Ion-Induced Stuck Bits in 1T/1C SDRAM Cells

Larry D. Edmonds, Steven M. Guertin, Leif Z. Scheick, Duc Nguyen, and Gary M. Swift

***Abstract*--Radiation exposure of certain types of devices tends to stick bits, causing them to not be read out correctly after programming. Evidence of a linear trend in stuck bits in SDRAM memory cells is presented. This trend makes a cross section, as traditionally defined for single event effects (SEE), unambiguous. However, there is considerable part-to-part variations in the cross section.**

I. INTRODUCTION

Stuck bits created by heavy-ion or proton irradiation can be an important concern for some DRAMs. Because DRAMs are very susceptible to single event upset (SEU), they are often protected by an error detection and correction circuit (EDAC). The most common versions of EDAC will correct single upsets, but will not correct upsets when two or more occur in a common EDAC word and during a common EDAC cycle time. For EDAC to be effective, the device architecture is arranged so that multiple upsets created by a single particle hit are in different EDAC words. Given this arrangement, an observable error requires two or more ion hits within a common word and common cycle. A sufficiently short cycle time can produce an extremely small probability for an observable error, even when the raw (no EDAC) SEU rate is very large. That is, unless there are stuck bits. Stuck bits are not corrected by EDAC, so they persist from one cycle to the next (unless they anneal). If a word contains a stuck bit, a single SEU at any time in that word produces an observable error, i.e., the stuck bit effectively disables EDAC from the point of view of the afflicted word. The observable error rate from that word is the raw SEU rate for that word, which is typically many orders of magnitude larger than the observable error rate when EDAC is able to function. Calculations for some specific cases have indicated that stuck bits will be the dominant failure mechanism for some DRAMs.

Stuck bits in the memory elements of some devices, such as DRAMs and some SRAMs, are believed to occur from either of two types of mechanisms. One mechanism is single event gate rupture, while another is micro-dose [1]. The issue of stuck bits has been examined to some extent for SRAMs [2-8] and the TID radiation response of SDRAMs have similarly been studied [9-10]. The present paper is concerned with those stuck bits that are believed to be caused by micro-dose, as evidenced by a certain degree of annealing.

Manuscript received July 17, 2001. This work was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under Contract with the National Aeronautics and Space Administration, Code AE, under the NASA Microelectronics Space Radiation Effects Program (MSREP).

The Authors are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA.

Micro-dose is a hybrid between total ionizing dose (TID) and a single event effect (SEE). Like SEE (but unlike TID), the disturbances created by heavy-ion hits are spatially non-uniform. Like TID (but unlike SEE), damage created by different hits at the same location is expected to be cumulative. This distinction between micro-dose and SEE raises a question regarding the meaning of a cross section for stuck bits. The experimental definition of a device cross section, associated with a given ion, is taken here to be an increment of counts divided by an increment of fluence. For conventional SEE, this experimental definition gives an unambiguous result, in the sense that the cross section does not depend on previous irradiation history (assuming that TID does not significantly alter the characteristics of a device). For brevity, the term "linearity" will refer to the property that the device cross section, experimentally defined as above, is independent of irradiation history. An equivalent definition of linearity is that the number of counts is proportional to fluence. However, it is not obvious whether linearity applies to stuck bits, so there is a question as to whether traditional SEE rate calculation methods can be used to estimate stuck bit rates in a space environment. In particular, Dufour *et al.* [11] reported that data obtained from a particular SRAM exposed to large LET (>90 MeV-cm^2/mg) heavy ions appear consistent with the postulate that stuck bits are caused by two hits to the same transistor gate. If true, then linearity cannot be assumed when calculating rates in space. Oldham *et al.* [12] later suggested that most of these stuck bits were actually produced by single hits. However, this still leaves a question as to whether it is the single hits or the multiple hits that are most important at smaller LET. The objective of this paper is to experimentally determine, for a particular device (the Hyundai 16Mx4 SDRAM), the validity of linearity for LET values ranging from very small (0.375 MeV-cm^2/mg) to moderately large (37.9).

The experimental data in the next section show that linearity does apply to the tested device, for all tested LETs and with fluences up to the largest used in the tests (which are expected to exceed the fluences encountered in natural space environments). Therefore, for this device, cross sections can be measured, and rates in space can be estimated, using the same methods used for other types of SEE. However, additional tests and/or modeling efforts are needed to determine whether this conclusion is universal, or limited to special families of devices. Also, the phrase "same methods used for SEE" should be qualified, because the directional dependence of device susceptibility may or may not be typical of other types of SEE (e.g., single event upset). We attempted to quantify the directional dependence of the cross section but we were not successful (all data having enough quality to be shown in this paper were measured at normal incidence), so this is a subject for future work. Our primary concern was merely to answer the question of whether cross section is or is not a function of irradiation history.

II. Experimental Data and Observations

The experimental method used for the first series of measurements starts with one device and exposes it to only one type of heavy ion (one LET value). The device is in an inverse bleed-down bit pattern, which results in all memory elements being susceptible to soft errors as discussed in [13]. This is a habit developed from soft error tests of DRAMs. It is sometimes not known whether it is the "1" or "0" at the device output that corresponds to the charged state of a memory capacitor, so by using this pattern (or its inverse) the charge state of all memory capacitors becomes known. The pattern used for these tests puts all memory capacitors in the charged state. The detection method could only count events in which the capacitor became stuck in the discharged state. It was thought that all stuck bits would be in this state, but we did not experimentally determine whether any bits were stuck in the opposite state. The device is biased at 3.3 volts. The fluence is applied in increments, and the number of stuck bits is recorded after each fluence increment. For experimental convenience, most tests monitored only a fraction of the bits in the device. The data are used to plot the cumulative number of stuck bits as a function of the cumulative fluence. The test is then repeated, using a fresh device and a different LET. This produces the curves in Fig.1, which also indicates the fraction of the device that was measured for each test. A similar test method, but using 200 MeV protons instead of heavy ions, produced Fig. 2 (the entire device was measured). One exception to this experimental method is the test represented in Fig. 1c. This device was pre-exposed to carbon (to produce Fig. 1b), producing an initial (prior to the fluorine test) number of stuck bits (which is slightly smaller than the number at the end of the carbon test because some annealing occurred). Another exception is the test represented in Fig. 1e. This device was also pre-exposed to carbon. The carbon test results are not shown because the test was terminated after only one data point. Without attempting to suggest implications from this, we simply point out that Co-60 tests performed on a separate set of devices found that the devices remained functional for TID levels (SiO_2) up to 60 krads (two devices) and 75 krads (a third device).

The most striking observation from Figs. 1 and 2 is that a strong linear relationship is seen in all cases, which include very small LET up to moderately large LET. We suspect that linearity would continue to larger LET, but this was not tested. The slopes of the curves in the figures were used to calculate cross sections. The cross sections were normalized to account for the fraction of the device that was tested, so the final cross sections represent a complete device. The 200 MeV proton cross section was found to be 1.4×10^{-9} cm^2/device. The heavy-ion cross sections are shown in Fig. 3.

Fig. 3 leads to another observation, which is part-to-part variations, as seen by comparing different points (which are from different devices) at the same LET. The fact that each device shows a strong self-consistency (i.e., linearity with minimal scatter in Fig. 1) indicates that variations between devices are true part-to-part variations, rather than an artifact of experimental scatter. Note that the points showing this variation are at LET values at which the cross section rises rapidly, i.e., the cross section is very sensitive to LET, so we might not be surprised to see some scatter here. But again, the absence of scatter in Fig.1 indicates that this is not just a sensitivity issue, there is a real part-to-part variation. Note that most of the tests looked at only a portion of the device. There might be variations between different portions of the device, but we did not investigate this.

The part-to-part variations obscure the shape of the cross section curve in Fig. 3. To see the shape of the curve, at least for one device, we used a more conventional test method that exposes the same device to different ions. Linearity has already been established from the first series of measurements, so the use of this conventional method is now believed to be meaningful. The result is shown in Fig. 4. An observation from this figure is that the cross section gradually increases with increasing LET.

III. III. A Suggested Explanation for the Observations

Observations from the previous section are part-to-part variations, linearity, and a gradual increase in cross section with increasing LET. Oldham *et al*. [12] and Poivey *et al*. [14] have already suggested an explanation for the part-to-part variations. Their work also provides a suggested explanation for the other observations. The authors point out that the threshold adjust implant dose exhibits both microscopic variations and macroscopic variations. The authors present quantitative arguments indicating that the macroscopic variations can account for considerable part-to-part variations. The microscopic variations produce cell-to-cell variations. The authors (and others, e.g., [7]) point out that there is also another type of microscopic variation associated with ion hit location. Specifically, a hit at one location in a gate oxide can have a much greater tendency to produce a stuck bit than a hit at another location.

The two microscopic variations (or even one without the other) imply that we can define different areas in the device in terms of different susceptibilities. In particular, for a given LET, one area (which could be zero, but assume it isn't) is defined by the condition that one hit, by the selected LET, will stick a bit. Another part of the device can be partitioned into sub-areas having the property that two or more hits in a single sub-area, by the selected LET, are needed to stick a bit. This is consistent with the increase in cross section with increasing LET because these areas depend on LET (e.g., via an increasing number of contributing cells due to cell-to-cell variation). This is also consistent with linearity if we make the additional assumption that the areas compare in such a way so that a single hit to the first area is much more probable (for the fluence used in a test) than two or more hits to a common sub-area of the second type. If this assumption is correct, then nearly all contribution to the measured cross section is from the first area, so the fact that damage is cumulative is irrelevant to the cross section.

The above explanation of linearity relies on a very speculative assumption regarding the way various areas compare, so the explanation is not very convincing unless it can be shown to be quantitatively consistent with the data. Specifically, the suggested explanation states that the number of stuck bits created by double hits is less than the number created by single hits. To show consistency, we should estimate these two numbers so that we can verify (for each LET) that the first number really is smaller than the second number. It is important to recognize that the first number is not the total number of bits that received two hits. It is the number that became stuck as a result of receiving two hits. An analogous statement applies to the second number.

To estimate these numbers, we use two approximations that are probably not very accurate but may still be adequate for order of magnitude estimates. The first approximation ignores variations within a bit (all variations are between bits), so each bit can be assigned a well-defined sensitive area A which (for now) is taken to be the area of the gate oxide. The second approximation recognizes that the bits are not all identical in terms of susceptibility, but assumes that the collection of bits that are potentially able to become stuck from two hits at a given LET L (that is, they will become stuck if they actually do receive two hits at an LET L) is the same as the collection that are potentially able to become stuck from a single hit at an LET of $2L$. This approximation can be stated as

$$N_2^{(P)}(L) = N_1^{(P)}(2L) \qquad (1)$$

where the left side is the number of bits that are potentially able (hence the superscript) to become stuck from two hits at an LET L, and the right side is the number that are potentially able to become stuck from a single hit at an LET of $2L$. The right side can be estimated from the device cross section. The device area that is potentially able to produce a stuck bit from a single hit at an LET of $2L$ (that is, will produce a stuck bit if the area is actually hit) is $\sigma(2L)$ where σ is the cross section function. The right side of the above equation is this device area divided by the bit sensitive area A, i.e.,

$$N_1^{(P)}(2L) = \frac{\sigma(2L)}{A}. \qquad (2)$$

The expected number of bits that will actually become stuck from two hits at an LET L is the number that are potentially able to become stuck from two hits multiplied by the probability that a given bit will actually receive two hits. This gives

$$N_2(L,F) = N_2^{(P)}(L)\, P_2(F) \qquad (3)$$

where $N_2(L,F)$ is the expected number of stuck bits created by double hits when the fluence is F and the LET is L, and $P_2(F)$ is the probability of two hits from a fluence F. Combining the above equations gives

$$N_2(L,F) = \frac{\sigma(2L)}{A} P_2(F). \qquad (4)$$

The number that $N_2(L,F)$ must be compared to is $N_1(L,F)$, which is the expected number of stuck bits created by single hits when the fluence is F and the LET is L. In analogy with (2) and (3) we have

$$N_1^{(P)}(L) = \frac{\sigma(L)}{A}, \qquad N_1(L,F) = N_1^{(P)}(L)\, P_1(F)$$

where $N_1^{(P)}(L)$ is the number of bits potentially able to become stuck from a single hit, and $P_1(F)$ is the probability of receiving a single hit. Combining the last two equations gives

$$N_1(L,F) = \frac{\sigma(L)}{A} P_1(F). \qquad (5)$$

The probabilities are calculated from the Poisson distribution and are given by

$$P_1(F) = e^{-AF}\,(AF), \qquad P_2(F) = \frac{1}{2} e^{-AF}\,(AF)^2.$$

Combining the last equations with (4) and (5) gives

$$\frac{N_2(L,F)}{N_1(L,F)} = \frac{\sigma(2L)}{2\,\sigma(L)} A F. \qquad (6)$$

The gate oxide area is not precisely known, but is estimated by noting that the die area is about 0.25 cm^2, the cell area (summed over cells) is about half of the die area, and the oxide area is about one-tenth of the cell area. Therefore, the total (summed over cells) oxide area is estimated to be 0.0125 cm^2. There are about sixty four million cells, so the area A of one gate oxide is estimated to be about $2\text{x}10^{-10}$ cm^2 (or 0.02 μm^2). If we use Fig. 4 to compare cross sections at different LET values, we find that $\sigma(2L)/\sigma(L) \le 5$ for all L in the plotted range. Using these numbers, the fluence at which linearity is expected to be noticeably violated (which is the fluence that makes the left side of (6) equal to 1, and is calculated by setting the right side equal to 1) is about $2\text{x}10^9/cm^2$. Some heavy-ion test fluences slightly exceeded this value (up to a factor of three) and linearity was still observed. However, the analysis was based on approximations intended only for order of magnitude estimates. For example, Poivey *et al.* [14] point out that, at least for some devices, only a fraction of the gate area is sensitive. If we assume that the sensitive area A is only a portion of the oxide, we would predict linearity up to a larger fluence. Therefore, the suggested explanation for linearity, that single hits are more important than double hits, appears credible. At least there is no clear contradiction with the quantitative observations.

IV. A Lesson Learned

Some of the stuck bits anneal, so it should be mentioned that we believe that the data presented here are not significantly corrupted by annealing. The test time was not excessive, and the tests that produced these data avoided a problem that was seen in other data sets that are not shown. A lesson learned is that we should avoid a test sequence in which a given run produces a small number of counts (e.g., a low LET with a fluence comparable to that used for other runs) after a large number of stuck bits have accumulated from previous runs. Only a small fraction of the stuck bits anneal during a run, but a small fraction of a large number can be significant compared to the small number of stuck bits acquired from the run.

V. The Influence of Annealing on Rates in Space

The annealing of stuck bits in this SDRAM was found to be significant, so the prediction of the number of stuck bits accumulated in space should include this effect. We did not obtain annealing data having sufficient quality for presentation. However, we did establish linearity for the tested device, and this is enough to give some credibility to a suggested method of calculation. This section focuses on the method. Future work is needed to obtain the required input data for the device investigated here. It is believed that the method will be useful for those individuals that already have the required data for a device that is of interest to them and that exhibits linearity. It must be acknowledged, however, that the method discussed below has not been experimentally verified and is offered here only as a suggestion.

We first consider a function K that is experimentally defined as follows. A device is irradiated for a short time (short enough so that significant annealing does not occur during the irradiation) using a flux large enough to create a large number (large enough for statistical significance) of stuck bits. We then record the number of stuck bits that remain as a function of time after the irradiation has been turned off, and define $K(t)$ to be the fraction of the initial stuck bits that still remain after annealing for a time t. A fundamental assumption is that this same K also describes the probability that a selected stuck bit will anneal after a prescribed time even if the device is still being irradiated (e.g., in space) during this time. Linearity, which is assumed to apply to fluences at least as large as encountered in space, gives credibility to this assumption because linearity suggests that we can ignore those instances in which the same sensitive area is hit more than once. From the point of view of a selected stuck bit, its probability of recovering after a prescribed time is independent of whether the device is still being irradiated because this bit will almost certainly not be hit again anyway. Therefore, the same K measured in the laboratory should also be relevant to conditions in space, with some qualifications. The qualifications recognize that K will depend on various parameters, such as temperature. It presently seems plausible that K might even depend on the type of irradiation that created the stuck bits. If this is found to be true, then either of two approaches might be used. One approach is to look for a worst-case K that can be applied to all particle types for conservative estimates. Another approach is to treat different particle groups (characterized by species, energy or LET, and direction of travel relative to the device) independently and then sum results over particle groups. The analysis below assumes that either one function K represents (conservatively if not accurately) all particles, or we are considering a particular particle group with the intention of summing over groups later.

Before giving a general equation, we discuss the method via an illustration. In this illustration, a mission duration is four hours long, and the objective is to estimate the number of stuck bits present at the end of the four hours. First consider the stuck bits created during the fourth hour. Those that were created near the beginning of the hour had almost an hour to anneal, but a conservative assumption is that there was no anneal time for any of them. The contribution to the mission total from this set of stuck bits is conservatively estimated to be all of them. Now consider the stuck bits created during the third hour. Those created near the beginning of the hour had almost two hours to anneal, but a conservative assumption is that they all had a one-hour anneal time. The contribution to the mission total from this set is conservatively estimated to be the number created during the third hour multiplied by the fraction of stuck bits that remain after a one-hour anneal. Similarly, the contribution from those created during the first hour is conservatively estimated to be the number created multiplied by the fraction of stuck bits that remain after a three-hour anneal. Generalizing this illustration to a mission that is m hours long, a conservative estimate is given by

$$N(m\,\Delta t) = \Delta N(m\,\Delta t) + \sum_{i=1}^{m-1} K(i\,\Delta t)\,\Delta N(m\,\Delta t - i\,\Delta t)$$

where Δt is one hour, $N(m\,\Delta t)$ is the number of stuck bits at the end of the mission, $\Delta N(i\,\Delta t)$ is the number created during the i^{th} hour, and $K(i\,\Delta t)$ is the fraction of stuck bits that remain after annealing for i hours. The numerical conservatism can be removed by allowing Δt to be arbitrary (instead of one hour) so that we can take the limit as $\Delta t \rightarrow 0$. Multiplying and dividing the sum by Δt and then taking this limit gives

$$N(t_0) = \int_0^{t_0} K(t)\,R(t_0 - t)\,dt$$

where t_0 is the time of observation and $R(t)$ is the creation rate, at time t, of stuck bits. The creation rate is distinguished from the net rate that includes losses from annealing. The assumed linearity implies that the creation rate can be calculated from the customary (for single event effects) method that combines cross section data with environmental data.

A special case will probably apply often. This case applies when: (a) the environment can be adequately (or conservatively) approximated as constant in time for rate calculations, (b) virtually all stuck bits eventually anneal (the number of permanent stuck bits, which accumulate for all time, is negligible if not zero), and (c) the anneal time is short compared to the mission duration. For this case, the number of stuck bits approaches a limiting (or saturation) value. This is because almost all of the stuck bits were created in the recent past, so it does not matter whether the launch date was a moderately long time ago or a very long time ago. The saturation value is calculated by letting $t_0 = \infty$ in the above equation and factoring out the R (which is now constant) to get

$$N(\infty) = R \int_0^\infty K(t)\, dt \, .$$

If the integral is finite, the number of stuck bits approaches a saturation value. If not, there is no saturation, i.e., the number continues to increase regardless of the mission duration. If the integral is finite, we can define the characteristic anneal time t_c by

$$t_C \equiv \int_0^\infty K(t)\, dt$$

so the equation for N(∞) becomes

$$N(\infty) = R\, t_C \, .$$

If $K(t)$ is a decreasing exponential function of t (i.e., a semi-logarithmic plot of K versus t is a straight line), the t_c defined above is the same as the time constant in the exponential function. More generally, t_c is calculated from the above equation.

VI. CONCLUSIONS

Stuck bits can be created in the Hyundai 16Mx4 SDRAMs by protons and very low LET heavy ions, and can effectively disable EDAC. Three definite conclusions have been reached regarding the tested device. One is linearity over a broad range of LET values, even for very large fluences. This is important because it justifies the use of traditional SEE methods for cross section measurements and rate calculations in the absence of annealing. If there is annealing but it did not corrupt the cross section measurements, traditional methods are expected to be valid for calculating the creation rate of stuck bits. However, it is not yet known whether linearity over such a broad range of LETs is typical or a property of special families of devices, so additional work is needed. Also, we did not successfully determine whether the directional dependence of the cross section is typical of some other types of SEE, so this is another subject for future work. In the meantime, some worst-case assumption (e.g., the cosine law) might be used for conservative estimates. The second conclusion, which is not a surprise in view of work previously done by others [12], [14], is that there is considerable part-to-part variation. This is important because it implies that lot testing is needed. The third conclusion is a lesson learned. Cross section data are most corrupted by annealing when the accumulated number of stuck bits prior to a run is large while the number acquired during the run is small. The corruption is least when the sequence of LETs and fluences are selected so that each run creates more stuck bits than the previous run.

Other suggestions were given, but they are speculative. It was argued that the dominance of single hits over multiple hits might be consistent with the quantitative data (at least there is no clear contradiction), so this presently appears to be a credible explanation for linearity. A method for including annealing when estimating the number of stuck bits obtained in space was suggested but not tested against observations. Also, additional work is needed to obtain the annealing data needed for the calculation.

VII. REFERENCES

[1] G. M. Swift, D. J. Padgett, and A. H. Johnston, "A new class of single event hard errors," *IEEE Trans. on Nucl. Sci.*, vol. 41, no. 6, pp. 2043-2048, Dec. 1994.

[2] R. Koga, W. R. Crain, K. B. Crawford, D. D. Lau, S. D. Pinkerton, B. K. Yi, and R. Chitty, "On the suitability of non-hardened high density SRAMs for space applications," *IEEE Trans. on Nucl. Sci.*, vol. 38, no. 6, pp. 1507-1513, Dec. 1991.

[3] J. P. David, J. G. Loquet, and S. Duzellier, "Heavy ions induced latent stuck bits revealed by total dose irradiation in 4T cells SRAMs," *RADECS 99,* pp. 80-86, 2000.

[4] J. L. Barth, J. W. Adolphsen, and G. B. Gee, "Single event effects on commercial SRAMs and power MOSFETs: final results of the CRUX flight experiment on APEX," *IEEE Radiation Effects Data Workshop*, pp. 1-10, 1998.

[5] M. A. Xapsos, "Hard error dose distributions of gate oxide arrays in the laboratory and space environments," *IEEE Trans. on Nucl. Sci.*, vol. 43, no. 6, pp. 3139-3144, Dec. 1996.

[6] B. F. Cockburn, and Y.-F. N. Sat, "A transparent built-in self-test scheme for detecting single V-coupling faults in RAMs memory," *Records of the IEEE International Workshop on Technology, Design and Testing*, pp. 119-124, 1994.

[7] R. Gaillard, and G. Poirault, "Numerical simulation of hard errors induced by heavy ions in 4T high density SRAM cells," *IEEE Trans. Nucl. Sci.*, vol. 41, no. 3, pp. 613-618, June 1994.

[8] S. Duzellier, D. Falguere, and R. Ecoffet, "Protons and heavy ions induced stuck bits on large capacity RAMs," *RADECS93*, pp. 468-472, 1994.

[9] B. G. Henson, P. T. McDonald, and W. J. Stapor, "SDRAM space radiation effects measurements and analysis," *IEEE Radiation Effects Data Workshop*, pp. 15-23, 1999.

[10] S. Bertazzoni, G. C. Cardarilli, D. Piergentili, M. Salmeri, A. Salsano, D. Di Giovenale, G. C. Grande, P. Marinucci, S. Sperandei, S. Bartalucci, G. Massenga, M. Ricci, V. Bidoli, D. De Francesco, P.G. Picozza, and A. Rovelli, "Failure tests on 64 Mb SDRAM in radiation environment defect and fault tolerance in VLSI systems," *DFT '99 International Symposiumon*, pp. 158-164, 1999.

[11] C. Dufour, P. Garnier, T. Carriere, J. Beaucour, R. Ecoffet, and M. Labrunee, "Heavy ion induced single hard errors on submicronic memories," *IEEE Trans. on Nucl. Sci.*, vol. 39, no. 6, pp. 1693-1697, Dec. 1992.

[12] T. R. Oldham, K. W. Bennett, J. Beaucour, T. Carriere, C. Polvey, and P. Garnier, "Total dose failures in advanced electronics from single ions," *IEEE Trans. on Nucl. Sci*., vol. 40, no. 6, pp. 1820-1830, Dec. 1993.
[13] S. M. Guertin, L. D. Edmonds, and G. M. Swift, "Angular dependence of DRAM upset susceptibility and implications for testing and analysis," *IEEE Trans. on Nucl. Sci*., vol. 47, no. 6, pp. 2380-2385, Dec. 2000.
[14] C. Poivey, T. Carriere, J. Beaucour, and T. R. Oldham, "Characterization of single hard errors (SHE) in 1M-bit SRAMs from single ion," *IEEE Trans. on Nucl. Sci*., vol. 41, no. 6, pp. 2235-2239, Dec. 1994.

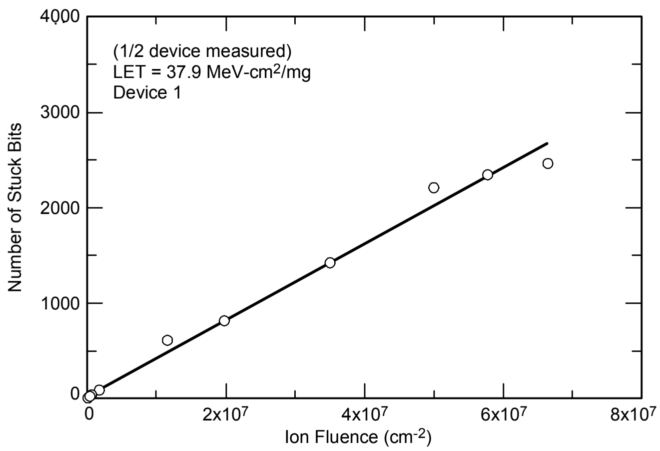

Fig. 1a: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Xenon Ions)

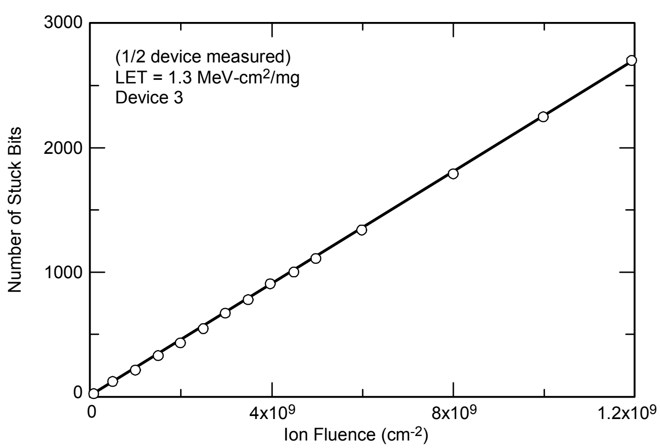

Fig. 1b: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Carbon Ions)

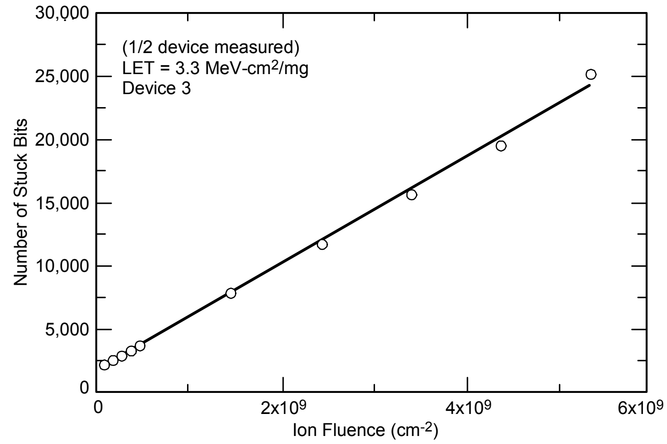

Fig. 1c: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Fluorine Ions)

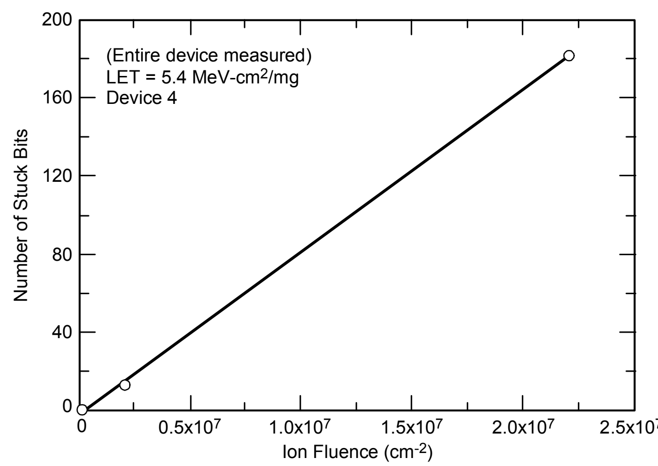

Fig. 1d: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Argon Ions)

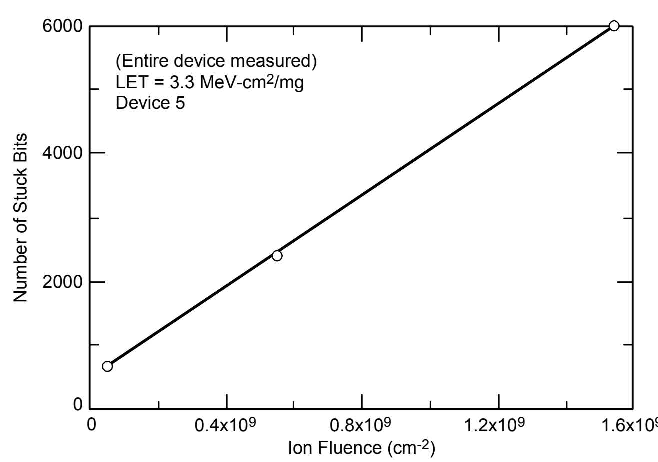

Fig. 1e: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Fluorine Ions)

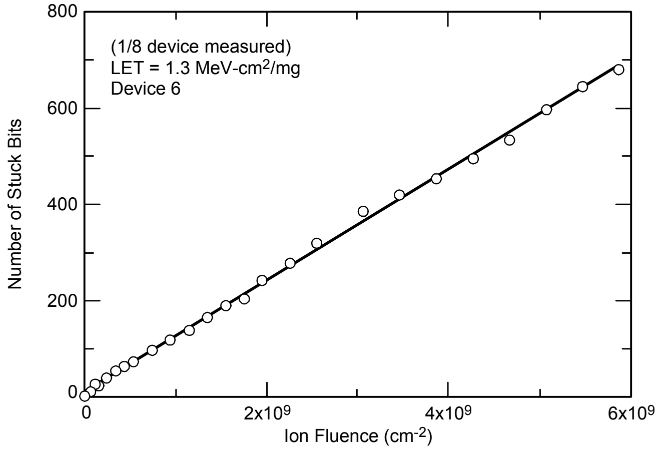

Fig. 1f: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Carbon Ions)

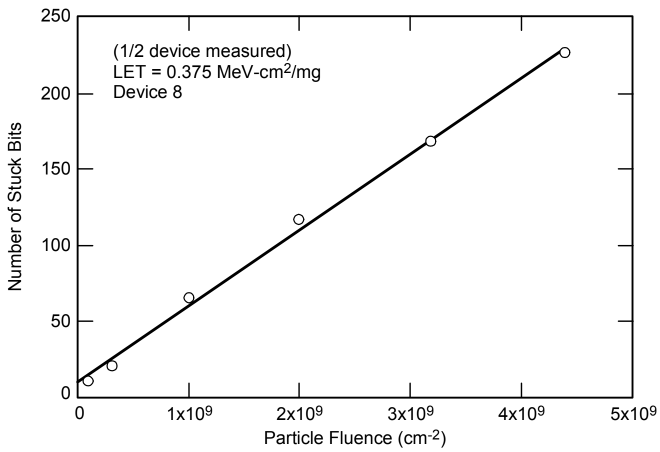

Fig. 1g: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (Lithium Ions)

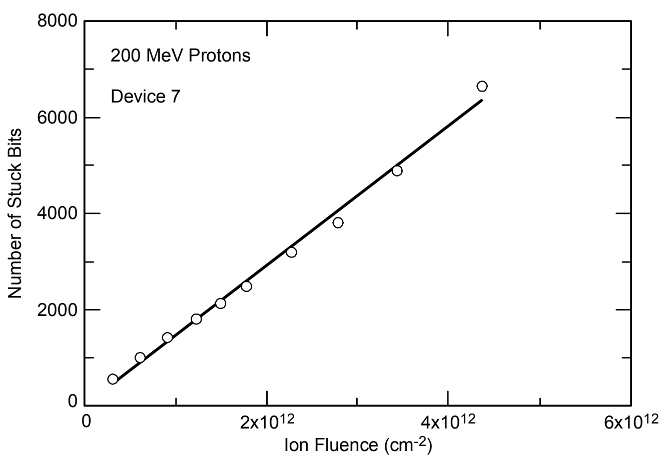

Fig. 2: Stuck Bits vs. Ion Fluence for the Hyundai 64-Mb SDRAM (200 MeV Protons)

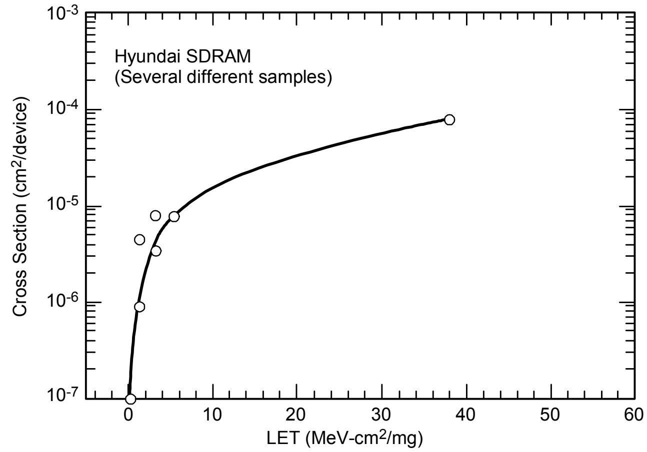

Fig. 3: Cross Section Data from Different Samples of the Hyundai 64-Mb SDRAM

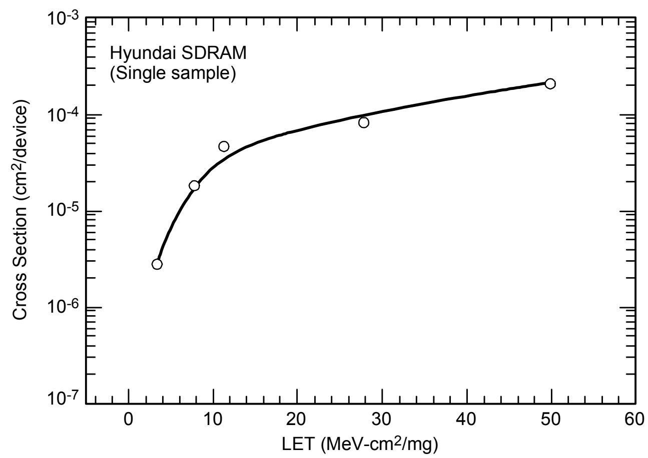

Fig. 4: Cross Section Data from One Sample of the Hyundai 64-Mb SDRAM