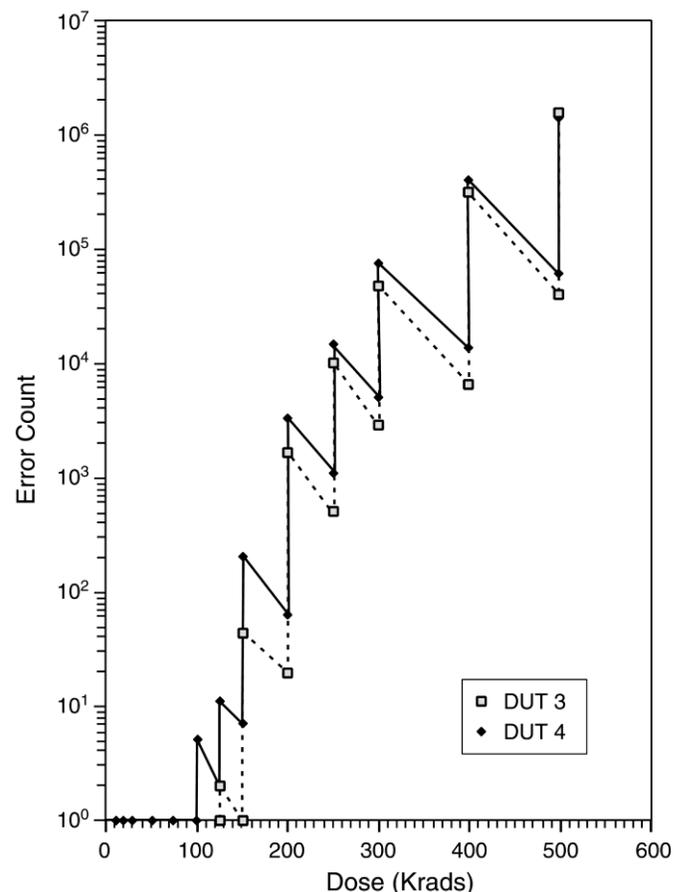


Fig 2. Error count for Samsung 8G NAND flash as a function of dose, showing recovery during irradiation, and new errors from reprogramming.

## V. SEE RESULTS AND DISCUSSION

During testing, the DUTs were irradiated with the ions indicated in Table I. The DUT was oriented normal to the incident beam, except as noted. The errors observed in static testing are shown in Fig. 3a for the Micron part, and Fig. 3b for the Hynix part.

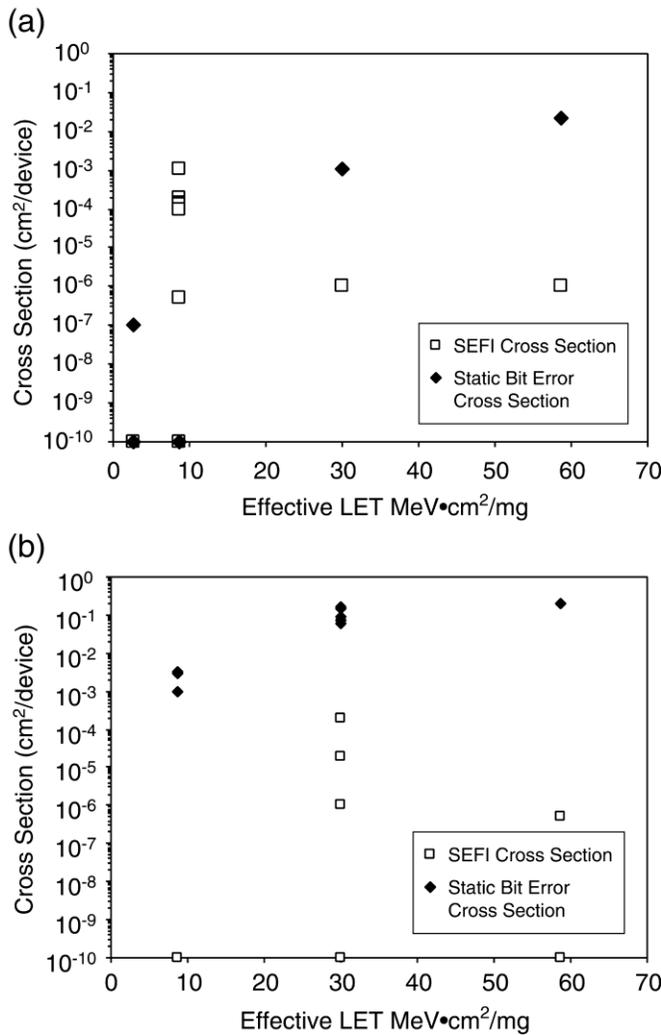


Fig 3. Static mode bit error cross-section, and SEFI cross-section (a) Micron; and (b) Hynix.

Even for the static case, bit errors and Page/Block errors were evident in the patterns of upsets observed. It is likely that the Page/Block errors arise due to upsets in configuration registers in the memory array. Because the DUT was not actively exercised during the exposure, we could not determine exactly when a page/block error occurred, so cross sections are approximate for these error modes. Here and in the following discussion, bit errors are taken to be single bits, which are flipped, as a result of the interaction with incident ions, normally from zero to one. We do not have the physical to logical address mapping, which would allow us to look for multiple bit errors (error clusters) for these parts. However, in the overwhelming majority of cases of bit errors, there is only one error in a page, or one error in an entire block, which makes it extremely unlikely that there will be multiple errors from a single ion. This result is consistent with previously published results on the upset mechanism in flash memory—an ion passing through a floating gate creates a dense charge column, which creates a conducting path between the gate and substrate, which allows charge to leak off the floating gate. Since the ion only hits one gate at normal incidence, only one bit is affected. This situation is far different from that in

volatile memories, where charge generated in the Si substrate can be shared across multiple nodes. The only apparent multiple bit errors are cases where an entire page or a block (or a large part of one) upsets simultaneously—these page and block errors are attributed to errors in the control logic, rather than to the individual bits. These are counted as SEFIs (Single Event Functional Interrupt). In general, a SEFI is any event where the entire DUT, or a large part of it, stops working, presumably from an interaction with a single ion. As a practical matter, most of the SEFIs recorded here are either page errors or block errors, although some involve multiple pages or multiple blocks. Some are also watchdog errors, where the DUT simply stopped responding to commands.

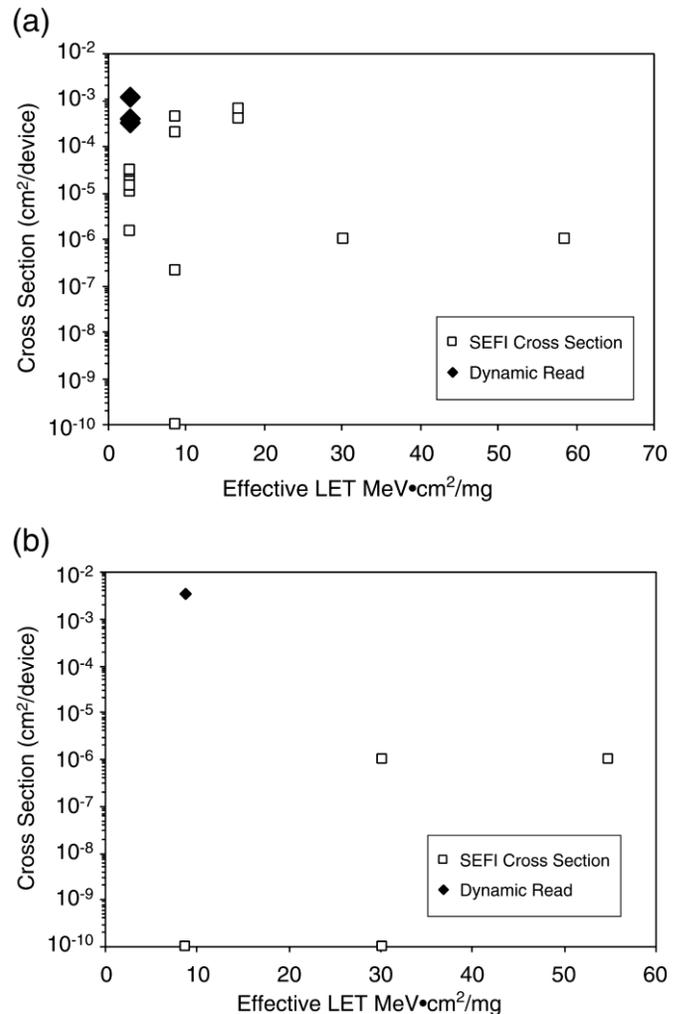


Fig 4. Dynamic read error cross-sections: (a) Micron; (b) Hynix.

For the Dynamic Read condition, the parts showed exhibited transient read errors in addition to the bit and Page/Block errors, and other SEFIs, which are plotted in Fig. 4. For Ne and Ar ions (LET up to 9.74), there were no static bit errors at any LET, detected after the beam was turned off. There were transient read errors, as shown in Fig. 4, at these LETs, which are thought to be due to noise in the read circuit. At higher LETs (Cu, Kr, and Xe ions), SEFIs were observed on all shots for the Micron part, and many shots for Hynix parts, which made it difficult to count the transient errors—if

the DUT stops responding to commands, it is hard to say what errors were not counted. And most of the errors that were counted, were due to control logic, and not associated with individual bits. After the DUT was restored to operating condition, it could still be checked for static bit errors, however. The transient results are shown in Fig. 4. Static results are not shown in Fig. 4, but they are consistent with results in Fig. 3. In Fig. 4, we have attempted to count SEFI events, despite the obvious difficulties of doing so. For example, one can count block errors, but it is often unclear whether multiple events are independent or not. We have assumed that block or page errors at widely separated addresses are independent events, and block or page errors at consecutive addresses are one event. Of course, if the DUT stops responding to commands, there may be other events that were missed completely. The number of SEFI events is small in any case. As always, the statistical uncertainty associated with rare events is large.

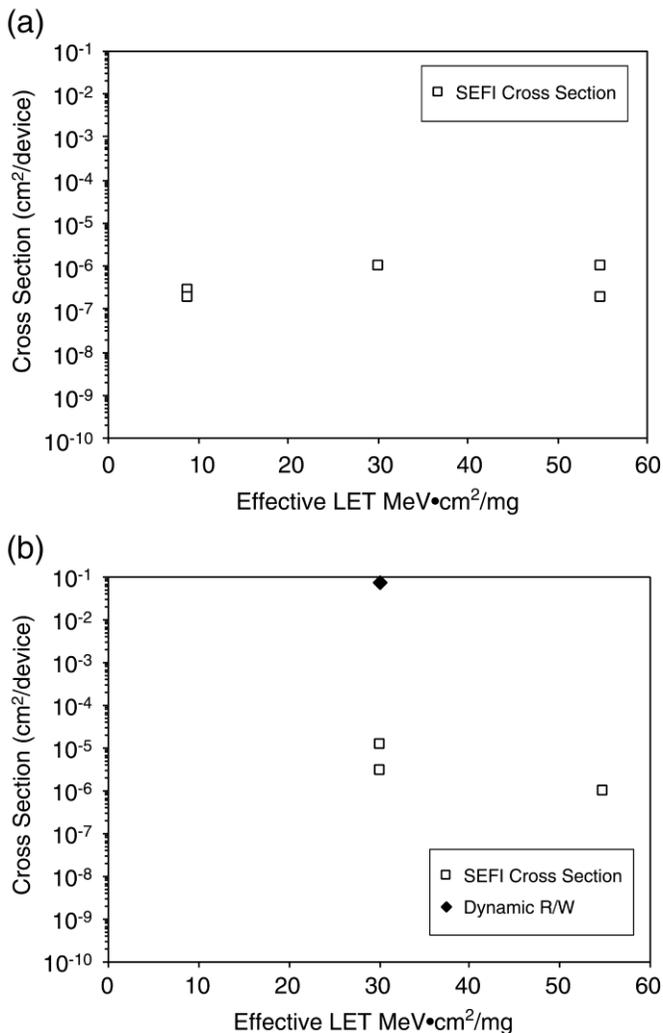


Fig 5. Dynamic R/W error cross-sections: (a) Micron; (b) Hynix.

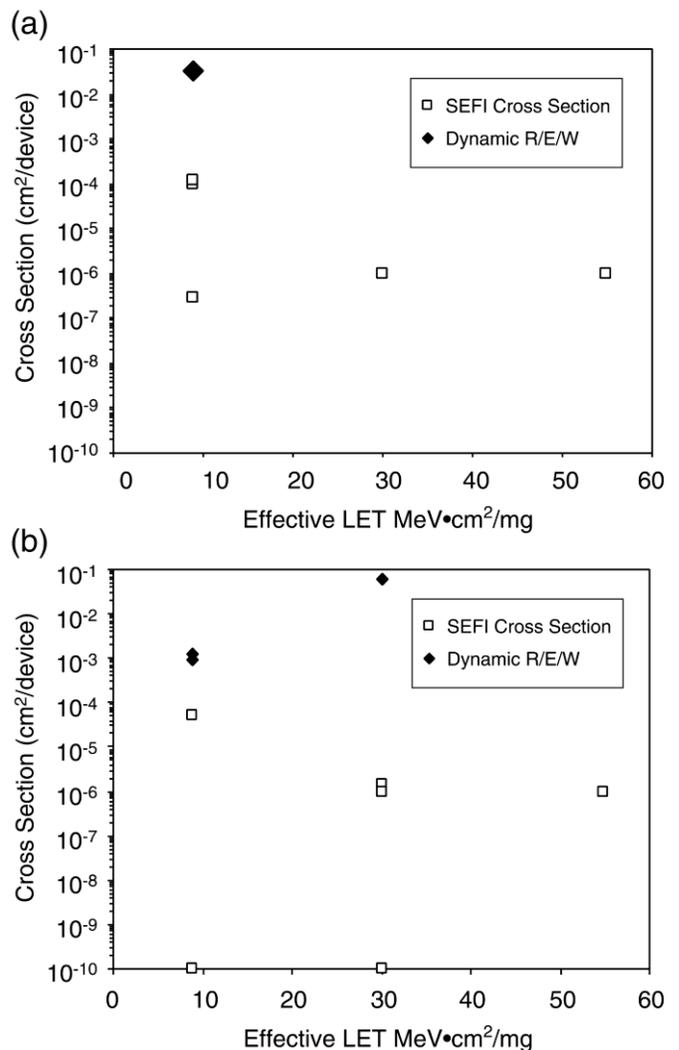


Fig 6. Dynamic R/E/W error cross-sections: (a) Micron; (b) Hynix.

Results of the dynamic R/W tests are shown in Fig. 5. Generally these results are unremarkable, because the usual zero-to-one errors are rewritten as they occur. For this reason, there are fewer errors indicated than in Fig. 4. 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, but this appears not to have happened, at least not on a large scale. Probably, this is because the write operation is performed only when a zero-to-one error is detected. For this circuit, four thousand such errors are still only one part per million of the entire memory, so the write circuit duty cycle is a very small number. Where a static cross section is given, it is based on the number of errors detected after the exposure and resetting of the DUT. The transient cross section is based on errors detected during the exposure. But some of the transient errors are probably really static bit errors that were rewritten during the test.

Results for the dynamic R/E/W tests are shown in Fig. 6, to the extent that they can be determined. For this condition, there were many more SEFIs than in the R/W (without erase) mode, which is probably due to the fact that every block is erased and rewritten on every cycle, so that the duty cycle for

high voltage operations is much higher. As a practical matter, there are many page and block errors, which usually appear to be independent, on almost every shot with LET at or above 9.74 (Ar). With many large chunks of the memory completely knocked out, it becomes impossible to determine static or transient errors affecting only single bits.

The Micron and Hynix parts were tested for SEL, Single Event Latchup, along with the Micron 4G, which had been tested previously for other SEE. Latchup testing was done at 70° C, and 3.6 V, which are the maximum rated operating voltage and temperature for all of these parts. Heating was accomplished with a strip heater inserted under the chip package in the test socket. Thermal contact was achieved with conducting grease. No latchup was observed for any part with any incident ion. Power supply current was monitored throughout the test, and some current increase was observed on nearly every shot. Many of these events were bus contention, where the DUT corrected itself during the exposure, without operator intervention. There was no case of a high current condition, requiring a power cycle to restore normal operation, which would have been the signature of a true latchup.

In Table 2, we show Weibull parameters for the static bit error cross-section along with the Crème 96 calculated error rate for geosynchronous orbit at solar minimum. Although there is some difference in response between manufacturers, they all have excellent results compared to advanced volatile commercial memories, which rarely achieve error rates better than  $1 \times 10^{-9}$  errors/bit-day. The largest error rate here is for the Hynix part, which corresponds to about three errors per year for a 4G. We note that these error rates are obtained without error correction, even though all the manufacturers recommend the use of error correction. With properly implemented error correction, the bit error rate for all three of these parts would be too low to ever measure.

The SEFI (Single Event Functional Interrupt) rate is of greater concern for space applications than the bit error rate, however. Typically, a SEFI occurs when a control circuit malfunctions as a result of a single ion interaction, and the entire memory, or a large part of it, fails. This may mean reloading the entire memory from a backup, or even rebooting the entire system. Although the Micron part had the best bit error response, it was the one most sensitive to SEFIs. For example, there was one point in the test where there were eight consecutive exposures of the Micron part with Ar ions (LET=8.7), which is a relatively low LET. Fluence on these exposures varied from  $10^3$  -  $10^4$  ions/cm<sup>2</sup>. There was a SEFI on seven of the eight exposures. Fluence was  $10^4$  particles/cm<sup>2</sup> on the one shot with no SEFI. Therefore, one can conclude that a SEFI occurs, on the average, every few thousand incident particles at this LET or higher. According to the Crème 96 input spectrum, the flux at LET = 9 or greater is  $1.9 \times 10^{-5}$  particles/cm<sup>2</sup>-sec. If the interval between SEFIs is about one thousand particles, this is equivalent to a SEFI every one-to-two-years. Probably, the actual interval,

per chip, is a few times greater. Depending on the number of chips in the system, this may, or may not, be a low enough rate to be considered acceptable for a system. But the SEFI cross-section for the Samsung part is about  $10^{-6}$  cm<sup>2</sup>, at the LET of Xe. At this LET, the flux is less than one particle/cm<sup>2</sup> per hundred years, so the expected SEFI rate per chip would be about one every  $10^8$  years. Clearly, this is good enough for most systems. The Hynix SEFI rate is intermediate.

TABLE II  
WEIBULL PARAMETERS FOR STATIC BIT ERRORS, CREME96 ERROR RATES,  
AND FIGURE OF MERIT

	Hynix 4G	Micron 4G	Samsung 4G
Threshold (LET)	1	3	3.5
Width (LET)	29	55	30
Exponent	4	5	5
Saturation (cm <sup>2</sup> /bit)	$5 \times 10^{-11}$	$4.7 \times 10^{-12}$	$4.5 \times 10^{-11}$
Geosynch. error rate (e/b-d)	$2.3 \times 10^{-12}$	$1.5 \times 10^{-16}$	$3.7 \times 10^{-13}$
Figure of Merit (FOM)	$1.1 \times 10^{-13}$	$2.2 \times 10^{-15}$	$6.2 \times 10^{-14}$
Error rate (40K ft)	$3 \times 10^{-13}$	$5 \times 10^{-15}$	$1 \times 10^{-13}$
Error rate (sea level)	$3 \times 10^{-15}$	$5 \times 10^{-17}$	$1 \times 10^{-15}$

Once the heavy ion cross section is known, the error rate can be estimated for other environments, using the FOM (Figure of Merit), as defined by Petersen [1,2]. He defines the FOM in terms of the Weibull parameters used to fit the cross section curve. Petersen himself established the correlation between FOM and bit error rate for heavy ions [1], and protons [2]. Normand [4] established the correlation between bit error rate and FOM for atmospheric neutrons at aircraft altitudes. He calculated expected avionics error rates for several parts using both the Bendel Two Parameter model, and the Burst Generation Rate model. The models agreed reasonably well with each other, and the error rates correlated well with the FOM. Then he did a least squares linear fit (on a log-log plot), which is shown in the top-right part of Fig. 7. If we simply extend his least squares fit, and plot points for these flash memories at the appropriate FOM, we get the results shown in the lower left part of Fig. 7, for parts operating at 40K ft. Although Fig. 7 looks unusual, it is necessary to plot the results in this way, because the error rate for flash memories is so much lower than for conventional volatile memories. The expected error rate is also given in Table 2.

According to Ziegler [5], the neutron spectrum at sea level is very similar to the spectrum at aircraft altitudes, but the flux at sea level is reduced by a factor of about 100. If we apply this correction to the aircraft altitude results, we get the sea level results, also indicated in Table 2.

We note that the flash cross sections used in this analysis are several orders of magnitude smaller than for any of the parts in Normand's original study. Therefore, the FOM and predicted error rates are also orders of magnitude smaller. For this reason, the procedure we have outlined here involves a significant extrapolation from previous work. However, the only assumption made in determining the FOM is that a Weibull curve can be fitted to the measured cross-section.

Clearly, this assumption is not violated here. However, it has been pointed out [6] that the Weibull curves in Table 2, do not reach a true saturation, which may introduce some uncertainty into the FOM calculation. While this is true, we have attempted to estimate the magnitude of the effect by

making plausible assumptions of the “true” saturation, and the uncertainty seems to be a factor of two or less. On a log-log plot like Fig. 7, covering nine decades on each axis, this difference would be barely perceptible.

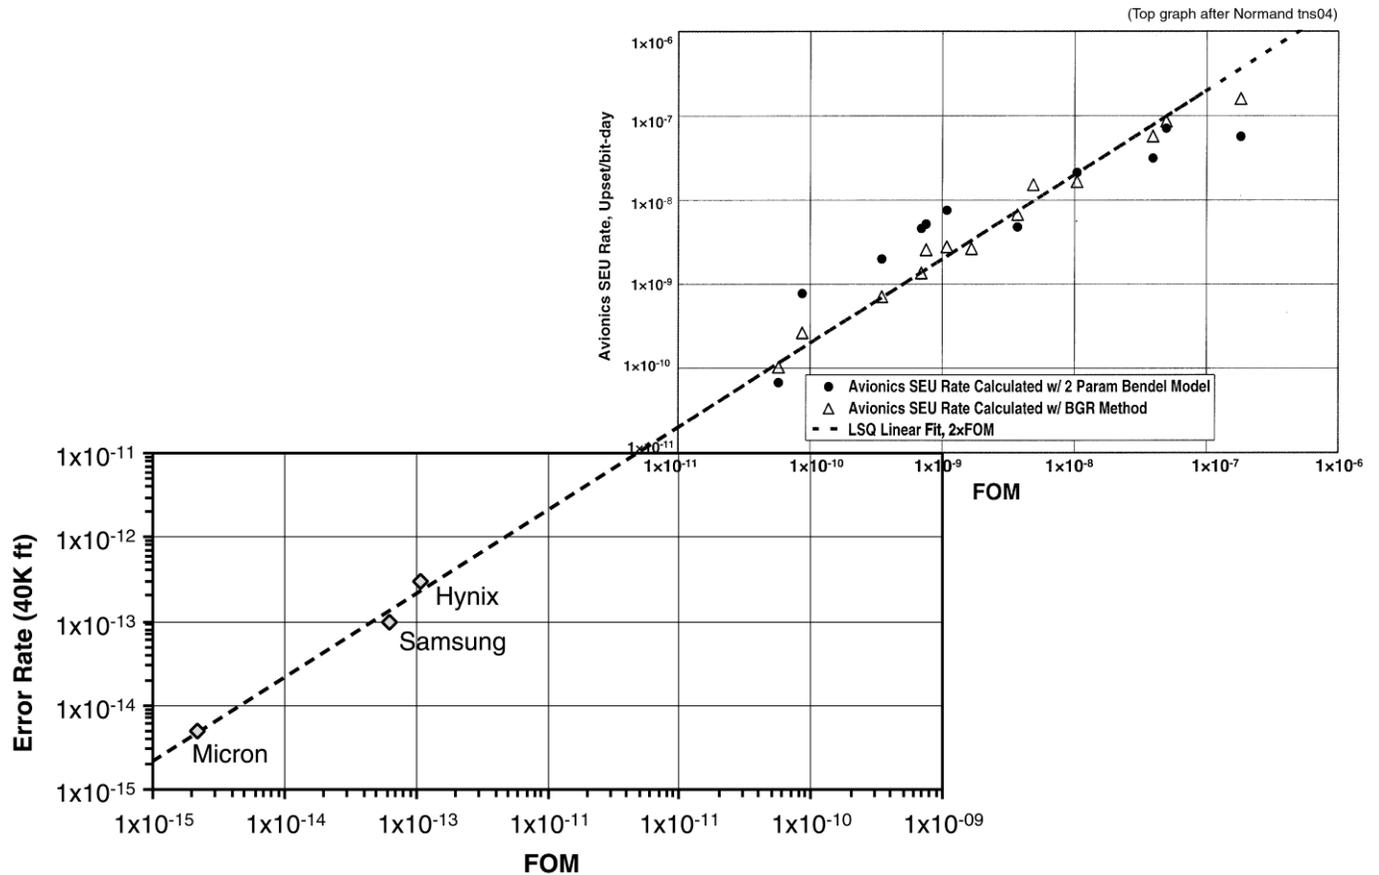


Fig. 7. Floating Gate Technology Radiation Damage.

### VI. CONCLUSIONS

Typically, NASA systems require TID hardness to 50 krad (SiO<sub>2</sub>), with a 2x margin highly desirable. At high dose rate, the Samsung part was the only one of the three to achieve this margin, although all met the basic requirement. SEE response is generally excellent for all flash products, with error cross-sections orders of magnitude lower than for standard volatile memories. None of the parts suffered SEL, but there were other destructive effects, usually failure of the erase circuit. The SEFI rate is also a concern with some flash memories.

### VII. ACKNOWLEDGMENT

The authors wish to thank Martha O’Bryan for technical assistance in the preparation of this manuscript. The authors also wish to thank Michael McLelland of Southwest Research Institute, who supplied some of the samples used in this work, and Steve Weinger of Samsung Semiconductor, who supplied the single chip 8G samples.

### VIII. REFERENCES

- [1] E.L. Petersen, J.B. Langworthy, and S.E. Diehl, Suggested single event figure of merit, *IEEE Trans. Nucl. Sci.*, vol. NS-30, no. 6, pp. 4533-4539 (1983).
- [2] E.L. Petersen, The SEU figure of merit and proton upset calculations, *IEEE Trans. Nucl. Sci.*, vol. NS-45, no. 6, pp. 2550-2565 (1998).
- [3] T.R. Oldham, M. Friendlich, J.W. Howard, M.D. Berg, H.S. Kim, T.L. Irwin, and K.A. LaBel, TID and SEE response of an advanced Samsung flash memory, *IEEE Radiation Effects Data Workshop Record* (2007).
- [4] E. Normand, Extensions of the FOM method—proton SEL and atmospheric neutron SEU, *IEEE Trans. Nucl. Sci.*, vol. NS-51, no. 6, pp. 3494-3504 (2004).
- [5] J.F. Ziegler and H. Puchner, *SER—history, trends and challenges*, Cypress Press, 2004.
- [6] E.L. Petersen, private communication.