Heavy Ion SEE Test Report for the Current Spike Experiment

Tim Oldham, Dell-Perot Systems, Inc. /NASA-GSFC
Mark Friendlich, MEI Technology Inc./NASA-GSFC
Farokh Irom, NASA-JPL

Test Date: 7-8 Dec 2010 Report Date: 24 Jan 2011

I. Introduction

This study was undertaken to resolve differences in test results reported by different investigators, concerning high current spikes. Specifically, the occurrence of destructive high current spikes has been reported in SEE testing of advanced flash memories [1, 2], but not confirmed in other tests [3]. Although the same, or similar, parts were used in these different tests, the test equipment was different, and beam conditions were also quite different. These spikes are described in [1] as less than 400 ms in duration, and as destructive at high LETs for some manufacturers, although the spikes occur for all manufacturers, and also at lower LETs. The purpose of this test was to match as nearly as possible experimental conditions where spikes had been reported, and, if spikes were confirmed, to investigate possible mechanisms that might cause them.

II. Devices Tested

We tested a total four Samsung parts, and five Micron parts. The Samsung parts were 8G NAND flash, part number K9F8G08U0M, Lot Date Code (LDC) 807. The Micron parts were 4G NAND flash, part number MT29F4G08AAAWP, LDC 0748. Both of these parts were tested in [1] or [2], and the parts were selected for this test because we were trying to duplicate those results.

The device technology is 73 nm minimum feature size for the 4G parts and 60 nm for the 8G parts. All the parts are single die, SLC (single level cells). The chips came in a 48-pin TSOP package, but the plastic had been dissolved on the topside to expose the chips, allowing the beam to reach the chip surface. The 4G Micron parts had 4096 blocks, with block size 128K x 8. The blocks had 64 pages, each 2K x 8. For the Samsung 8G parts, the page width was doubled, to 4K x8, but still 64 pages/block and 4096 blocks. In both cases 80 bad blocks were allowed, although the actual number is usually much less, in our experience.

III. Test Facilities and Beam Conditions

Facility: Texas A&M University Cyclotron
Flux: (3×10^2 to 1×10^5 particles/cm^2/s).
Fluence: All tests were run to 1E3 to 1E7 p/cm^2, or until destructive or functional events occurred.

In the Dec 2010 test, only Au and Xe ions were used. This was done because we were trying to duplicate beam conditions in [1,2], in order to duplicate results. Fluence was usually intended

to be 10^7 ions/cm^2, but many runs were terminated early because the DUT stopped working. Flux was initially 10^4 ions/cm^2-sec, or higher to keep run times tractable. However, one of the goals of the study was to look for a flux dependence, because the suggestion had been made that the current spikes could be due to multiple ion hits. Therefore, shots at 10^3 ions/cm^2-sec and even 3 x 10^2 ions/cm^2-sec were also included in the test. Another goal of the test was to identify the LET dependence of the current spikes, if they were observed. Time did not allow any lower LET ions to be used in the Dec run, but the Micron 4G had been tested extensively at lower LETs for the MMS program. SEE results have been reported previously [3], but the current traces were not reported, because they were considered unremarkable at the time. They will be included in this report, however. The ions listed in Table I reflect both the Dec 2010 test, and the earlier MMS test.

Table I: Ions/Energies and LET for these tests.

TAMU Ions	Energy/ AMU	Energy (MeV)	Approx. LET on die (MeV•cm^2/mg)	Angle	Effective LET
Ne	15	300	2.7	0	2.8
Ar	15	600	8.4	0, 45	8.4, 11.8
Kr	15	1260	30.1	0	29.3
Xe	15	1965	54.8	0, 50	54.8, 84
Au	15	2955	87.5	0, 50	87.5, 136

IV. Test Conditions

Test Temperature: Room Temperature in the December 2010 test, although limited SEL testing was done at 70° C in the earlier MMS test

Operating Frequency: (0-40 MHz).

Power Supply Voltage: (3.3V for SEU, 3.6V (3.3+10%) for SEL). Standard test methods for SEU testing (e.g., ASTM 1192) call for testing at nominal voltage less 10%, because SEU in standard volatile memories is caused by voltages being pulled down. However, flash memories are designed to retain information even at zero volts, so the upset mechanisms are clearly different, here. In addition, we are also looking for control logic errors, which are thought to be due to things turning on when they are not supposed to be on. Reduced voltage would cause an underestimate of the rate for these events. Therefore, we used nominal voltage, 3.3 V, in all tests except latchup tests, which were done at 3.6 V, and also at elevated temperature.

V. Test Methods

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.

Generally, 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 all of the testing reported here 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, the part is normally power cycled and then read to determine static errors from the previous shot. Then it is erased, and read to verify that the erase was successful. Then the AA pattern is rewritten, and read to verify the write was successful. 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 irradiation, with bias applied, in which a pattern was written and verified, and the part was irradiated. Once the irradiation reached the desired fluence, it was stopped, and the memory contents were read and errors tallied. Normally, we also do a similar test with no bias, but this was not done here because the focus was current spikes, which could not happen with no bias.

2) Dynamic Read, in which a pattern was written to memory and verified, then subsequently read continuously during irradiation. 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.

3) Dynamic Read/Erase/Write, which again was similar to the Dynamic Read and Read/Write, except that a word in error was first erased and then rewritten. In this mode, the words that are read are compared to an "expected" pattern, which is actually the complement of the stored pattern. For this reason, every word is erased, as if it were in error. Because the Erase and Write operations use the charge pump, it is expected that the Flash could be more vulnerable to destructive conditions during these operations.

4) Latchup testing, was not done in the Dec 2010 test, but was conducted at 70° C, and 3.6 V in a few cases in the earlier test. It was expected that high voltage, dynamic test modes would be most likely to result in latchup, so these were emphasized in the latchup testing.

5) 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 at 50 degrees in the Dec test, which was close to the maximum possible angle, because the socket would have blocked the beam at angles much higher.

The Block diagram for control of the DUT is shown in Figure 1. The FPGA based controller interfaces to the FLASH daughter card and to a laptop, allowing control of the FPGA and uploading of new FPGA configurations and instructions for control of the DUT. Power for the flash is supplied by means of a computer-controlled power supply. The National Instruments Labview interface monitors the power supply for over-current conditions and can shut down power to the DUT if desired. However, the goal of this study was to study high current events, so the current limit was usually set very high in this run. The use of this FPGA-based test system is the main difference between this test and the earlier tests, describe in [1,2].

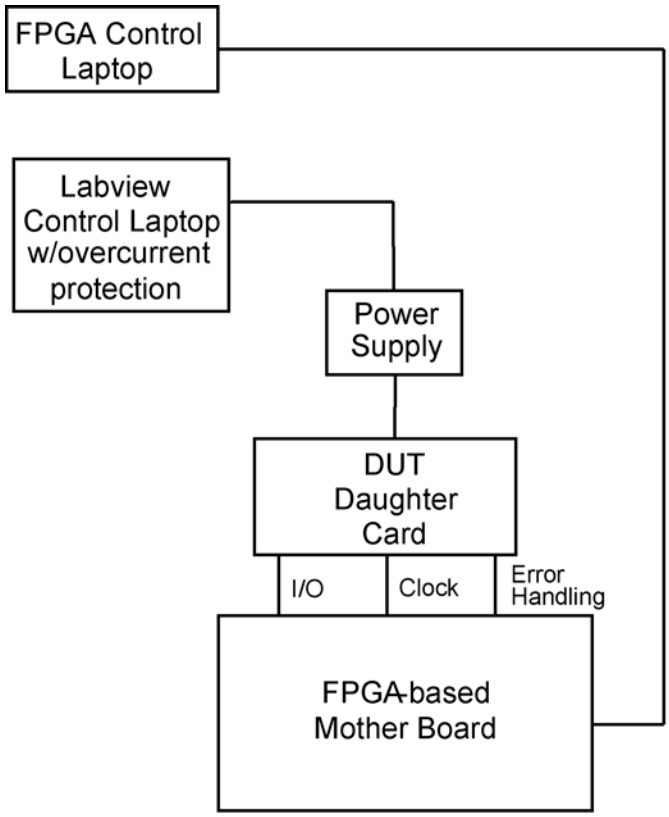

Fig. 1. Overall Block Diagram for the testing of the NAND Flash

VI. Results

In the December test, there were 35 exposures of different samples to either Xe ions or to Au ions, which resulted in numerous high current events. For definiteness, we will define a current level above 40 mA as a high current event. Although this level is somewhat arbitrary, it is roughly 50% above any current level that should be encountered in normal operation, and it is about 50% below the typical current spike level reported in [1,2]. By this definition, there were 47 high current events in this test run. In addition, some other events are considered to be high current, even though the current does not reach 40 mA. For example, in static mode, a Micron DUT should draw about 1 mA or less. If the actual current reaches 20 or 25 mA, it exceeds nominal current by more than an order of magnitude, and so is considered to be “high”. A Samsung DUT in static mode draws, nominally, 3-4 mA, which is also well below the 20-25 mA level. In all, then, we consider a total of 52 high current events.

The vast majority of these events appeared not to be current transients, but rather were changes in the DC current level. That is, the current would increase abruptly, and then be stable

at the new, higher value for an interval that varied, but was typically a few seconds to several minutes. Consider, for example, the example in Fig. 2, which has two high current excursions. (There is actually also a third event above 40 mA, but it occurred after the beam was off. The DUT had failed, and it exhibited an abnormally high current when we tried to reset it for the next exposure, but this event is not counted.) For both cases that are counted, the current level increased in one time step to a higher level, where it was stable for several seconds, before increasing again, to a second stable level. For the first current excursion, the interval from baseline to baseline was 37 sec, compared to 57 sec for the second excursion.

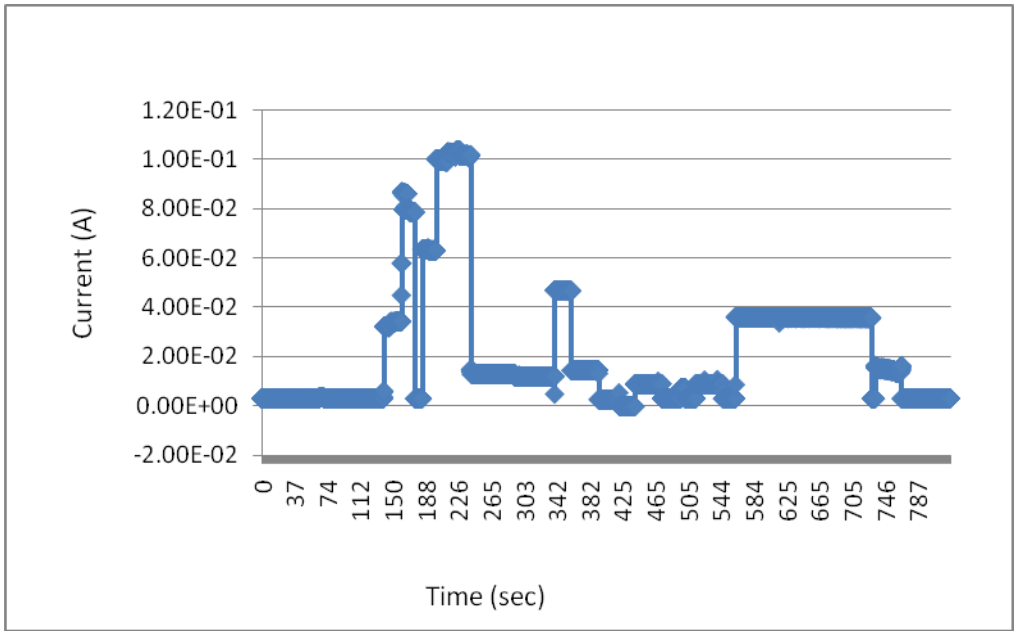

Fig 2. Samsung DUT 1, Run 1—Static mode. Fluence was 3 x 10^6 Xe ions/cm^2. Both erase and write functions failed. In the first current excursion, the current jumps from a few mA to about 32 mA, where it stays for about 21 sec, before increasing again to 80 mA or more for another 16 sec. Total duration is 37 sec. For the second excursion, the current jumps to more than 60 mA, for about 16 sec, then to about 100 mA for another 41 sec. Total duration is 57 sec.

Similarly, for Fig. 3, which is a shot in Dynamic Read mode, when the beam is turned on, the DUT starts to read, and the current increases to 10 mA, which is just the normal Read current. After 55 sec, the current jumps to about 38 mA, where it stays for 21 sec, before jumping again to almost 60 mA. The current stays at this value for about 32 sec, Then the current drops in three steps back to the quiescent level, 4 mA. Total duration of the interval where current was above the normal Read current level was more than two minutes. This was counted as one high current event, even though it had a complex time history, with many steps.

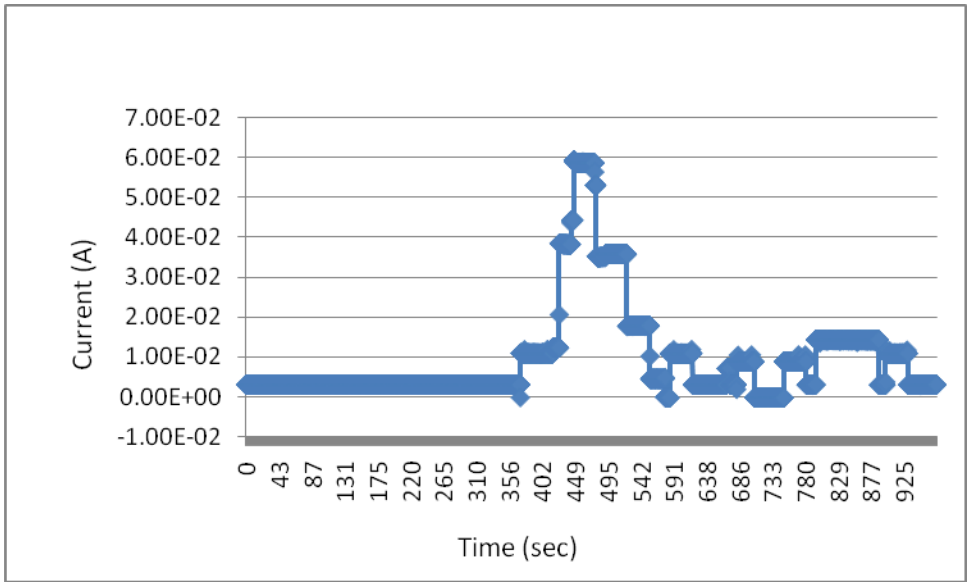

Fig. 3. Samsung DUT 3, Run 2—Dynamic Read Mode. Fluence was 106 Xe ions/cm2. There was a watchdog timer warning, meaning the DUT stopped responding to commands temporarily. But the DUT was fully functional for the next exposure.

There are many more current histories qualitatively similar to those shown in Figs. 2 and 3, but at this point, it is reasonable to consider what physical mechanism(s) can cause results such as those shown in Figs 2 and 3. To our knowledge, the best reference for this is by H. Shindou et al. [4]. They describe what they call "pseudo-SELs" and also "local SELs". A pseudo-SEL is defined as a rapid increase in power supply current, resembling a true SEL, but caused by a different mechanism. Typically, the actual mechanism is a logical conflict of the internal data bus lines. A localized SEL (LSEL) is a true SEL, but limited to a small part of the chip. The rest of the chip remains fully functional, because sufficient bias voltage is maintained. They say pseudo-SELs can often be cleared by a system reset, without cycling power. Presumably, a LSEL can also be cleared without a power cycle, if power is interrupted to just the affected part of the circuit. They also say that sometimes a power cycle is required, even without a true SEL. They use a photo-emission microscope to identify these cases. Since we did not have the use of such a microscope, we have counted as true SELs cases where a power cycle was required, and there were five such cases in the December test and two in the earlier MMS test. In Fig. 2 of [4], they show a current trace where three LSELs turn on sequentially. In each case, the current increases abruptly, but then is stable at the new value, until the next LSEL turns on. Qualitatively, their curve looks exactly like the rising current parts of Fig. 2 and 3. In our data, 48 of the 52 high current events appear to match the signature of LSEL events, indicating that they probably are due to LSELs. With this conclusion in mind, we now discuss the remaining current traces.

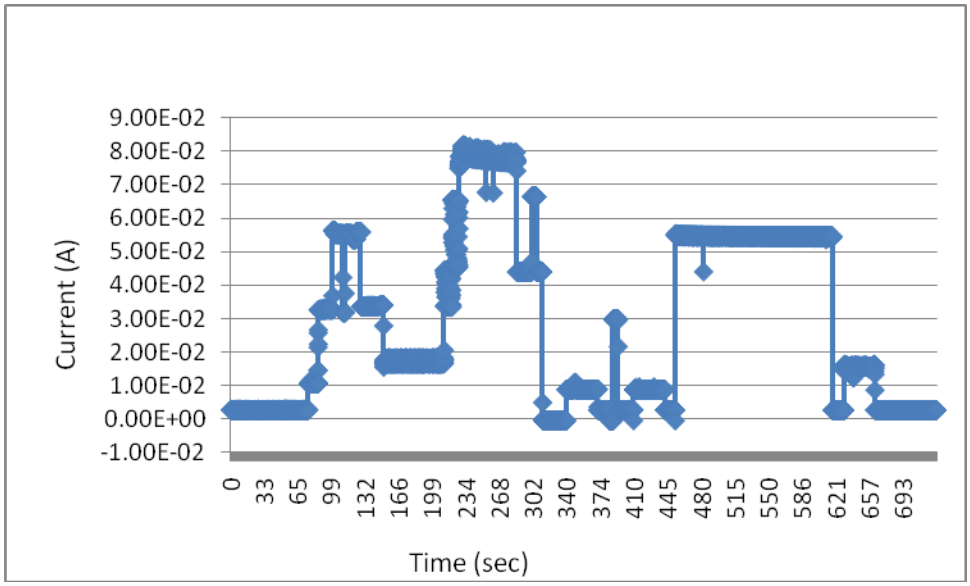

Fig. 4. Samsung DUT 3—Run 3. Dynamic R/E/W mode, with fluence 1.42×10^6 Xe ions/cm^2. Watchdog timer error, followed by self recovery, followed by functional failure.

In Fig. 4, we count three high current events, all of which have the step-wise signature discussed above. The first one, which peaks at about 55 mA, has duration slightly over one minute. The second one peaks at about 80 mA, with duration 1:13. The third one, which could perhaps be included in the second peak, goes from about 45 mA to more than 65 mA. It is the shortest in terms of duration, about 6 sec. After that, the beam was turned off and power cycled. The rest of the trace was due to efforts to reset the DUT for the next exposure, which were not successful. All three events have sharp increases in current, but the curves are flat on top, until the level changes again.

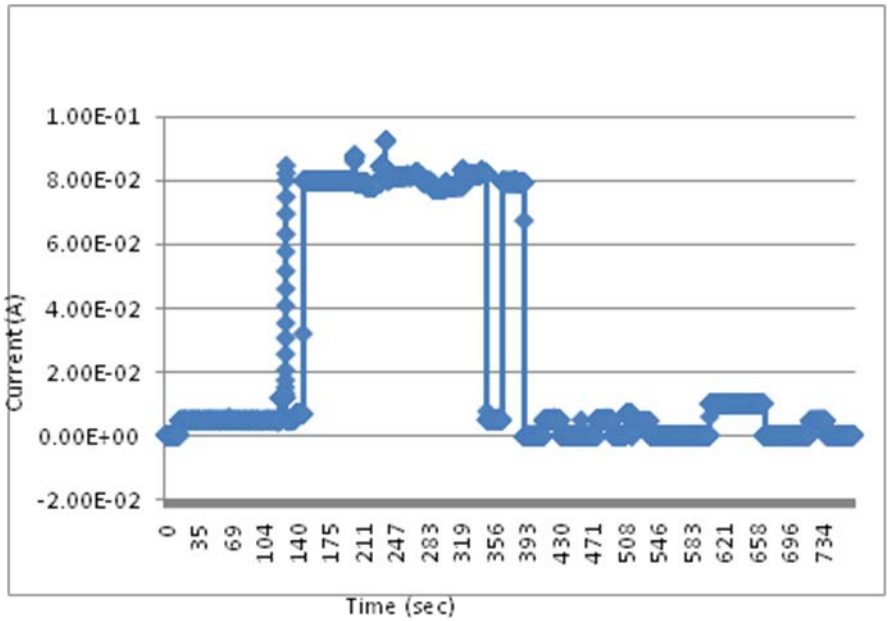

Fig. 5. Micron DUT 20—Run 10. Dynamic Read, with fluence 2×10^6Xe ions/cm^2. Watchdog timer error, but DUT was functional after reset.

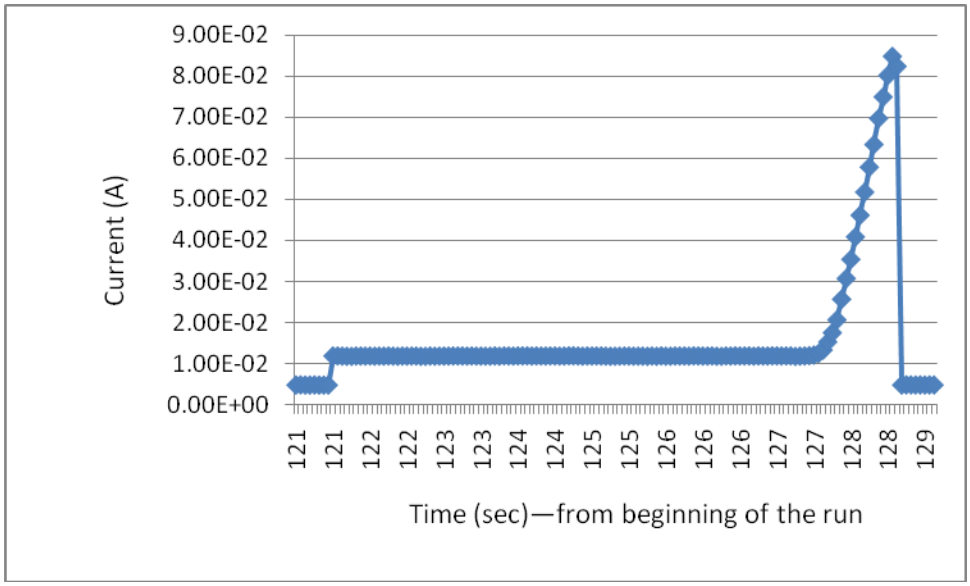

Fig. 6. First high current event from Fig 5, (Run 10) on an expanded time scale. Apparent example of pseudo-SEL.

In Fig. 5, there are three high current events, two of which appear to be additional examples of LSELs—each time, the current jumps to 80 mA, but is flat on top. The other high current event is more interesting when viewed on an expanded time scale, shown in Fig. 6. This appears to be an example of a pseudo-SEL, caused by bus contention, as discussed in [4]. The DUT is supposed to be in Dynamic Read mode, which means continuous reading with the beam on. The initial current in both Fig. 5 and 6 is 4-5 mA, which is the nominal Read current for this part. Then, in Fig. 6, the current jumps to 10 mA, which is the nominal write current. This suggests the DUT is trying to write, even though the commands are to Read. After about five seconds of apparently trying to write, the current ramps up abruptly, most likely because of contention between the Read and Write control logic, meaning the Read and Write circuits are fighting for control. The apparent contention is resolved abruptly, when the current drops back to the nominal Read current, as if a new Read command had the effect of turning off the Write circuit. Interval with current greater than nominal Write current is about one second

There are three other runs, where there is a high current event that does not have the step-wise LSEL signature, which we will discuss next. These are runs 5, 9, and 16. The complete current trace for Run 5 is shown in Fig. 7. There are two high current events that appear to be caused by LSELs, and one additional event. This last event is shown on an expanded time scale in Fig 8. Similarly, for Run 9 the entire trace is shown in Fig. 9, with one of the two high current events shown on an expanded scale in Fig. 10. The trace for Run 16, shown in Fig. 11, has three high current events, the first of which is shown on an expanded scale in Fig. 12.

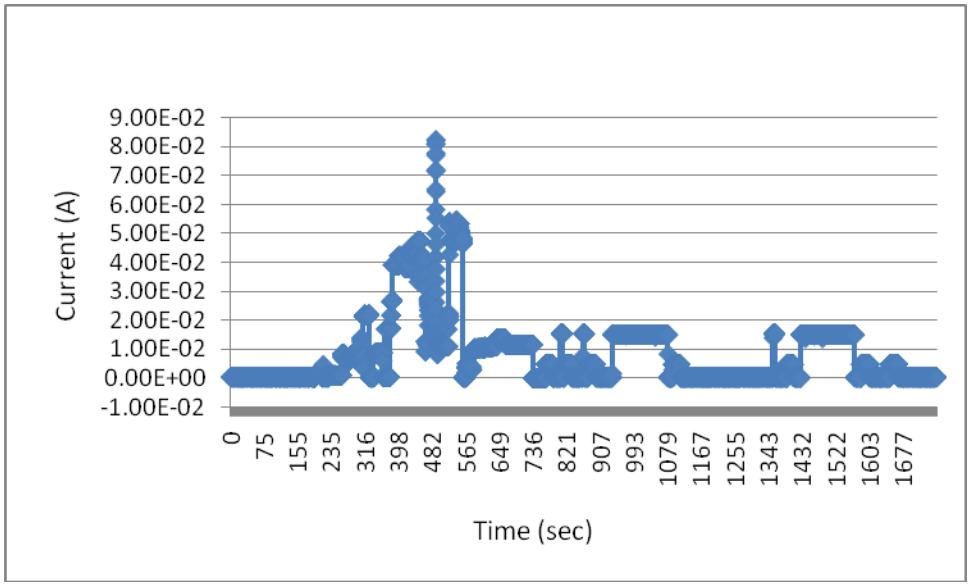

Fig. 7. Micron DUT 21—Run 5. Static mode, with fluence $5x10^6$Xe ions/cm^2. The DUT failed on this shot, losing both the Erase and Write functions.

These three events, from Runs 5, 9, and 16, could also be due to bus contention. However, for these last three events, the DUT was in Static mode, which means it was not receiving any commands from the test system. On Runs 5 and 9, the baseline current before the current rose was higher than the nominal static mode idling current (< 1 mA), indicating that there was activity not driven by commands from the test system. After the current dropped, the baseline current was lower than before, which suggests that some activity had stopped. There could have been bus contention associated with this change, which caused the momentary current rise. For Run 16, the baseline current before the current rose was the nominal static idling (< 1 mA) current initially, but it changed to a higher value after the current rose and fell, which suggests that some new activity started. If there was bus contention associated with this change, it could also explain the current rise. While this evidence is suggestive, it is not as conclusive as the case for Run 10, shown in Fig. 6, because the system does not send a command that clearly and quickly resolves the contention. In Static mode, the system sends no commands at all. For these reasons, bus contention is a plausible mechanism for the results of these three runs.

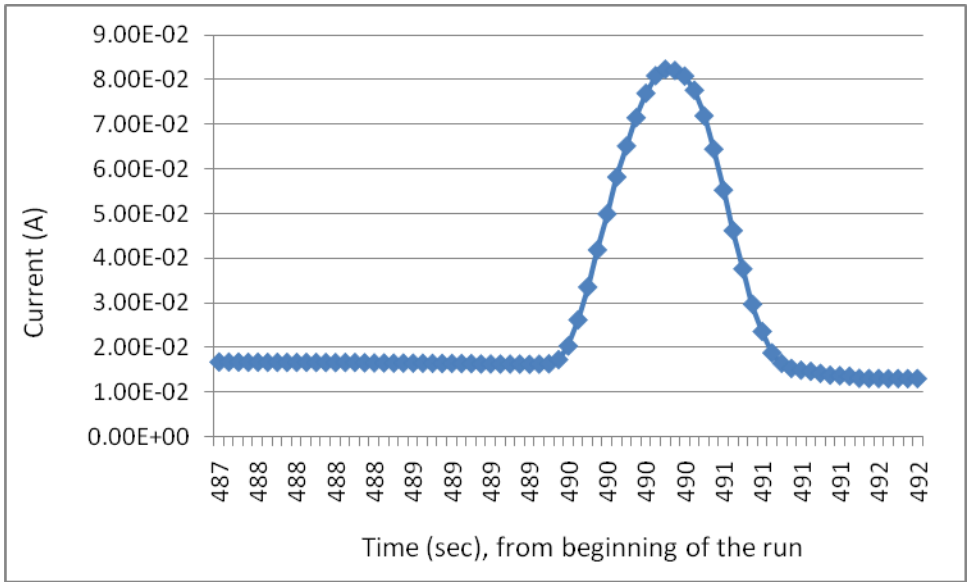

Fig. 8. High current transient from Run 5. Full width is about 1.7 sec, and FWHM is close to 1 sec.

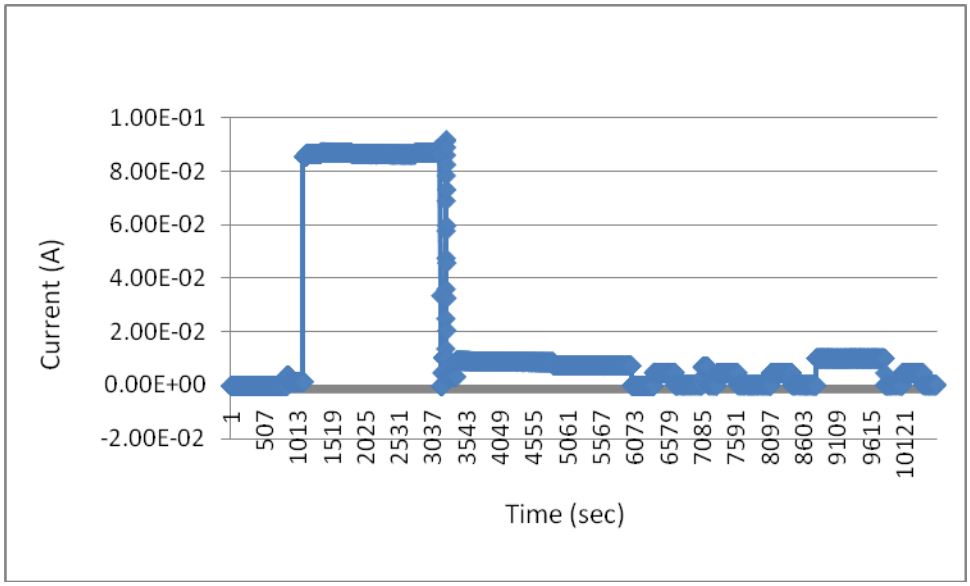

Fig. 9. Micron DUT 20—Run 9. Static mode, with fluence $1x10^6$ Xe ions/cm^2. There are two high current events on this Run, although they appear to run together on this scale. The DUT survived this Run and was used on several more Runs.

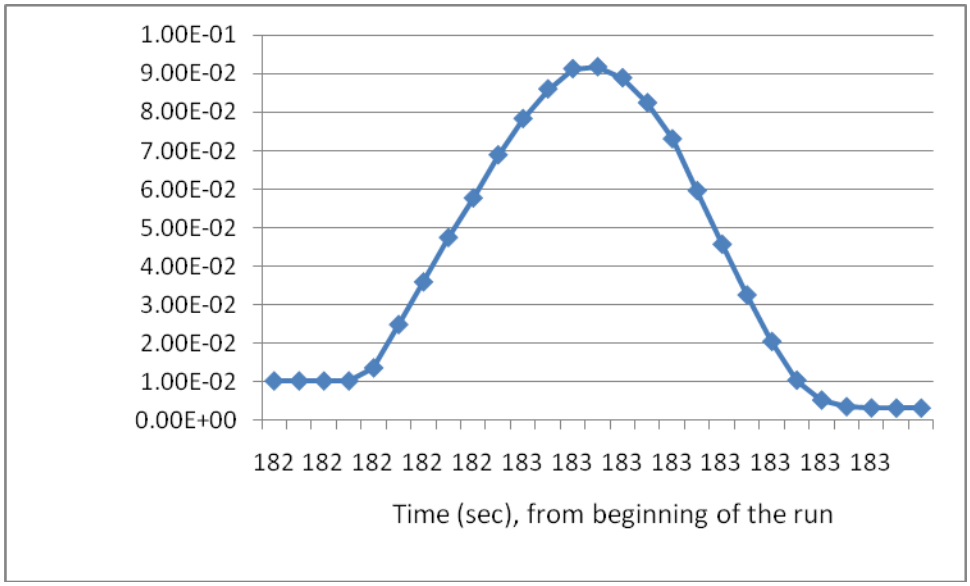

Fig. 10. High current transient from Run 9 expanded. Full width is about 1.2 sec, FWHM is about 0.7 sec.

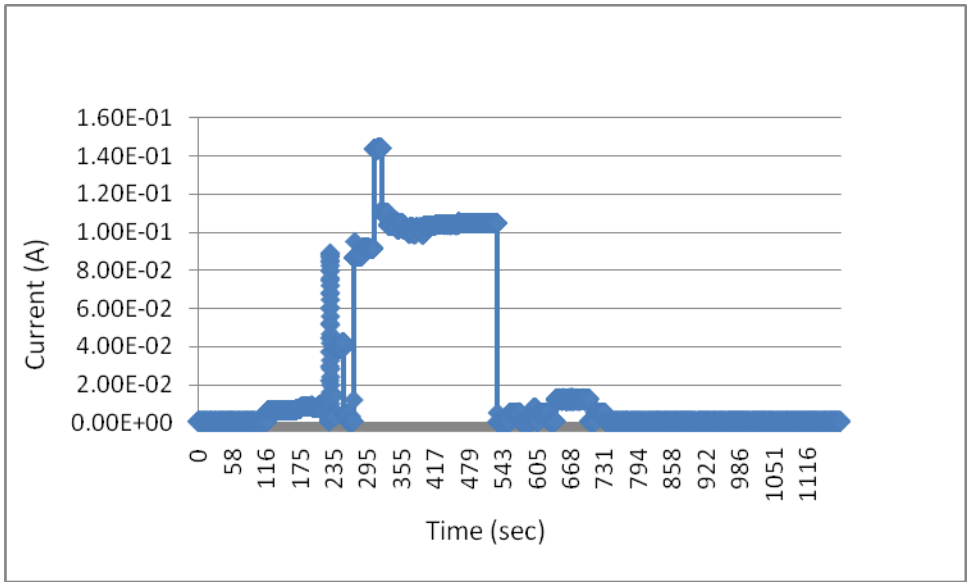

Fig. 11. Micron DUT 20—Run 16. Static mode, with fluence $3x10^6$Au ions/cm^2. DUT had a watchdog timer error, then an apparent true SEL, with power cycle required to bring down the current. DUT was fully functional after PC, however. There are two high current events (LSELs) besides the transient in the next Figure.

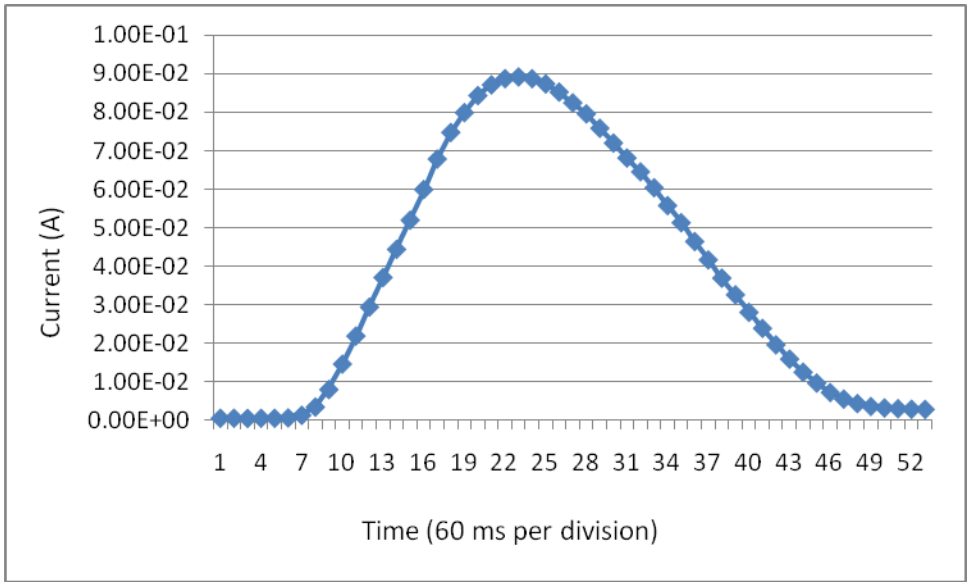

Fig. 12. Run 16 transient. Full width is about 2.5 sec, FWHM is about 1.4 sec.

We performed three exposures where the charge pump and other peripheral circuits were shielded from the beam. These were Runs 19, 20, and 21, where one run was in each of the three test modes. With the control logic shielded, there were no high current events on any of the three runs. Therefore, we will not show any of these current traces. On Run 22, we did the opposite, shielding the memory array and exposing only the control logic. Run 22 was actually broken into four parts (22 a, b, c, d) because it had been suggested that if we stopped the run after the current went high, and interrogated the chip, we could determine exactly when the chip failed. These traces are not very interesting because the runs were terminated by the operator as soon as the current went high. It had also been suggested that, in Dynamic Read mode, the Samsung parts would fail, probably at the first high current event. The actual result was that the DUT failed, but only after the fourth high current event. We also tested Micron parts in this manner—not with the array shielded, but stopping the run every time the current went high. According to [1, 2], the Micron parts should not have failed. The first Micron part , DUT 23, actually failed at the first high current event, on Run 23. Again, we did not let the run continue after the current went high, so the traces are not much to look at. Then we replaced the DUT with another Micron part, DUT 24. Instead of breaking one run into pieces, we assigned each piece a new run number, so Runs 24-28 are similar to the pieces of Run 22, except for the manufacturer—stop the run as soon as the current goes high, and check the part. This second DUT survived five runs, so we decided to stop the Dynamic Read test and use the other test modes.

Of the remaining runs, all the traces show some current above normal except one, which was Run 12. The trace is shown in Fig. 13. In this Run, the DUT is in R/E/W mode, but suffers a watchdog time interrupt, meaning that it does not respond to all the commands. For this reason, the trace does not look normal. But the maximum current in the entire run is only 14 mA, which is within the normal Write current range for this part. Since the test system was sending commands to Write, this current level is only to be expected. In light of the other results, the surprising thing about this run is that there were not higher current levels observed.

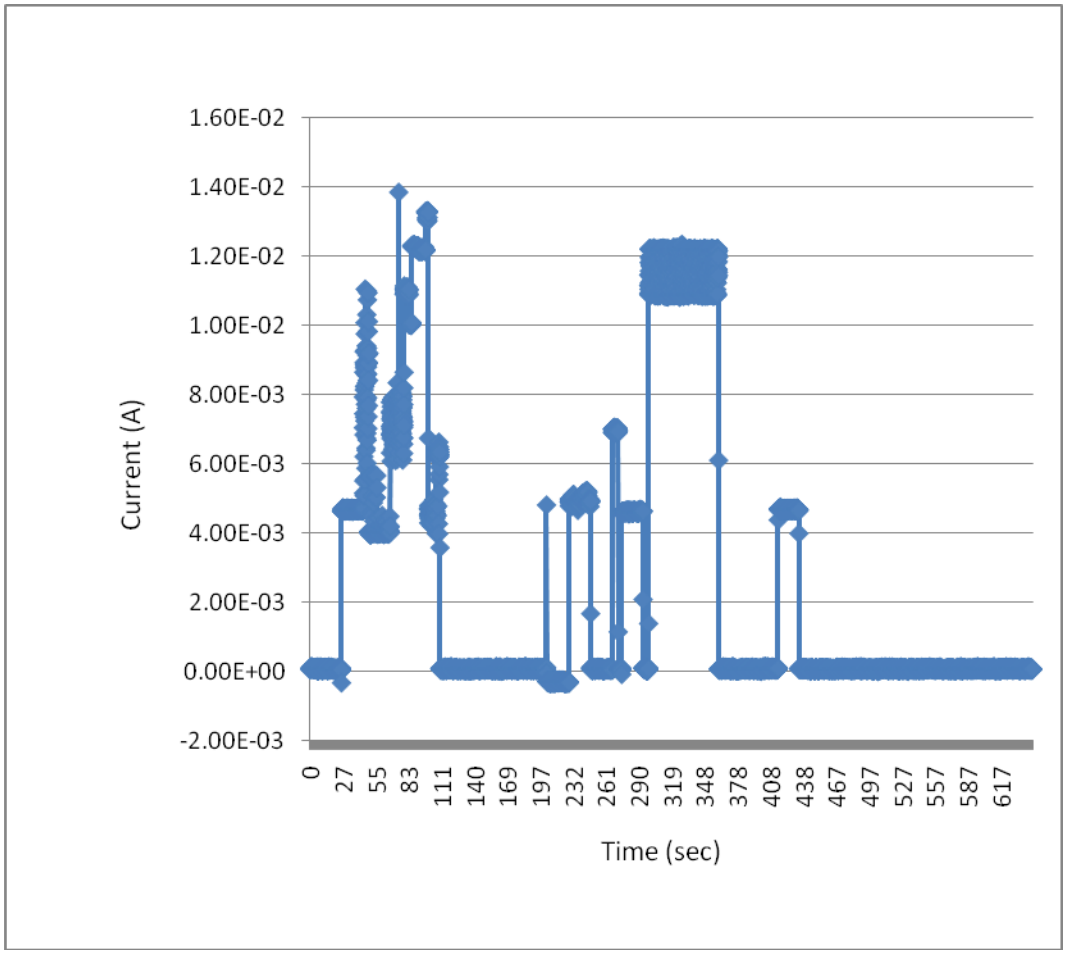

Fig. 13. Micron DUT 20—Run 12. R/E/W mode, with fluence 1.05×10^6 Xe ions/cm^2. Watchdog timer error, but DUT was fully functional after PC.

The remaining runs all produced current traces qualitatively similar to those already discussed, so we will not discuss them extensively. But the traces will be shown for the sake of completeness. Run 4 (Fig. 14) and Run 6 (Fig. 15) both have maximum currents of about 25 mA, which falls below our 40 mA standard. But both are counted as high current events, because both were Dynamic Read shots, and the current was at least at least 2x the nominal Read current level.

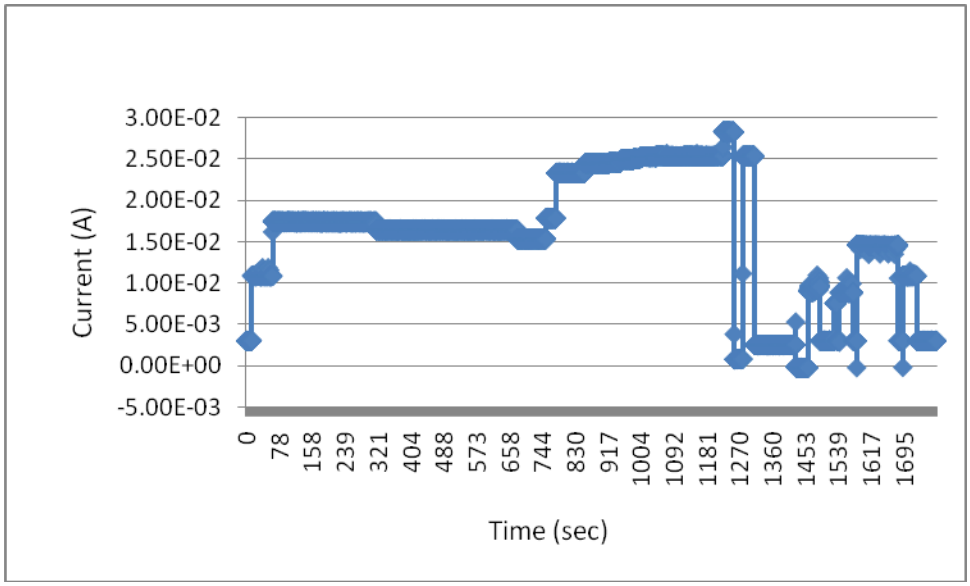

Fig. 14. Samsung Dut 4—Run 4. Dynamic Read mode, with fluence $1x10^6$ Xe ions/cm^2. Watchdog timer error, but DUT was fully functional after being reset.

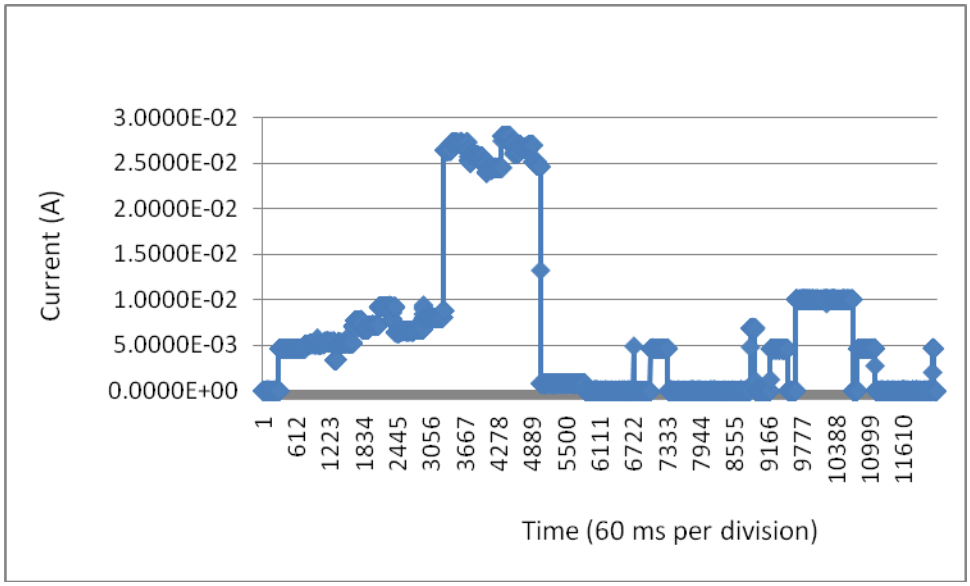

Fig. 15. Micron DUT 20—Run 6. Dynamic Read mode, with fluence $3x10^6$ Xe ions/cm^2. DUT was fully functional after reset.

For Run 7, the results are shown in Fig. 16. On this run, there are three events where the current exceeds 40 mA. The first two of these are shown on an expanded time scale in Fig. 17 and Fig. 18. Both clearly have the stair step structure typical of LSEL events. Run 8 results are shown in Fig. 19, showing one high current event with a rectangular shape of duration 3 sec. Run 11 results are shown in Fig 20, with two high current events, both apparently due to LSELs. Fig. 21 shows results for Run 13, with just one rectangular shaped high current event, of 48 sec duration.

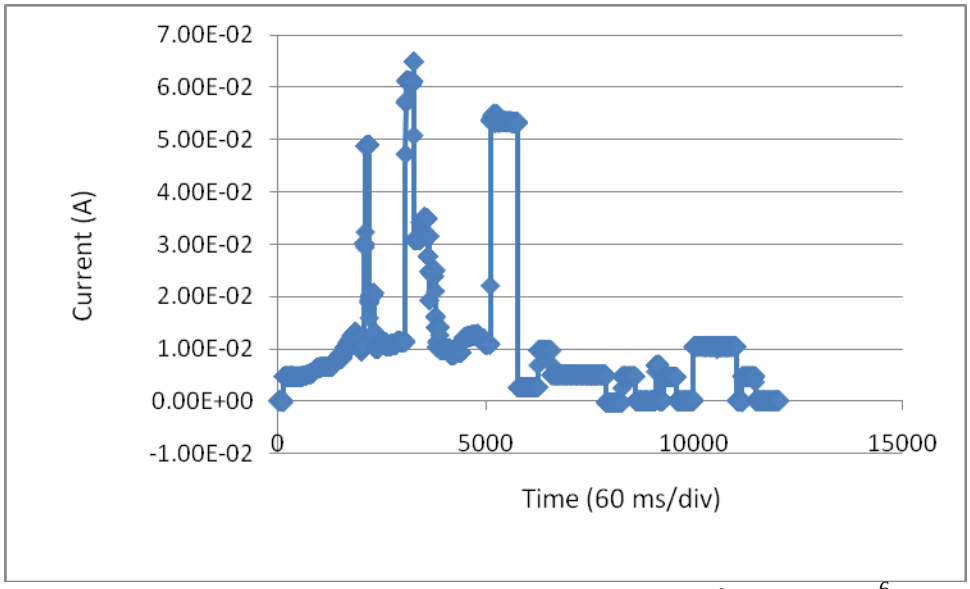

Fig. 16. Micron DUT 20—Run 7. Dynamic Read mode, with fluence $3x10^6$ Xe ions/cm^2. Watchdog timer error, but DUT was fully functional after PC.

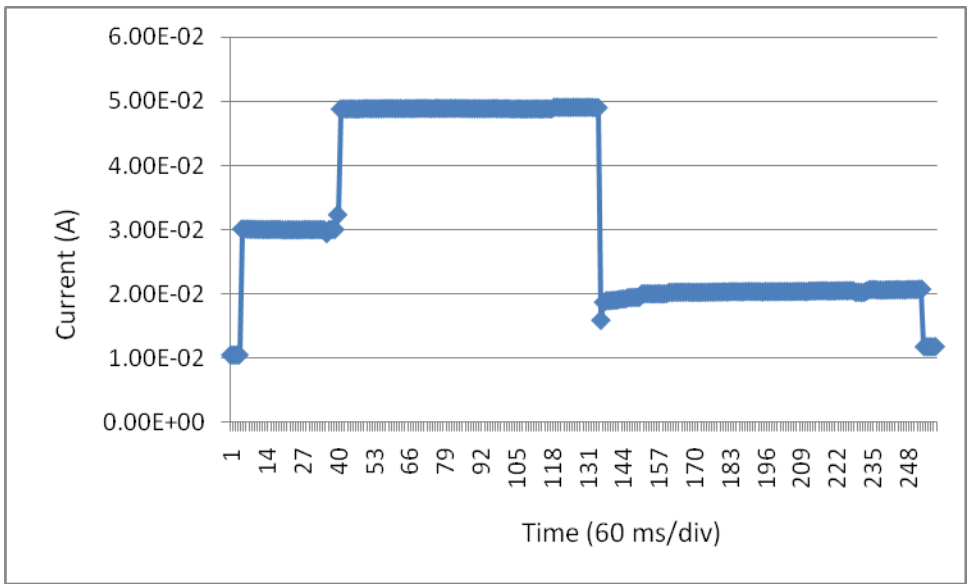

Fig. 17. Run 7, where first current peak is shown on an expanded time scale. Period shown is about 14 sec.

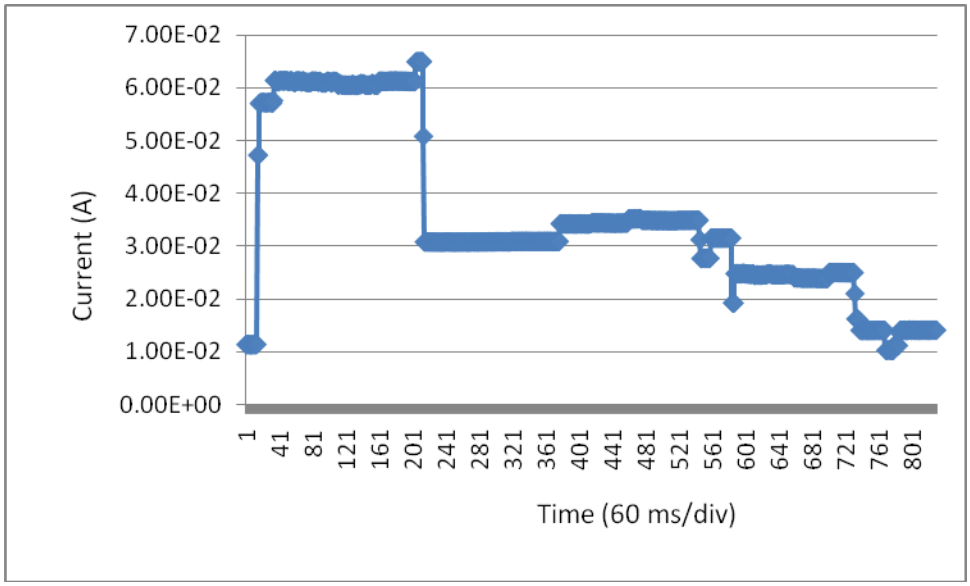

Fig. 18. Run 7 second peak on an expanded time scale. Total duration is about 50 sec.

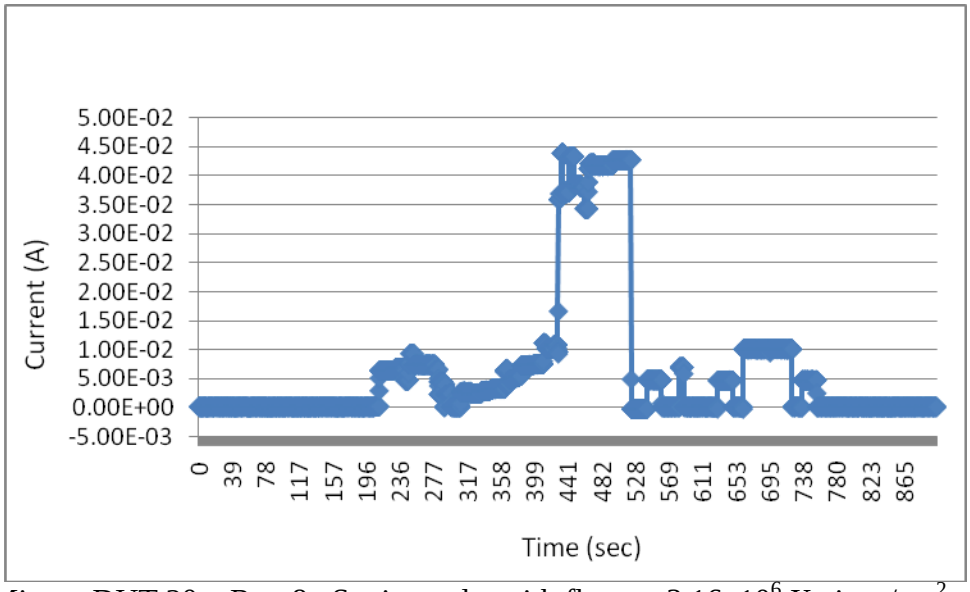

Fig. 19. Micron DUT 20—Run 8. Static mode, with fluence 3.16×10^6 Xe ions/cm^2. DUT was fully functional after PC.

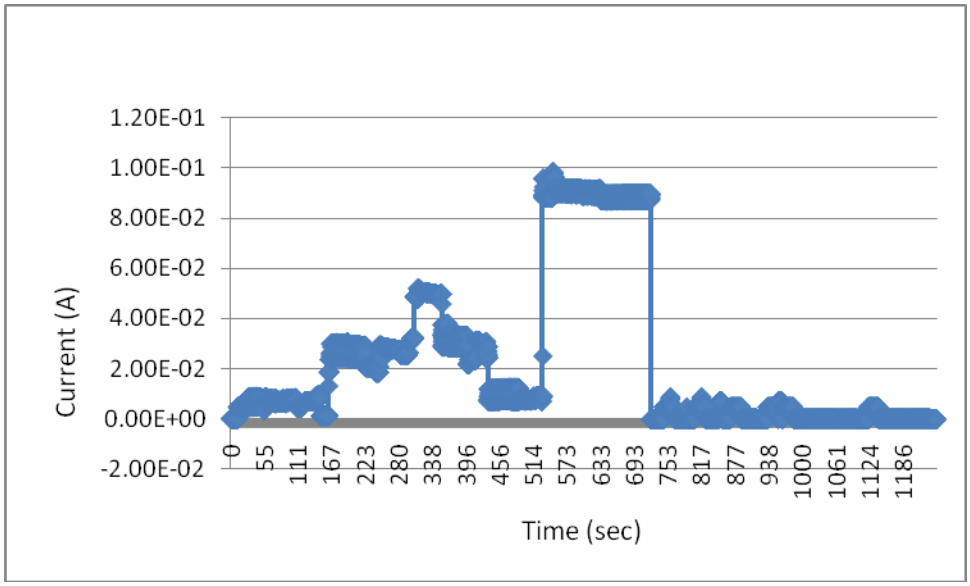

Fig. 20. Micron DUT 20—Run 11. R/E/W mode, with fluence $4x10^6$ Xe ions/cm^2. Watchdog timer error, but DUT was fully functional after PC, except for one block, which was screened out.

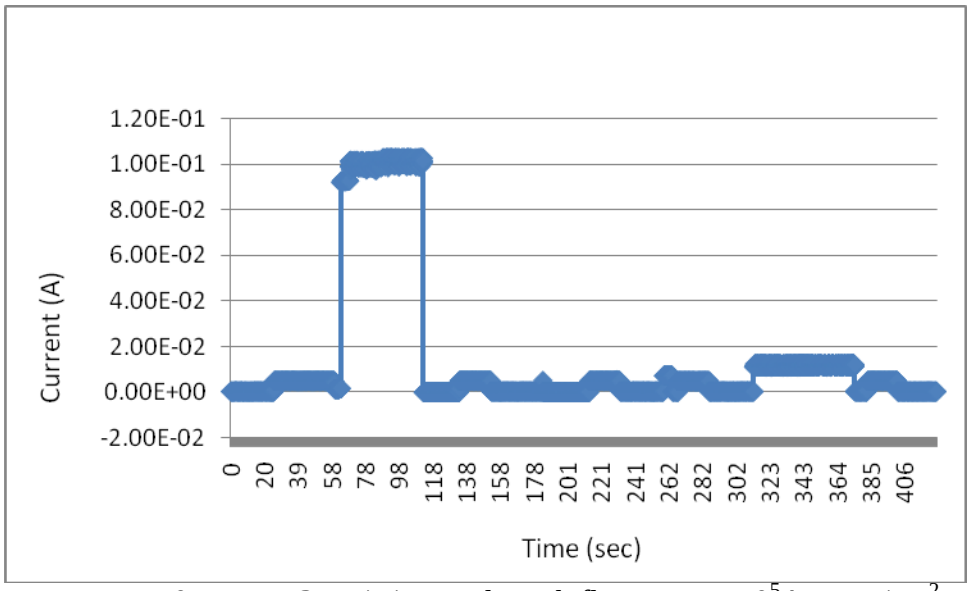

Fig 21. Micron DUT 20—Run 13. R/E/W mode with fluence $1.7x10^5$Au ions/cm^2. Watchdog timer error, but DUT was functional after PC.

Results for Run 14 are shown in Fig. 22, where the waveform has the stair step shape associated with multiple localized SELs (LSEL). The total duration is of the high current condition is about three minutes. Run 15 is shown in Fig. 23, where an apparent true SEL required a power cycle, and operator intervention, to restore normal current levels. The current transient peaking at 150 mA, which occurred on top of an already high baseline, is shown on an expanded scale in Fig 24. Run 17 is shown in Fig. 25, but high currents ended with a Read command. No PC was required, but the DUT failed, losing both erase and write functions.

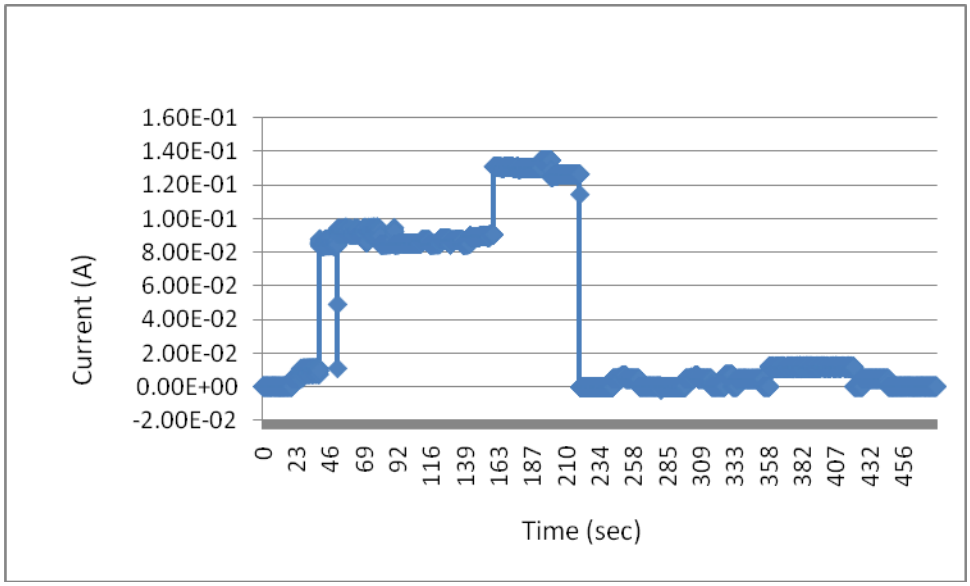

Fig. 21. Micron DUT 20—Run 14. Dynamic R/E/W mode, with total fluence $2x10^6$ Au ions/cm^2. Watchdog timer error occurred, but DUT was functional after being reset.

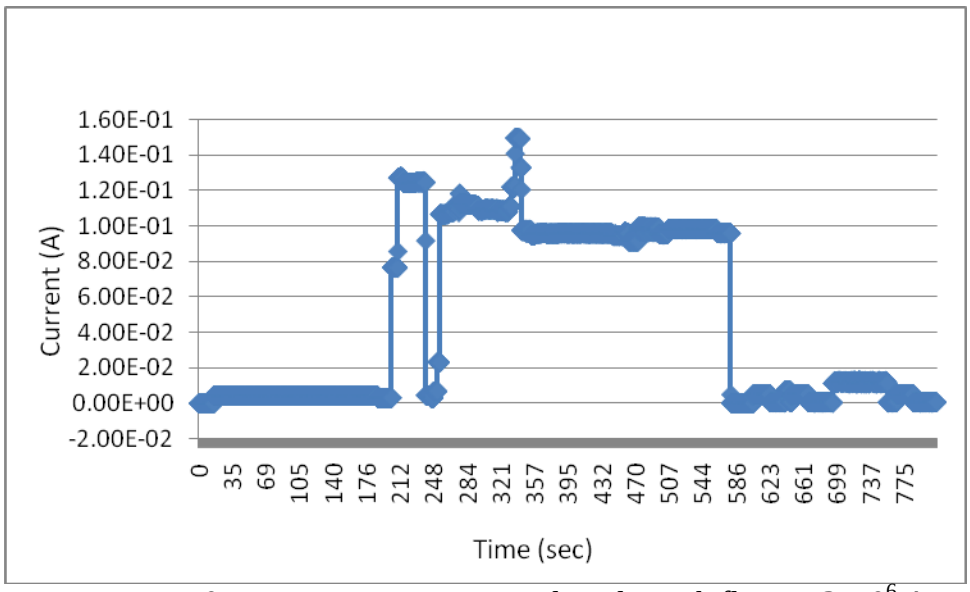

Fig. 22. Micron DUT 20—Run 15. Dynamic Read mode, with fluence $3x10^6$ Au ions/cm^2. Watchdog timer error, and DUT reset failed to restore normal current levels. Apparent true SEL with PC required to end high currents. DUT was still functional afterwards.

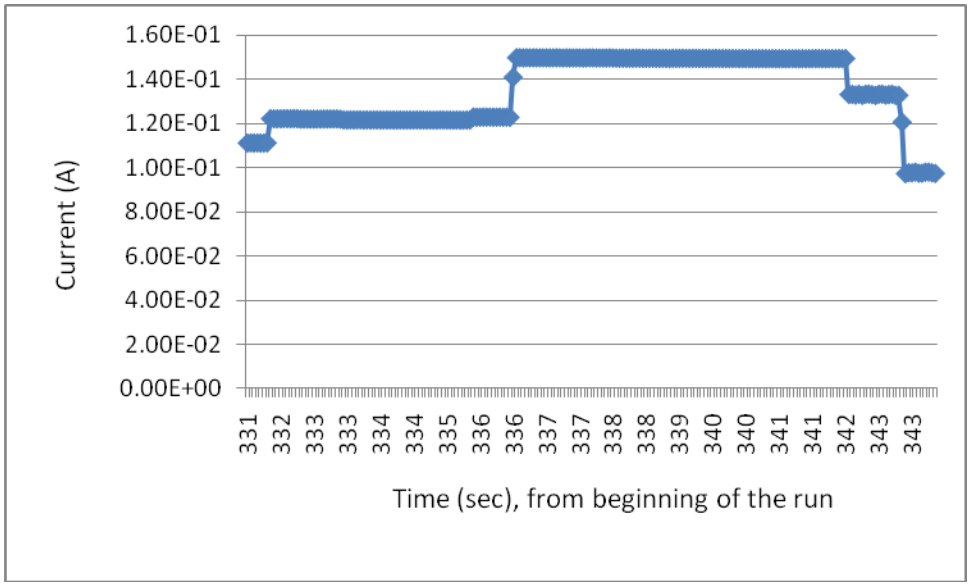

Fig. 23. Run 15 highest peak, 150 mA, with typical stair step signature. Total duration is about˚13 sec.

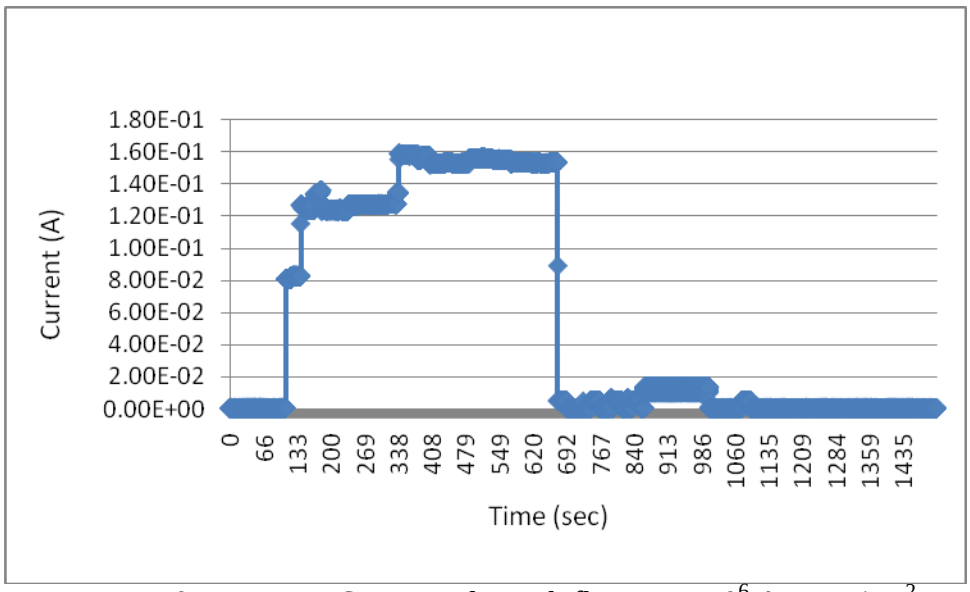

Fig. 24. Micron DUT 20—Run 17. Static mode, with fluence $2x10^6$ Au ions/cm^2. High current dropped on Read command. DUT failed on this run, losing both erase function and write˚function.

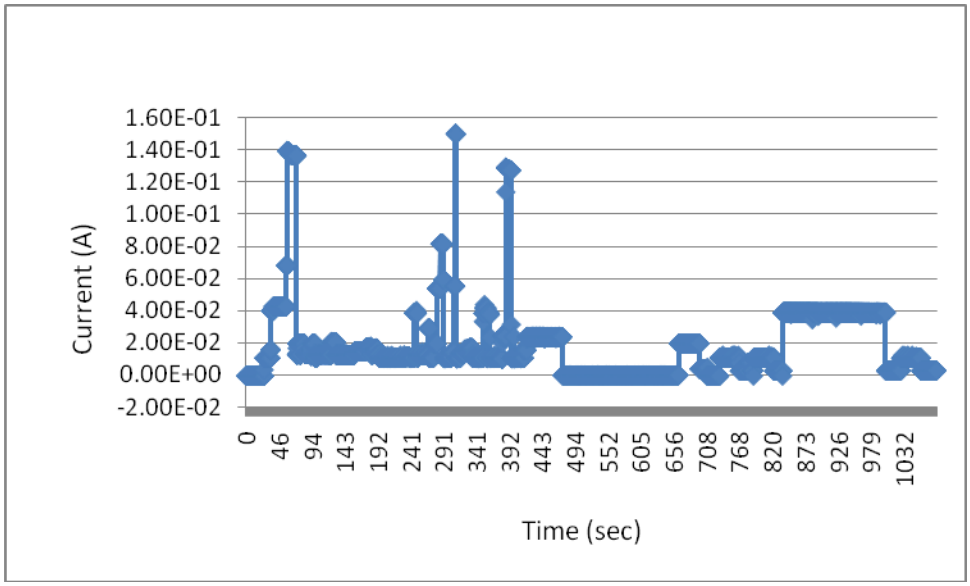

Fig. 25. Samsung DUT 4—Run 18. Dynamic Read mode, with fluence $3x10^6$ Au ions/cm^2. Watchdog timer error, and DUT failed, losing both erase and write functions.

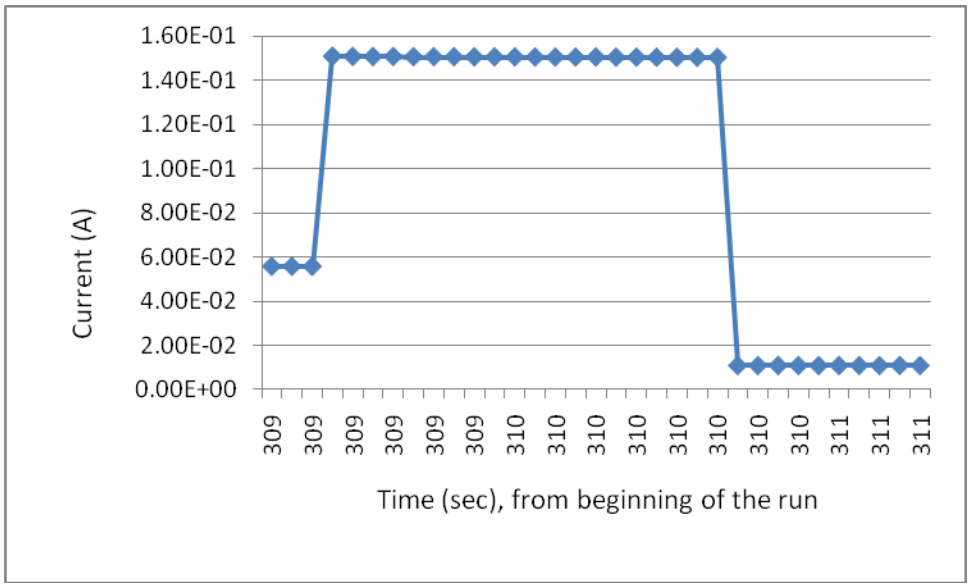

Fig. 26. Highest peak current event in Run 18, shown on an expanded time scale. Total duration of the entire trace is about 2 sec.

In Fig. 25 and Fig. 26, results from Run 18 are shown. In Fig. 25, there are five peaks that exceed 40 mA, although there is one more at about 39 mA that could also have been counted, perhaps. The sharpest of these is shown on an expanded time scale in Fig. 26. The interval where the trace is flat on top is about 1.4 sec.

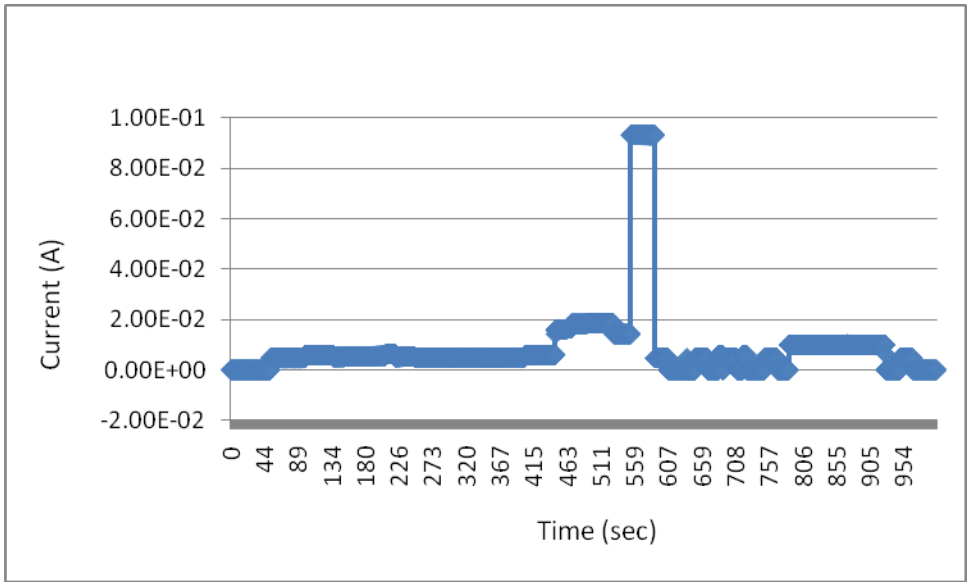

Fig. 27. Micron DUT 23—Run 23. Dynamic Read mode, with fluence 4.4×10^5 Au ions/cm^2. Write and Erase functions both failed.

In Fig. 27, we show the results for Run 23, for Micron DUT 23 in Dynamic Read mode. The part failed, losing both the Write and Erase functions.

Runs 29-35 were conducted at very low fluence, 300 ions/cm^2-sec, because one idea that had been considered was the current spikes in [1, 2] were due to multiple ion strikes. Results are shown in Figs. 28-34. In all cases, a high current event occurred, at relatively low fluence. When the high current was observed, the operator intervened to restore normal operation.

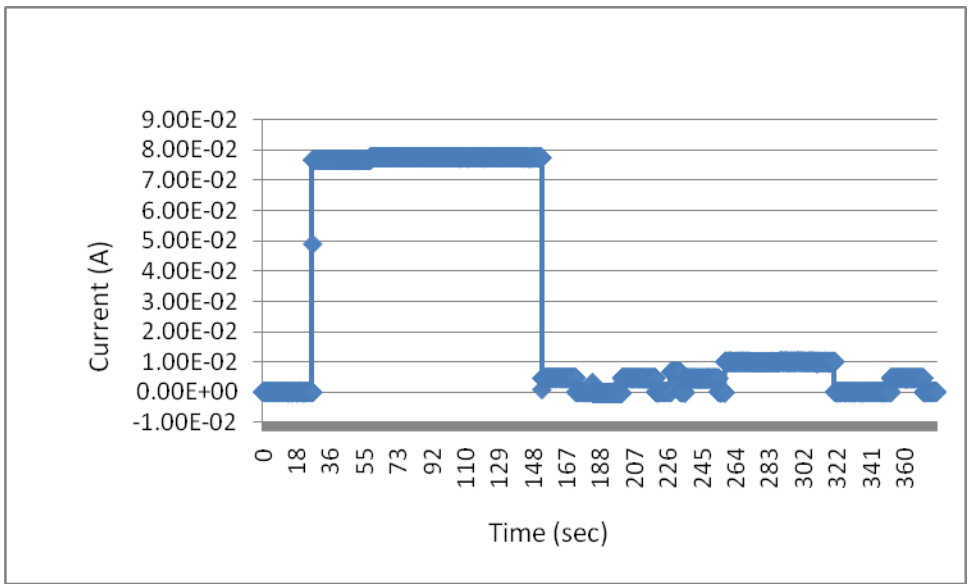

Fig. 28. Micron DUT 24—Run 29. Static mode, with fluence 5.8×10^4 Au ions/cm^2. DUT was fully functional afterwards. High current duration was a little over 2 minutes.

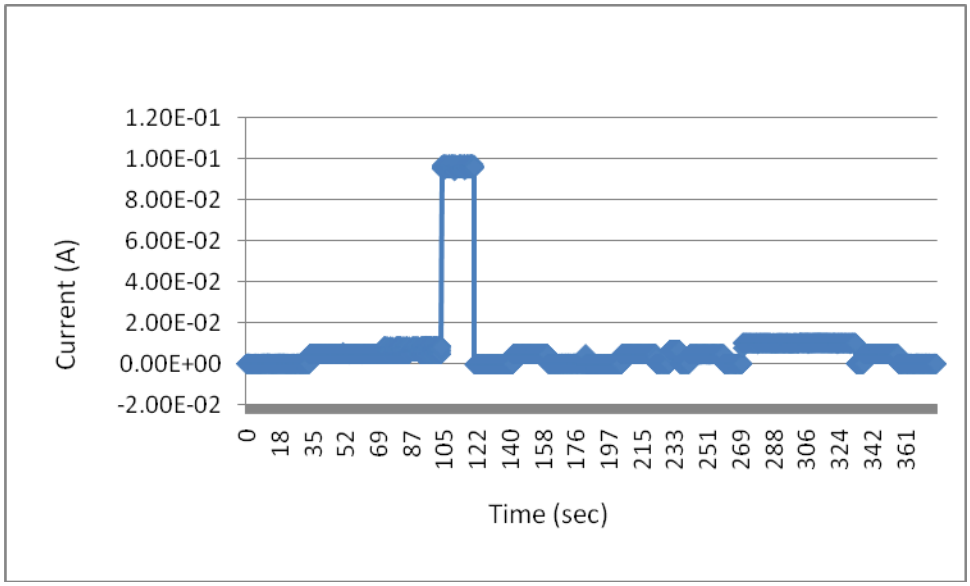

Fig. 29. Micron DUT 24—Run 30. Dynamic R/E/W mode, with fluence 1.64×10^4 Au ions/cm^2. DUT was fully functional after PC. Duration of high current was 17 sec, terminated by operator.

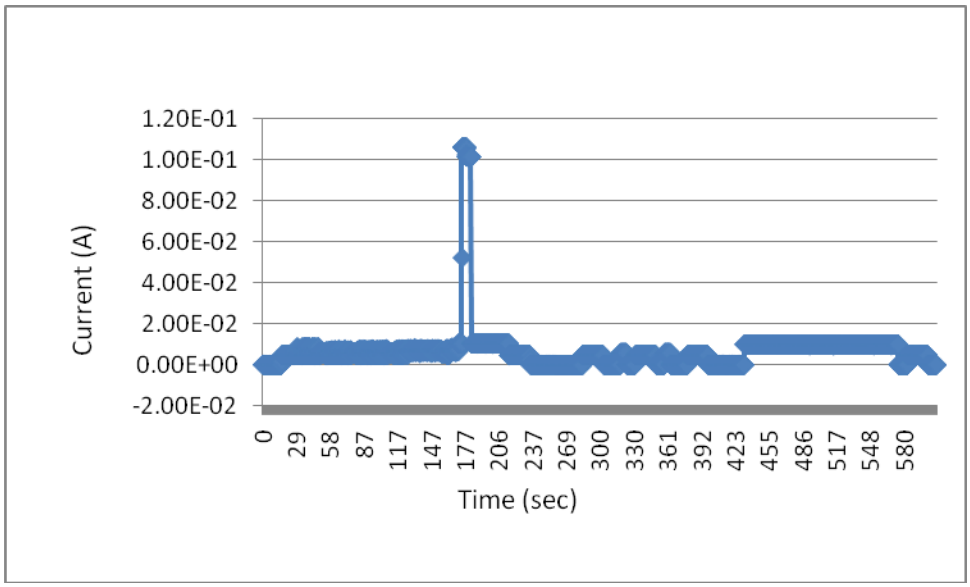

Fig. 30. Micron DUT 24—Run 31. Dynamic R/E/W with fluence 5.6×10^4 Au ions/cm^2. Current dropped when the beam was turned off, but DUT failed, losing both Erase and Write functions.

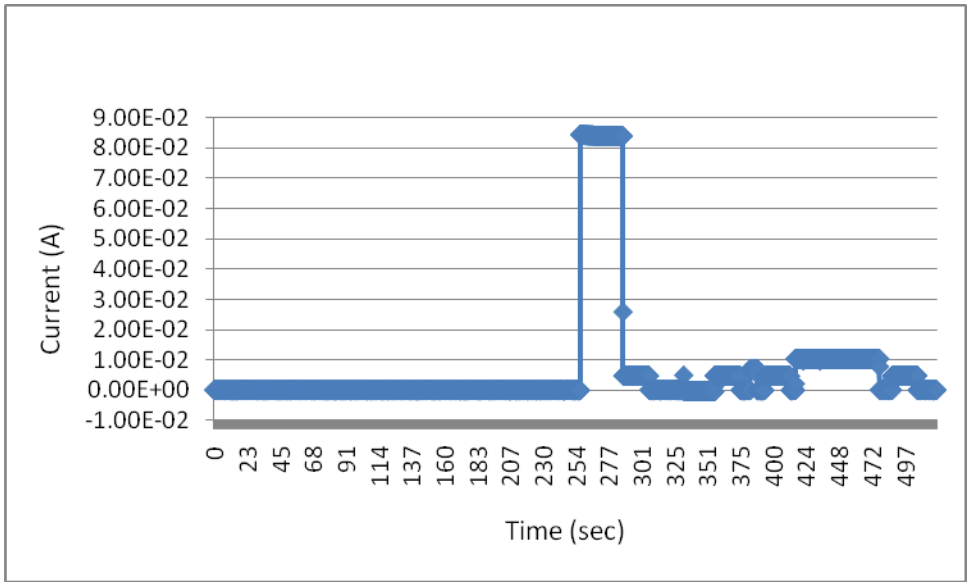

Fig. 31. Micron DUT25—Run 32. Static mode, where current went high after $7x10^3$ Au ions/cm^2. Current dropped to normal after 31 sec, from operator intervention.

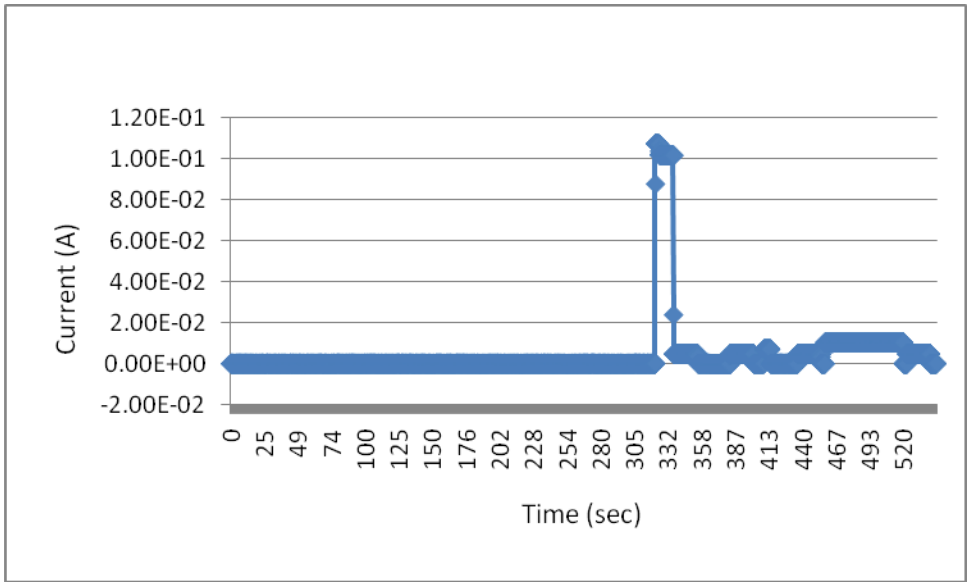

Fig. 32. Micron DUT—Run 33. High current after $9.7x10^4$ Au ions/cm^2, terminated by operator after 14 sec.DUT was fully functional afterwards.

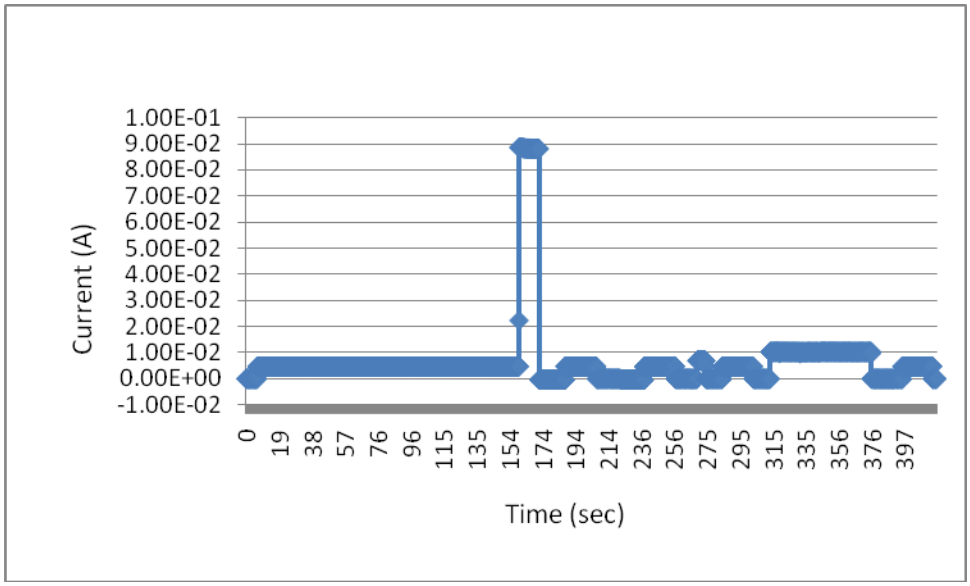

Fig. 33. Micron DUT 25—Run 34. Dynamic read mode, where current went high after 5.4×10^4 Au ions/cm^2. High current was terminated by the operator after 12 sec. DUT was fully functional afterward.

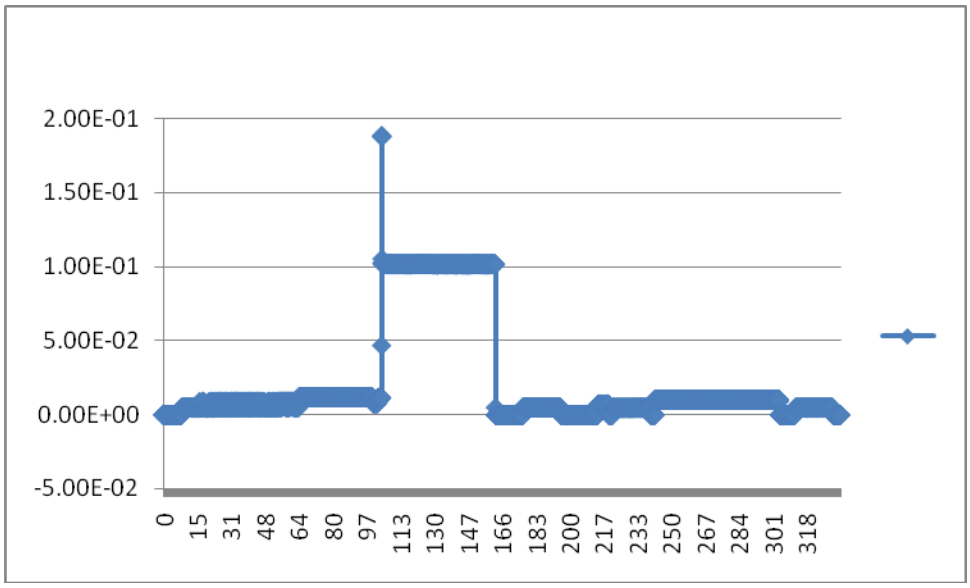

Fig. 34. Micron DUT 25—Run 35. Current went above 100 mA after 3.5×10^4 Au ions/cm^2 and beam was turned off, but current limiting brought it down to 100 mA. Operator intervened to cycle power after about 1 minute. DUT was fully functional afterwards.

VII. Discussion

The results reported in [1, 2] are qualitatively different from those reported in [3], and the purpose of this test was to try to resolve the differences. To put these differences in perspective, it is important to understand the different beam conditions that were used, and their consequences. In [1, 2], the authors commonly perform exposures of 1×10^7 Au ions/cm^2, with LET of about 87 MeV/mg/cm^2. According to the input spectrum for CREME96, for geosynchronous orbit in the Adams worst case environment, the flux at or above this LET is

about 1 particle/cm^2 every 7200 years, which is about 100 human lifetimes. Under present conditions in space, it would take 7.2x10^{10} years to accumulate a fluence of 10^7 ions/cm^2. Since this is more than 5x the interval since the Big Bang, it obviously would not be a good assumption to assume constant conditions for such an interval! If such an exposure produced 10 current spikes, which is a roughly typical result, they would then be about 7x10^9 years apart, which is greater than the age of the Earth and about half the interval since the Big Bang. Therefore, even if we had reproduced the reported current spikes in a ground test at the accelerator, it would not make them real in space.

In all accelerated stress testing, the goal is to accelerate the degradation mechanisms and failure mechanisms that will be significant in normal operation, or real life, if you will. The problem is that increasing the stress always has the potential to introduce new failure mechanisms, which would not ever be encountered in normal operation. In designing an accelerated test procedure, one always has to try to prove that the test method accelerates the right mechanisms, without introducing other, new mechanisms. Frequently, this is a significant challenge.

In the MMS testing of the Micron parts, there were only two high current events in 64 shots at lower LETs. There were two true SELs (apparently), requiring power cycling to restore normal current levels. Both occurred with Xe ions, and both DUTs failed. But the pseudo SELs and localized SELs, commonly observed in the December test run, were not observed with Ne, Ar, or Kr ions. Nor were they observed with Xe, although other high current events were.

As indicated above, the purpose of this experiment was to duplicate the experimental conditions in [1,2] to see if we could also duplicate the current spikes reported there. The current spikes were described as having a typical pulse width of 300-400 ms. Although we observed 52 high current events in this experiment, they lasted for widely varying periods. None was as short as 300-400 ms, but some of them can be characterized in different ways, and the exact difference depends on how they are characterized. Forty-eight of the 52 events are basically rectangular, stair step waveforms, where we have used the interval from baseline to baseline as the measure of the high current event period, because it is very clear where the vertical line rises from the baseline current. Many of those events have multiple stair steps, so another measure, such as FWHM, would miss some of the steps. Then, to be consistent across the whole data set, we have also used the baseline-to-baseline interval for the four events that do not have rectangular shapes. These are shown in Figs. 6, 8, 10, and 12. They are tabulated in this way in Table II, below. But it has been suggested that FWHM would be a better measure of the duration of these transients, so it is also given in the text above. Using the baseline-to-baseline interval, none has duration less than 1 sec, although two have FWHM less than 1 sec. Depending on how one looks at it, the difference is 2-3x, or less. While one can perhaps debate the significance of the difference, it is clear that, by any measure, there is <u>some</u> difference. It is also clear that the rate at which short duration, high current transients happen is at least one order of magnitude lower than in [1, 2], and perhaps two orders of magnitude less. For example, Fig. 6 of [1], shows ten high current spikes of short duration, in one exposure, lasting about 1000 sec. Here, if LSEL rectangular waveforms are excluded there were only four in the entire 35 exposures (or 38, if 22 a, b, c, and d are counted as separate events). Putting it another way, when we tried to duplicate Their Fig. 6, the closest we came was that nine of the 10 events disappeared, and one had a different pulse width. On most shots, all ten disappeared, which suggests that the mechanism is just different.

In [1, 2] it was reported not only that current spikes occur, but that (1) they are destructive at high LET, but not at low LET; (2) they are destructive in Dynamic Read mode for Samsung parts, but not for Micron parts; (3) Micron parts also fail in Write mode. In this test, only high

LETs were used. The results here were that four Micron parts failed: (1) one in static mode, on the first shot, and another in static mode but not on the first shot; (2) one failed in Dynamic Read mode, on the first shot; (3) one failed in R/E/W mode, but not on the first shot. Four Samsung parts also failed: (1) one in static mode, on the first shot; (2) two in Dynamic Read mode, neither on the first shot; and (3) one in R/E/W mode, not on the first shot. It had also been suggested that, with Samsung parts, if we stopped the Run as soon as the current went high, we could determine when the failure occurred, and that it would usually occur on the first high current occurrence. When we tried this procedure, the part did eventually fail, but not until the fourth high current event. We also note that all the failed parts were tested again, about three weeks after the TAMU run, and none of them had recovered.

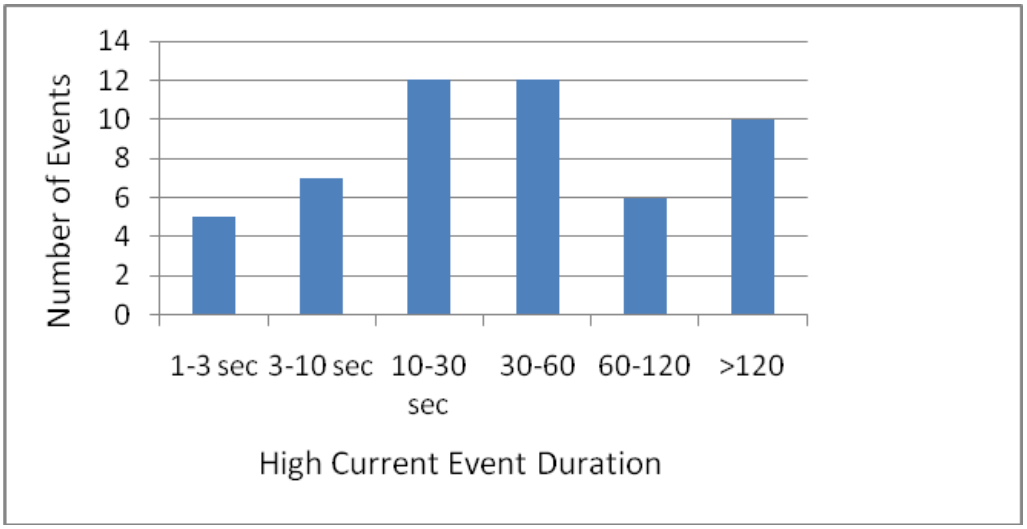

Table II. Number of high current events with given duration, in seconds.

As we have shown, the current events arise from two different known mechanisms, described in [4], although there could possibly also be others. The two mechanisms are localized SEL (LSEL), which gives rise to the stair step pattern evident in 48 of our 52 high current events. The other mechanism is referred to in [4] as a pseudo-SEL, caused by data bus contention. This seems to explain very nicely the current trace shown in Fig. 6. It probably also explains the results in Figs 8, 10, and 12, although this is less certain than for Fig. 6. Therefore, we conclude that bus contention is the leading candidate for the mechanism to explain other four events in this study. In fact, it is the only mechanism that has been proposed, so far. The authors of references [1, 2] report observing current spikes, and they report other observations which suggest to them that the spikes are connected to the charge pump. But they have not yet offered an explanation of how the charge pump might cause current spikes. If an alternative mechanism is proposed, we would certainly be willing to consider it, but until that happens, bus contention, as proposed in [4] is the most likely explanation for our data. We also note that the authors of reference [4] have obtained data very similar to what we observed in this test run. And they offer very convincing explanations for almost all our data. But their test circuits did not even contain charge pumps! It is also possible that bus contention will not be able to explain the results in [1, 2]. Similar results to what we report here have also been reported in other kinds of circuits [4, 5], which leads us to believe these results are not unique to flash memories, but common to advanced technologies of many kinds.

VIII. Conclusions

The purpose of this test was to try to reproduce results reported in [1, 2]. Although we observed 52 high current events, none of them matched exactly signature of the events reported in [1,2], although one or two of them might be considered “close”, depending on how “close” is defined. We have identified mechanisms which seem to account for the results reported here, and which are commonly observed in advanced technologies other than flash memories. Therefore, there is no reason to think the results reported here are unique to flash memories. One of the main conclusions in [1, 2] was that Micron parts would not fail except in Program (Write) mode. We were unable to confirm this result, observing three such failures.

IX. References

1. F. Irom and D. Nguyen, Single event effect characterization of high density commercial NAND and NOR nonvolatile flash memories, IEEE Trans. Nucl. Sci., vol. 54, no. 6, pp.2547-2553 (2007).
2. F. Irom, D.N. Nguyen, M. Bagatin, G. Cellere, S. Gerardin, and A. Paccagnella, Catastrophic failure in highly scaled commercial NAND flash memories, IEEE Trans. Nucl. Sci., vol. 57, no. 1, pp. 266-271 (2010).
3. T.R. Oldham, M.R. Friendlich, A.B. Sanders, C.M. Seidleck, H.S. Kim, M.D. Berg, and K.A. LaBel, TID and SEE response of advanced Samsung and Micron4G NAND flash memories for the NASA MMS mission, IEEE Radiation Effects Data Workshop Record, pp. 114-122 (2009).
4. H. Shindou, S. Kuboyama, T. Hirao, and S. Matsuda, Local and pseudo SELs observed in digital LSIs and their implication to SEL test method, IEEE Trans. Nucl. Sci., vol. 52, no. 6, pp. 2638-2641 (2005).
5. C. Poivey, H. Kim, and R.L. Ladbury, SEL Testing of the 512 Mb SDRAM from Maxwell, T121804_SDRAM_Testchip, http://radhome.gsfc.nasa.gov.