Single-Event Upset in Flash Memories†

H. R. Schwartz, D. K. Nichols and A. H. Johnston
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Abstract

Single-event upset was investigated in high-density flash memories from two different manufacturers. Many types of functional abnormalities can be introduced in these devices by heavy-ions because of their complex internal architecture. Changes in the stored memory contents sometimes occurred, even when devices were irradiated in a read mode with the internal charge pump inactive. For one device technology, unusually high currents were observed during post-irradiation cycling that were high enough to cause catastrophic failure.

I. Introduction

Flash memories provide an alternative high-density commercial storage technology that can be used in selected applications that do not require frequent write operations. Flash memories are expected to be inherently more resistant to soft errors than DRAMs, but newer flash devices use complex internal architectures to improve write and erase time, as well as to make device operation more transparent to the user. An internal write-state machine is used to control write sequences, which require erasure at the block or whole-chip level before writing. The erase process is inherently slower than in other memory technologies. Flash memories are designed to provide fast read access, comparable to that of DRAMs. However, writing is a more complex operation in flash technologies, requiring much longer time intervals [1-4]. A high voltage – 12 to 20 V -- generated with an internal charge pump, is required to write to the storage cells. Read operations require only the normal supply voltage (3.3 or 5 V, depending on the device design).

Total dose effects on older nonvolatile technologies have been investigated [5-7], but much less is known about how single-event upset will affect such devices. This paper discusses single-event upset effects in two basic types of flash memories that use fundamentally different storage cell designs. The dominant problem for these devices is upsets in the complex internal control circuitry rather than upset in the memory elements. Many of these upsets produced conditions that could only be restored by removing power and reinitializing the devices. Frequently the functional errors also caused part of the array to be rewritten, even when the device was irradiated in the read mode, with no external voltage applied to the write circuitry and the charge pump inactive.

From the standpoint of testing and applications, flash memories have some parallels with microprocessors because they use complex internal architectures, and the visibility of internal changes induced by single-event effects can depend on the way that the part is applied. Special design features of flash devices include the use of thin oxides and high electric fields, and the need for some sections of the device to provide logic and control of the high internal voltages required for erasure (this requires internal high-voltage logic). Thus, some facets of their design and operation are unique, and may present new issues from the standpoint of radiation vulnerability.

II. Flash Memory Technologies

The basic structure of a flash-memory cell uses a dual sandwiched gate structure, interposing a floating gate between the body of the device and the control gate. It is similar to the structure of an EEPROM, but uses a much thinner oxide between the floating gate and channel region. The thin oxide allows charge to be transferred to and from the floating gate by either of two mechanisms: Fowler-Nordheim (F-N) tunneling from the source or body, or hot-electron injection from the channel region. The cell technologies of the two manufacturers both use F-N tunneling for erasure, but use different write mechanisms. Further details of their respective device technologies are discussed below.

A. Cell Structures and Fabrication Technology

Intel NOR Structure

Devices produced by Intel use a NOR cell configuration, which allows random access of individual cells. This approach minimizes access time compared to the NAND structure. Erasure is done at the block (or entire chip level in older Intel flash memories) by applying the high programming voltage (V_{PP}) to the source, grounding the control gate and allowing the drain to float, as shown in Figure 1(a). Charge in the floating gate is transferred to the source by Fowler-Nordheim tunneling. Programming is done by grounding the source, and applying V_{PP} to the control gate. During programming, charge is transferred to the floating gate by hot electrons from the channel (CHE), as shown in Figure 1(b). The oxide between the floating gate and channel is very thin, ≈ 100 Å, to enhance the tunneling mechanism. The Intel technology requires a programming voltage of 12 V.

†The research in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration, Code Q, under the NASA Microelectronics Space Radiation Effects Program (MSREP).

0018-9499/97$10.00 © 1997 IEEE

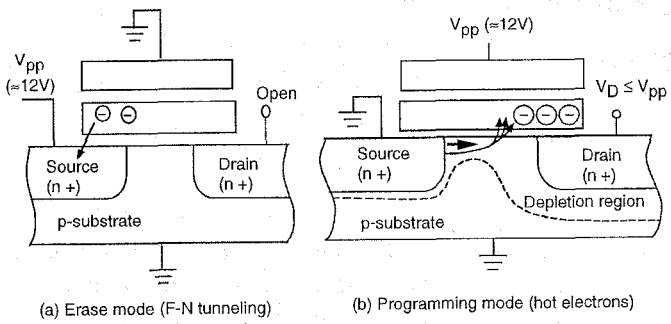

Figure 1. NOR Flash Structure Used in the Intel Devices

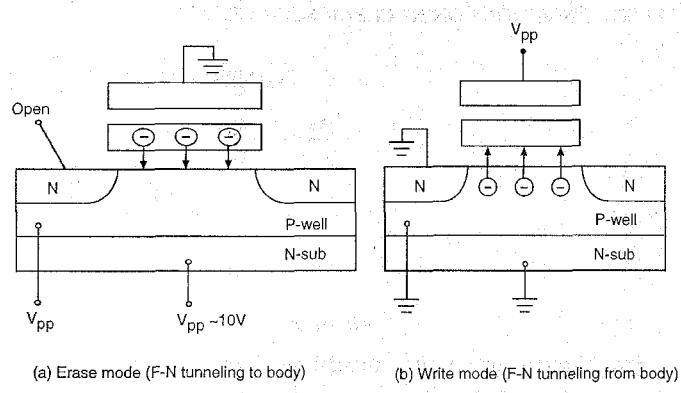

Figure 2. NAND Flash Structure Used in the Samsung Devices.

Because of variations in the tunneling current and threshold voltage, a single write sequence does not always transfer sufficient charge to the floating gate to ensure a "1" condition. An internal comparator is used to read each block or byte at the end of the write sequence, repeating the write operation several times if necessary (verify circuit). This section of the device must withstand the entire programming voltage, and uses high-voltage transistors [2,8].

Newer flash memories produced by Intel can operate at either 3.3 or 5 V. A boosted word-line voltage is required with the lower power supply voltage; internal circuitry detects the voltage and automatically applies the boost voltage [8]. Some versions use "smart voltage" logic that automatically detects the power supply and programming voltage. The internal charge pump is automatically activated when a voltage is applied to the programming pin.

Older Intel flash memories were fabricated on bulk substrates, and were sensitive to latchup from heavy ions. However, newer devices are fabricated on epitaxial substrates. Some older devices (notably the 1-Mb 28F010) have been shifted to more advanced processes, and are no longer sensitive to latchup. Flash memory fabrication technologies change rapidly to meet performance and cost pressures in the commercial marketplace. Archival data on radiation response has only limited value for these classes of devices because of the frequent changes in technology.

Samsung NAND Structure.

The Samsung devices use NAND logic, stacking eight cells in series with a common bit line. This allows a more compact cell structure to be used, because separate contacts are not required to each source and drain region. However, the read time is inherently slower in this technology because cells cannot be accessed individually; the read path goes through other cells in the stack. In order to deal with this, the device architecture divides the memory into pages. A page buffer is used to improve read time.

The Samsung NAND structure uses Fowler-Nordheim tunneling for erasing and writing, as shown in Figure 2. The oxide between the floating gate and body is ≈ 250 Å in this technology. This requires a higher programming voltage – 18-20 V -- compared to the NOR cells used by Intel [9]. In the NAND structure,, the programming voltage must be applied to the substrate and well regions to erase the cells, and then to the control gate during writing. This requires internal high-voltage logic circuits. The write path applies a uniform field across the entire oxide, with more uniform charge transfer compared to CHE writing (Intel). This can potentially improve the number of allowable write/erase cycles.

B. Device Architectures

The basic architecture used by the Intel 16-Mb devices is shown in Figure 3. Most of the complex architecture, contained in the command and write state machines, is devoted to improving the apparent write time by interposing page buffers and successive execution commands [2,8]. The command state machine is microprogrammed, making it potentially susceptible to single-event upset, with possible conflicts in internal device operation and sequencing. Because of the complex internal architecture, single-event upset effects in flash memories can be very difficult to interpret, similar to upset modes in microprocessors. The upset response also depends on operating conditions. Not shown in Figure 3 are page buffers that are used to allow writing at the page level, as well as registers that control the data stream and addressing.

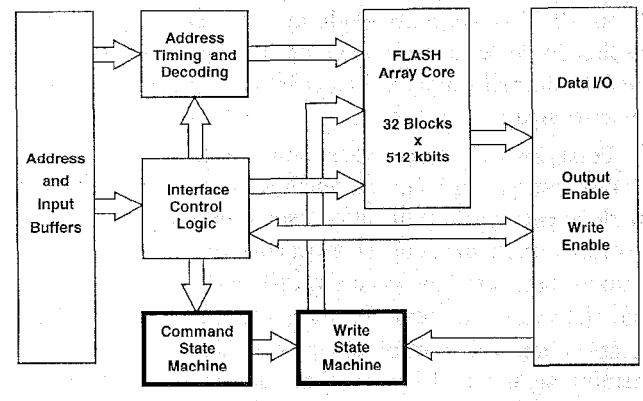

Figure 3. Block Diagram of the Intel 28F016SA Flash Memory

Fortunately the read path for these devices is much more straightforward, and is not affected by changes in the write state machine. Figure 4 shows a simplified diagram of the read path. Although one could get some apparent errors in the array because of changes in the state of the latch circuits or from row/column address changes, this type of error would not persist beyond a single read cycle. Thus, errors that occur on several successive reads can be clearly identified as changes in

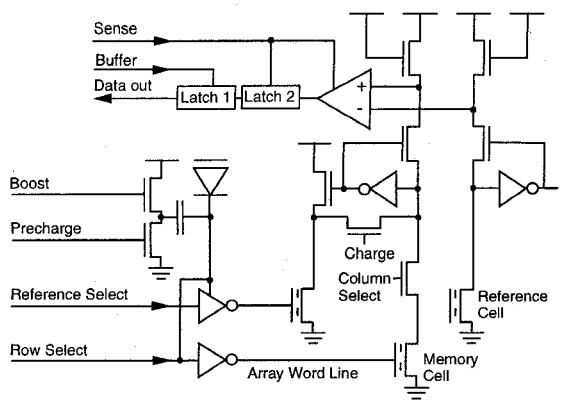

Figure 4. Block Diagram of Read Path Used in Intel Devices

the state of internal memory locations. The NOR cells allow individual bits to be accessed during read cycles. Writing is done at the page or block level, and erasing is done at the block level.

The architecture of the "smart-voltage" (SV version) of the 16 Mb Intel part is more involved, and details are not available from the manufacturer. Several additional functions are added in this version of the part, including the ability to queue a sequence of commands, and an automatic power saving feature which reduces the power during periods when the addresses are not switching. The area of the control circuitry is about 50% larger for the SV part compared to its SA counterpart. However, the array size appears identical for the two Intel device types.

The architecture of the Samsung devices is shown in Figure 5. Less detail is provided for the Samsung devices than the architecture provided by Intel. As noted earlier, the NAND cell structure does not provide access to individual cell locations. The internal architecture provides serial access to cells within a given structure, using pages of 264 bits (eight extra bits are provided in each page). All programming and read operations are done at the page level; individual cell access is not possible for NAND flash architectures.

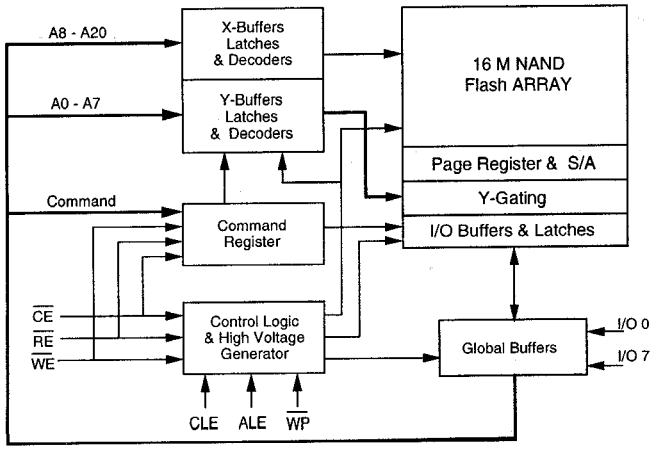

Figure 5. Block Diagram of the Architecture of the Samsung Flash Memories

The entire memory array is subdivided into blocks of 16 pages (each block has 4k bytes of memory). Erasure is done at the block level.

Address lines are multiplexed in the Samsung architecture. Another important difference is that the control logic in the Samsung architecture also affects the read path because of the requirement for serial access. This makes it possible to have control-like errors during the read mode with the Samsung device which lock up the read sequence, and in fact this occurred in some instances (read-lockup never occurred for the Intel devices). Resetting or repowering the Samsung device provides a way to restore normal read operation, and can generally be used to distinguish between cases where the storage array is altered by single-event effects from cases where the apparent errors are caused by lockup during the read mode.

For both manufacturers, the read mode is relatively straightforward, with less dependence on the complex internal control sequences that are used for erasing and writing. This makes it possible to distinguish between single-event errors in the array and decoding sections, and errors in the command state and write state machines.

C. *Specific Devices Used*

Four different flash memories were used in this study. Their basic characteristics are shown in Table 1, along with that of an earlier 1-Mb flash memory from Intel.

The 16-Mb Intel devices can be used with several different combinations of power supply and programming voltage. The chip operating voltage can be either 3.3 or 5 volts, selected by logic on the SA-version, or automatically in the SV-version. Both device types can be programmed with either 5 or 12 V, resulting in four different combinations of operating and programming voltage (the internal charge pump boosts the voltage to 12 V when the external programming voltage is 5 V). In order to simplify this, tests were only done with a 12-V write condition. However, tests were done with both 3.3 and 5 V logic operating voltages. In addition to the different combinations of power supply and programming voltage, these devices can be operated in a reduced power condition (automatic in the SV-version), as well as in a "deep-sleep" mode with total power of only about 20 μA. No tests were done with devices in the latter mode. Both Intel devices were fabricated on epitaxial substrates, with an epi-layer approximately 3 μm thick.

The two Samsung devices use similar architectures, but the 32 Mb device uses a much more compact structure. The chip size is nearly identical for the two parts, despite the factor of two difference in storage capacity. Unlike the Intel devices, the Samsung parts use only a single 5-V power supply voltage. The internal charge pump generates the high voltage required to erase and write to the cells. The Samsung devices were fabricated on bulk substrates.

Table 1. Flash Memories Used in this Study

Device	Description	Manuf.	Date Code	Cell Technology	Erase/ Write	Special Features
28F010	1-Mb flash	Intel	- - -	NOR	Whole chip	
28F016SA	16-Mb flash	Intel	none (1996)	NOR	Block or page	
28F016SV	16-Mb flash	Intel	9524/9534	NOR	Block or page	Smart voltage" allows transparent operation with several power supply voltage options.
KM29N16000	16-Mb flash	Samsung	9530	NAND	Block	
KM29N32000	32-Mb flash	Samsung	none (1996)	NAND	Block	Scaled device with twice the density

III. Testing Approach

Flash memories can be used in five basic ways: (1) an essentially unpowered mode, which only applies power during the relatively short duration that the memory contents are used; (2) a continually powered standby mode, in which the device is ready to begin reading, but the address lines are static; (3) a read-only mode, which applies continuous power to the device, along with address, clock and control sequences for reading, but never applies power to the write circuitry; (4) a read-mostly mode, or powered static mode, which is similar to the previous mode, but applies voltage to the charge pump and may also include brief periods for active writing; and (5) a mixed read and write mode (EWR), which involves many write cycles so that write duty cycle is a significant fraction of the total use period. These modes are briefly summarized in Table 2 below.

Mode 5 is very unlikely for flash memories, because they are only guaranteed to operate for a limited number of write/erase cycles. Consequently less emphasis was placed on characterizing SEE effects during active write cycles than in Modes 2 and 3, which are the most common application conditions. Many of our heavy-ion tests were done with the

Table 2. Basic Operating Modes for Flash Memories

Mode	Description	Supply Voltage	Write Voltage
1	Unpowered except for brief "read" periods	Zero, except for brief intervals	Not applied
2	Static	Continuous; static addressing	Not applied
3	Active read	Continuous; active addressing	Not applied
4	Read-mostly	Continuous; active addressing	Zero, except for brief intervals
5	Mixed read and write (EWR)	Continuous	Frequent

device in a static mode during the irradiation period. This simulates a very low duty cycle read-only mode, eliminating the possibility that upsets in decoding or write logic might interfere with cell upsets during active-reads. Reading the memory contents after the irradiation is stopped provides a direct way to determine whether portions of the memory were changed during irradiation. Some tests were also done in the EWR mode (Mode 5).

Tests were also done with devices unpowered to determine whether heavy ions could introduce cell upsets or operational errors in unpowered devices. This never occurred, even for effective LET values up to 120 MeV-cm^2/mg.

Two different radiation facilities were used for tests of flash memories, the Texas A&M cyclotron, and the Brookhaven Van de Graaff. Testing was done in vacuum, using ions with LET values from approximately 7 to 60 MeV-cm^2/mg. Some tests were done at non-normal incidence to increase the effective LET, up to 120 MeV-cm^2/mg at a 60 ° angle.

IV. Test Results for Intel Devices

A. *Upset Affecting Operational Conditions*

When the Intel devices were irradiated in the static mode (charge pump voltage not applied), several types of functional errors occurred. The threshold LET was approximately 7 MeV-cm^2/mg. The cross section for functional errors was between 10^{-7} and 10^{-6} cm^2 for the various error modes. The cross section remained at about this same order of magnitude even when ions with much higher LET were used, indicating that the response is caused by upsets in a small localized region (or regions) of the microcontroller. The finite time period required to go through operational cycles to detect incorrect operation introduces a latency period that makes it difficult to determine the precise time at which the internal error occurred. This latency, combined with counting statistics, results in larger uncertainty in the cross section for this type of functional error than for conventional upsets involving arrays of registers or storage cells. However, the small cross sections are consistent with upsets in individual control bits or small control registers in the internal write state machine. Figure 6 shows how the cross section for functional errors depends on LET, within the visibility limits discussed above.

Typical functional error modes that were observed are described in Table 3. In the majority of cases these functional conditions interfered with normal device operation and continued until power was temporarily removed from the device and reapplied, after which normal operation could be resumed. Most of these conditions did not change the contents of the internal memory array, but locked up the internal controller. However, one error type (designated "row/column

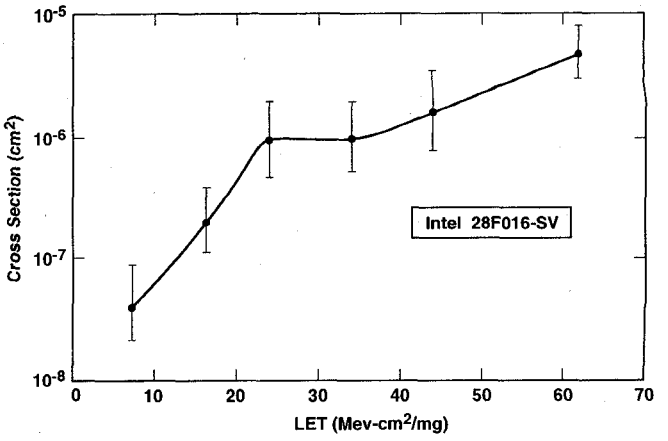

Figure 6. Cross Section for Functional Errors of the Intel SV Devices (Irradiated in Static Mode)

changes") caused portions of the array to be rewritten, even though the device was irradiated in the read mode.

The "slow block clear" errors initially caused the device to be stuck in a mode where the block clear status for a specific block could not be obtained. However, unlike "block clear lockup," the "slow block clear" condition could eventually be overcome by making numerous attempts -- typically 5 to 20 - to clear it. Clearing appeared to be related to the number of passes rather than the time interval after the functional condition was first observed.

Changes in internal memory locations can only be determined after the irradiation sequence has ended when devices are tested in the static mode. Memory errors could occur at any time during the irradiation run, along with complex functional errors. This prevents exact measurement of the cross section. However, by selecting a fluence which is low enough so that memory errors do not occur on all runs, it is possible to estimate the cross section for memory upset in a series of runs with reasonable counting statistics. The dependence of the estimated cross section for write errors on LET is very weak, and it flattens out at $\approx 10^{-6}$ cm^2, which is small compared total chip area. This strongly suggests that these internal memory changes are due to control circuitry (possibly inadvertent write commands), not fundamental interactions with the floating gate. Results with the control section shielded, discussed later, provide further corroboration

Tests were also done with active writing (erase, write, read, or EWR mode) during the time that the device was irradiated. The duty cycle for these sequences was as follows: approximately half this time was required to erase the memory, 30% for reading, and 20% for writing. Note that the charge pump is active during a significant part of the time in this mode. Somewhat surprisingly, the response of the Intel devices were not very different in the EWR mode compared to tests done in the static mode. However, there was one important difference: inadvertent block erasures sometimes occurred in the EWR mode, but were never observed when devices were tested a static or read mode.

Table 3. Functional Error Modes Observed for the 28F016SV Flash Memory (Irradiated in Static Mode)

Error Type	Description	Recovery Method
Block clear lockup	Block clear complete status never appears	Power cycling
False block clear	One or more blocks show block clear, even though they are not cleared	Power cycling
Slow block clear	Many passes are required to establish block clear for one or more blocks	Wait (power cycling not req.)
Row/column changes	Large portions of the memory array change state within a short time period, accompanied by block clear lockup	Power cycling
Slow first address programming	After successful block clear, the first address takes many passes and a long time to complete. Subsequent addresses work OK.	Wait
Read lockup	Status bits indicate internal modes and instructions are active, when device is expected to be in the ready state.	Power cycling
Write lockup	DATA WRITE status bit stuck in write-error mode during write sequence	Power cycling

No stuck bits or permanent errors (other than the catastrophic high current problem discussed below) were observed for either of the Intel devices, even when they were irradiated with an effective LET of 120 MeV-cm^2/mg. This was true for both the static and EWR modes. In addition, no errors were ever observed when the Intel devices were irradiated in an unpowered mode, using the same ion species.

B. High-Current Conditions

Currents Observed During Irradiation

During irradiation, the power supply current would exhibit a sudden increase, followed by similar steps in current (either up or down) as the irradiation continued. In many cases, complex functional errors occurred when the device was functionally tested after irradiation, but this did not always occur. Figure 7 shows an example of the steps in power supply current that occurred during a sequence of two runs. In this example, clear lockup occurred in both runs. The last current value persisted after the irradiation was over until the device was functionally tested, or in some cases shut down and reinitialized. None of these current conditions, which typically ranged from 3 to 70 mA, caused permanent damage.

The steps in current may be caused by functional changes in the control and/or write state machine, which turn on conflicting logic sections of the device. Similar effects have been seen in FPGA devices [12], where the current steps are caused by conflicting states of two logic elements. Note however that in the FPGA case the currents do not recover, because they arise from permanent changes in the antifuses

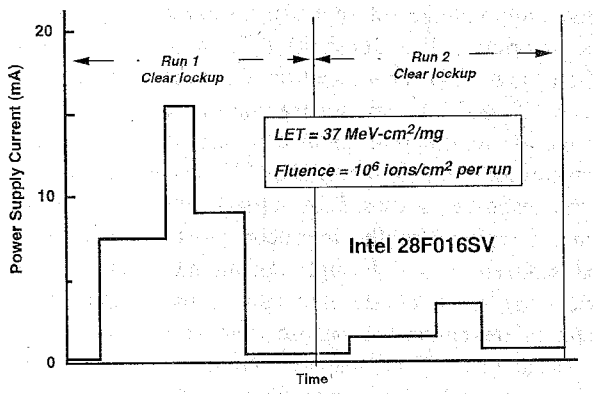

Figure 7. Examples of Steps in Power Supply Current for Intel Devices

that establish the logic pattern. For the flash devices, these types of currents always recovered, although in some cases power cycling was required to allow recovery.

Currents Observed After Irradiation

In addition to functional interrupt conditions, a high-current condition was sometimes observed after the device was irradiated in a static mode (no read or write operations during irradiation). The power supply current exhibited only the small changes discussed above during irradiation. After the irradiation, the power supply current jumped to very high levels at specific, reproducible address locations during the read cycle that was initiated after the beam was turned off. The current at these address locations exceeded 200 mA, and caused three devices to be destroyed.

For devices where the effect was not catastrophic, the high-current condition could be eliminated by power cycling. The effect was seen for both the SA and SV versions of the Intel devices. Although the precise cause of this effect was not determined, it is most likely due to logic conflicts in the address and buffers, which are affected by the command-state machine (see Figure 3).

C. Memory Errors (Write State Machine and Control Section Shielded)

In order to separate effects in the control section from memory upset, the microcontroller section of the memory was masked with 200 mils of copper. Figure 8 shows the location of the internal operating blocks, along with the section of the device that was shielded during tests of the array. Input buffers, address latches, interface control logic, and address/timing decoding were all shielded by the mask. However, the high voltage circuits used to write and erase were not covered by the shield, nor were the sensing circuits. None of the functional abnormalities discussed earlier occurred when the shield was in place, and it was then possible to observe single-event upset in the memory array.

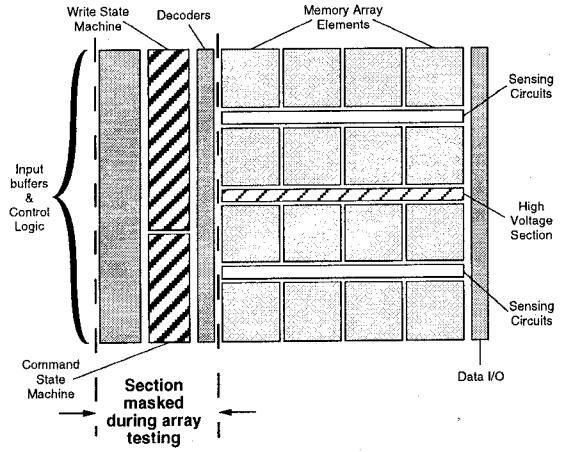

Figure 8. Diagram of the Intel SV Device Showing Location of Shield

Both types of Intel devices were tested with shields in place. No errors were observed for any of the SA-series, but errors in random locations, similar to "classic" memory errors, were observed for the SV (smart-voltage) devices. Memory upsets in the SV-type parts occurred at random address locations. Upsets were first observed at an LET of 44 MeV-cm^2/mg. These results are shown in Figure 9. The upper curve shows results with a power supply voltage of 5 V. The lower curve shows results with a power supply voltage of 3.3 V. Unlike most devices, this device is more susceptible to single-event upset when a high voltage (5 V) is used than when a low power supply voltage is used; this may be due to the way that the SV device is automatically reconfigured to erase, write, and read with different voltage conditions rather than a true difference in the upset sensitivity for the cells.

Because this effect only occurred for the SV device, it appears unlikely that it is caused by fundamental upsets in the memory array. The SV version contains additional registers and features which allow writing to be interspersed with erase cycles, increasing the effective speed of operation. These include data queue registers, additional page buffers, and automatic detection of voltage conditions which requires a more complex write-state machine compared to the SA-version. The memory arrays appear to be identical for the two devices, and the manufacturer's data sheets state that they are fabricated with the same process. Thus, it appears likely that the cell upsets that are observed in the SA-version occur because of the architecture. Although most of the control circuitry was covered by the shield, there are strips between the cell arrays that contain high-voltage logic and sensing circuits (see Figure 8). The diagram of the chip architecture shows that the sensing circuits provide bidirectional connections to the data queue register and the page buffers, providing a possible explanation for the difference in (apparent) memory cell upsets for the two device types.

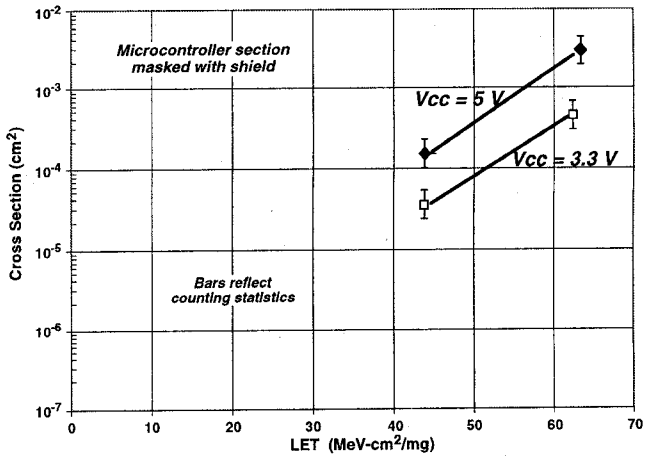

Figure 9. Upset Cross Section of the Intel SV Device, Write State Machine and Control Circuitry Shielded

VI. Test Results for Samsung Devices

A. General Results

Test results for the two Samsung devices were qualitatively very similar to the test results observed for the Intel devices,. Similar types of complex errors were observed, and in many cases it was necessary to cycle the power in order to reset the device. Table 4 shows the types of functional errors that occurred for the Samsung devices. Because of the page-mode architecture, less distinction can be made about some types of failure modes for the Samsung parts.

Read-mode lockup sometimes occurred for the Samsung devices, which never occurred for either Intel device. However, as discussed earlier, this is consistent with the different architecture used by Samsung. The cross section for functional errors was similar to that exhibited by the Intel devices, $\approx 10^{-6}$ cm^2.

Table 4. Examples of Functional Errors in the Samsung 16-Mb Flash Memory

Error Type	Description	Recovery Method
Row or column flips	Same values appear in multiple locations for the same bit position	Reinitialize and rewrite
†Lockups (self clearing)	Inability to progress beyond read, write, or block clear modes	Reprogramming (erase/write/read); power cycling not required
†Lockups (non-clearing)	Inability to exit read, write, or block clear modes	Power cycling followed by reprogramming
Stuck bits	Small number of bits permanently altered	None (permanent effect)

†Read lockups only occurred for the Samsung parts, not for the Intel devices. This is consistent with the page-mode architecture of the Samsung devices.

Write errors frequently occurred for the 16-Mb device when it was tested in the static mode, but errors were never detected in the 32-Mb device under a static or read condition. When the 32-Mb Samsung parts were tested in the EWR mode, errors did occur in the memory array, just as for the other three types of devices.

Significant changes in power supply current occurred for the Samsung devices during irradiation, just as for those from Intel. Figure 10 shows and example of current steps during four different irradiation sequences. Note that functional errors did not always occur when abrupt steps in current were observed. Neither of the Samsung devices exhibited the very large current during post-irradiation read cycles that were sometimes observed for the Intel devices. This may be related to differences in device architecture.

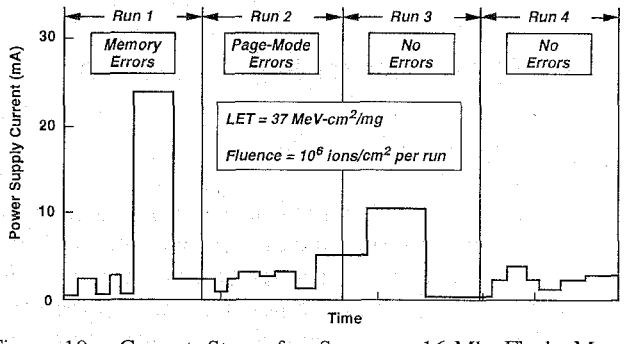

Figure 10. Current Steps for Samsung 16-Mb Flash Memory During Irradiation with Heavy Ions

B. Shielded Results

Both Samsung devices were also tested with the control section shielded. When tested in this manner in the read mode, no errors occurred in either device, even at LET of 60 MeV-cm^2/mg (only normal incidence could be used with the shield in place).

Tests were also done in the EWR mode. In that mode, random errors occurred in the 16-Mb device, as well as row-column bit flips. All of these errors could be restored by normal sequencing, and did not require power cycling. The cross section for such errors was much lower than that for the unshielded device, and it may indicate that not all of the control logic was completely covered by the shield.

No random errors were observed for the 32-Mb device even when it was irradiated in the EWR mode with the shield in place.

C. Stuck Bits

A small number of permanent errors -- 19 -- were detected in the Samsung 16 Mb devices after one run with bromine at Brookhaven (normal incidence; LET = 37 MeV-cm^2/mg). These errors were observed each time the device was read, even after the power supply was cycled several times. They were still present one week later, and could not be restored even by application of numerous erase/write/read sequence.

The number of errors dropped to 3 after a 24-hour annealing period at 100 °C. However, with repeated read cycles the number of errors increased, and the 19 errors that had originally been observed eventually reappeared. Some of the errors were intermittent, indicating a possible shift in threshold voltage for those specific cells. This indicates that the errors may be related to ionization damage rather than gate rupture. Several factors could produce such errors in these technologies, including microdose from heavy ions [13], and global total dose degradation which could make device operation uncertain for a small number of storage cells with marginal threshold voltage, either because of shifts in threshold voltage or slight degradation of the charge-pump circuitry.

Stuck bits only occurred for the 16-Mb devices, and did not occur for the more advanced 32-Mb part type.

VII. Discussion

The complex internal architecture of flash memories causes a great deal of difficulty in planning and interpreting single-event testing in these devices. Although the two device technologies examined in this work use very different cell technologies, individual cell upsets did not appear to be a significant factor for either technology, and consequently the usual way of looking at cross sections for memory devices cannot be applied to these devices. There are several reasons for this. First, the floating gates are isolated from the rest of the structure unless they are actually selected. Second, the high voltage required to write (purposefully or inadvertently) is not available unless the charge pump is active. Third, the magnitude of the charge stored on the floating gates is $\approx$ 1 pC [8,9,14], and write times of present technologies are quite long, tens to hundreds of microseconds. The floating gates cannot be accessed directly, only through the indirect processes of F-N tunneling or hot-electron injection. Consequently, it is unlikely that direct interactions of heavy particles in the thin gate regions (or in the underlying silicon regions) can discharge (or charge) the gate. However, such mechanisms may be possible if these devices are scaled further. First-generation flash structures required time periods $\approx$ 1 ms for writing, and this has decreased by two orders of magnitude as devices have evolved during the last five years.

One key point is that in three of the four device types, regions of the memory can be altered by SEE effects, even when the devices are irradiated in the static or read mode, with no voltage applied to the charge pump. Although the cross section is small, the total number of memory errors in a large array of devices could be a significant problem in system applications. Error detection and correction may be more difficult because of the number of complex functional errors that can occur, which may interfere with normal attempts to access devices in applications.

Although the threshold LET for upset in flash memories is significantly higher than that of DRAMs, the complex nature of the response of flash memories may be far more difficult to deal with from the standpoint of recovery, or application of error-detection-and-correction. Error correction would have to be implemented at the block level in order to correct for all types of errors, along with power cycling to recover from modes that caused functional errors, but did not produce memory errors.

The SEE response of older DRAM technologies was very straightforward, allowing error correction to be easily implemented. However, 4-Mb and newer generation DRAMs have complex response modes that have some similariities to the response modes of flash memories. In DRAMs these conditions occur because of special internal test modes, which can be inadvertently selected if the latch circuit that controls the test modes is upset. Architectures have been proposed for more advanced DRAMs that incorporate block control and sequencing that is much like the architecture used in flash memories [15]. This may be required to reduce power dissipation in newer devices. Thus, although present flash technologies appear to be more difficult to use in space than DRAMs, this may not be the case for more advanced DRAMs, particularly if they use a block architecture.

Flash memories are evolving rapidly, and the design and architecture of these devices continues to become more complex. New designs which store and detect more than one logic level in the floating gate (multilevel flash) have been demonstrated [16-18], and will further complicate testing and interpretation of SEE effects in flash technologies. Detection margins may be lower in multilevel flash technology, which may affect both SEE and total dose effects. The work in this paper shows how the architecture of flash memories affects the radiation response, and it is likely that even more complex operational errors will occur for more advanced devices.

VIII. Conclusions

Tests of these two flash memory technologies have shown that single-event upset effects are dominated by the complex architecture of these devices rather than upset in the storage array. The many different operating conditions and complex functional upset conditions that occur make it difficult to interpret single-event upset in flash memories. Recovery from many of the functional upsets could only be accomplished by removing power, and reinitializing the device, even when they were irradiated in a standby or read mode.

For three of the four device types, SEE effects during static or read cycles could also cause part of the memory to be rewritten. This is most likely due to upsets in the controller and/or write-state machine.

Abrupt steps in power supply current were observed for all four device types during irradiation. For devices from one manufacturer, very high currents were occasionally produced after the irradiation stopped, during reading, associated with a particular address condition. This caused catastrophic failure in three samples. Although the mechanism for this response was not determined, it may be due to logic conflicts from upsets in the controller.

Fortunately the upset cross section for complex errors in these devices was small, consistent with the cross section that is expected from individual control bits. Hence, the overall sensitivity of these parts to SEE effects is far lower than that of DRAMs. Furthermore, none of the devices appeared to be affected by heavy ions when they were irradiated in an unpowered mode, which is of interest in many applications that rely on the nonvolatile nature of these devices, with infrequent reading or writing operations.

Flash memories are evolving very rapidly because of their applications in high-volume consumer technology. Future changes in architecture, feature size, and floating gate technology may cause even more complex effects to occur in advanced devices. Devices with multiple storage levels have been proposed, which will provide another dimension in complexity.

REFERENCES

1. S. K. Lai, V. K. Dham and D. Guterman," Comparison and Trends in Today's Dominant E^2 Technologies," IEDM Technical Digest, pp. 580-582 (1986).
2. V. N. Kynett, A. Baker, M. L. Fandrich, G. P. Hoekstra, O. Jungroth, J. A. Kreifels, S. Wells and M. D Winston, "An In-System Reprogrammable 32 x 8 CMOS Flash Memory," IEEE J. Solid State Ckts., 23, 1157 (1988).
3. A. Concannon, S. Keeney, A. Mathewson, R. Bez and C. Lombardi, "Two-Dimensional Numerical Analysis of Floating Gate EEPROM Devices," IEEE Trans. Elect. Dev., ED-40, pp. 1258-1262 (1993).
4. M. Lanzoni, J. Sune, P. Olivo and B. Ricco, "Advanced Electrical-Level Modeling of EEPROM Cells," IEEE Trans. Elect. Dev., ED-40, pp. 951-957 (1993).
5. P. J. McWhorter, "Radiation Response of SNOS Nonvolatile Transistors," IEEE Trans. Nucl. Sci., NS-33, 1414 (1986).
6. D. F. Sampson, "Time and Total Dose Response of Non-Volatile UVPROMs," IEEE Trans. Nucl. Sci., NS-35, 1542 (1988).
7. P. V. Dressendorfer, Sandia National Laboratories, "An Overview of Advanced Nonvolatile Memory Technologies," presented in a short course at the 1991 Nuclear and Space Radiation Effects Conference, San Diego, California, July 15, 1991.
8. A. Baker, R. Alexis, S. Bell, V. Dalvi, R. Durante, E. Bauer, M. Fandrich, O. Jungroth, J. Kreifels, M. Lansgraf, K. Lee, H. Pon, M. Rashid, R. Rozman, J. Tsang, K. Underwood and C Yarigadda, "A 3.3V 16 Mb Flash Memory with Advanced Write Operation," Digest of Papers from the 1994 International Solid-State Circuits Conference, pp. 146-147.
9. K.-D. Suh, B.-H. Suh, Y.-H. Lim, J.-K. Dim, Y. J. Choi, Y. -N. Jigm S.-S. Lee, S.-C. Kwon, B.-S. Choi, J. -S. Yum, J.-H. Choi, J. -H. Choi, J. -R. Kim, and H.-K. Lim, "A 3.3 V 32 Mb NAND Flash Memory with Incremental Step Pulse Programming Scheme," Digest of Papers from the 1995 International Solid-State Circuits Conference, pp. 128-129.
10. K. Imamiya, Y. Iwata, Y. Suguura, H. Nakamura, H. Oodaira, M. Momodomi, Y. Ito, T. Watanabe, H. Araki, K. Narita, K. Masuda and J. Miyamoto "A 35 ns Cycle Time 3.3 V-Only 32 Mb NAND Flash EEPROM," Digest of Papers from the 1995 International Solid-State Circuits Conference, pp. 130-131.
11. D. Mills, M. Bauer, A. Bashir, R. Fackenthal, K. Frary, T. Gullard, C. Haid, J. Javanifard, P, Kwong, D. Leak, S. Pudar, M. Rashid, R. Rozman, S. Sambandan, S. Sweha and J. Tsang, "A 3.3-V 50 MHz Synchronous 16 Mb Flash Memory," Digest of Papers from the 1995 International Solid-State Circuits Conference, pp. 120-121.
12. G. Swift and R. Katz, "An Experimental Study of Heavy Ion Induced Dielectric Rupture in Actel Field Programmable Gate Arrays (FPGAs)," RADECS95 Proceedings, p. 425 IEEE Document 95TH8147.
13. C. Dufour, P. Garnier, T. Carrière, J. Beaucour, R. Ecoffet and M. Lebrunée, "Heavy-ion Induced Single Hard Errors in Submicronic Memories," IEEE Trans. Nucl. Sci., NS-39, 1693 (1992).
14. K. T. San, C. Kaya and T. P. Ma, "Effects of Erase Source Bias on Flash EPROM Device Reliability," IEEE Trans. Elect. Dev., ED-42, pp. 150-159 (1995).
15. K. Itoh, K. Sasaki and Y. Nakagome, "Trends in Low-Power RAM Technologies," Proc. of the IEEE, 83, pp. 524-543 (1995).
16. M. Bauer, R. Alexis, G. Atwood, B. Baltar, A. Fazio, K. Frary, M . Hensel, M. Ishac, J. Javanifard, M. Landgraf, D. Leak, K. Loe, D. Mills, P. Ruby, R. Rozman, S. Sweha, S. Talreja and K. Wojciechowski, "A Multilevel-Cell 32Mb Flash Memory," Digest of Papers from the 1995 International Solid-State Circuits Conference, pp. 132-133.
17. B. Eitan, R. Kazerounian and A. Roy, "Multilevel Flash Cells and their Trade-Offs," pp. 169-172, Digest of Papers from the 1996 International Electron Device Meeting.
18. M. Ohkawa, H. Sugawara, N. Sudo, M. Tsukiji, K. Nakagawa, M. Dawata, K. Oyama, T. Takeshima and S. Ohaya, "A 98mm^2 3.3 V 4Mb Flash Memory with FN-NOR Type 4-level Cell," Digest of Papers from the 1996 International Solid-State Circuits Conference, pp. 36-37.