Microelectronics Reliability 49 (2009) 707–715

Contents lists available at ScienceDirect

Microelectronics Reliability

journal homepage: www.elsevier.com/locate/microrel

A method to eliminate the event accumulation problem from a memory affected by multiple bit upsets

Juan Antonio Maestro *, Pedro Reviriego

Departamento de Ingeniería Informática, Universidad Antonio de Nebrija, C/Pirineos, 55, 28040 Madrid, Spain

ARTICLE INFO

Article history:
Received 17 December 2008
Received in revised form 23 February 2009
Available online 28 May 2009

ABSTRACT

Errors caused by radiation are a major problem for memories that have to operate in harsh environments. To evaluate the effects of radiation, memories are normally tested by exposing them to known radiation sources and measuring the number and shape of observed errors. If during the testing, errors accumulate in the memory, some of them can occur on adjacent locations and be interpreted as a multiple bit upset (MBU) instead of several independent events. This error accumulation makes difficult to characterize the number of events and error patterns that have affected the system. This paper presents an analysis of the error accumulation problem that enables the selection of radiation time values which can ensure that the effects of error accumulation are below a predetermined threshold. A method is also proposed to extract the correct error information from data that include a limited amount of error accumulation effects, therefore eliminating the event accumulation problem. Both techniques can be helpful to optimize the radiation testing of memories.

1. Introduction

As technology advances, the effect of radiation [1–3] on microelectronic circuits is becoming more important [4,5]. One of these effects is the production of soft errors, which have a temporary influence in the system. Single event upsets (SEUs) [6–9] are one of the most studied effects in this category, due to its high presence in applications. Memories [10–12], due to their broad use and large area, are specially sensitive to radiation effects [13,14]. That is especially true for multiple bit upsets (MBUs) [15–18], whose impact accounts for a growing number of effects. Therefore, one critical issue for manufacturers is to categorize how a certain memory behaves within different radiation environments and study the sources of potential errors [19,20].

In this situation, it is important to characterize the distribution of the different events induced in the system. That is to say, how many errors have been generated in the memory by radiation, how many were SEUs and how many MBUs, how many errors were produced by each MBU, the topology of these MBUs (the error patterns), etc.

However, calculating this information is not a trivial problem. In principle, MBUs are easily identified, since the errors in an MBU affect physically adjacent cells [17,21,22]. Therefore, isolated errors will be interpreted as SEUs and a group of closely located errors as an MBU. The issue is that a memory radiated for a long time or with high energies can present a large number of errors, which would tend to group and even overlap. Such groups of errors will be interpreted as an MBU instead of as a number of independent events. The overall effect of the error accumulation will be an overestimation of the number of MBUs and an underestimation of the total number of events.

In order to deal with this event accumulation problem, two approaches are proposed in the following. The first one predicts the effects of error accumulation as a function of radiation time, so that for a given level of tolerated error in the memory characterization process, the optimal radiation time can be calculated. The second approach tries to use the prediction of the effects of error accumulation to correct them so that the characterization of the memory can be done even when error accumulation is present.

The proposed approaches have been validated by simulation. The simulations reproduce the expected error events and patterns after a given radiation time and can therefore be used to statistically check the accuracy of the proposed techniques. The objective of both approaches is to enable the use of larger radiation times or larger fluence in the experiments. This provides more choices for radiation testing experiments and therefore may aid in the configuration of the experimental setup for memory characterization.

The rest of the paper is organized as follows: in Section 2, the problem of the error-per-event calculation when error accumulation exists is presented. In Section 3, two different approaches to deal with the error accumulation problem are proposed, for which some results are offered and commented in Section 4. Finally, conclusions and future work are presented in Section 5.

* Corresponding author. Tel.: +34 914524400; fax: +34 914521110.
E-mail addresses: jmaestro@nebrija.es (J.A. Maestro), previrie@nebrija.es (P. Reviriego).

0026-2714/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.microrel.2009.05.002

2. Problem definition: error-per-event calculation with error accumulation

As it was mentioned in the previous section, the goal of this paper is to propose methods to calculate the MBU distribution (number and topology) produced in a certain memory by a radiation source. In this way, the characterization of the technology in that particular environment can be performed.

Formally, let us use the term "error-per-event distribution" to refer to this MBU distribution, and let us denote it with the set $\{p(1), p(2), p(3), \ldots, p(\mathrm{n})\}$, being $p(n)$ the probability that an event with n errors has happened. In this way, $p(1)$ would be the probability that an SEU happens (only one error per event), $p(2)$ the probability that a 2-error MBU happens, etc. [20]. Calculating this distribution is very useful to foresee how a certain technology would behave in a given radiation environment and the probability that this technology is affected by multi-bit upsets. This has important implications for memory design. For example, interleaving is commonly used to protect against MBUs and the minimum interleaving distance required is given by the largest n for which $p(n)$ is significantly larger than zero [23,24].

One way to calculate the distribution would be to perform a radiation experiment on a memory (using a well-characterized source), and then examine the final memory map. By comparing this final map with the initial data, the set of erroneous positions (affected by the radiation) can be determined. However, the problem lies in estimating the nature of the events that produced those errors. This is because it is not possible to rebuild how events arrived in the memory, and therefore, if the final map shows e.g. four errors in adjacent positions (Fig. 1A), this could be due to four SEUs or to a 4-error MBU (or any other combination). This event accumulation problem has a very direct impact in the error-per-event distribution.

A possibility to perform this, addressed in [18,25], would be as follows: The experiment should be real-time controlled using a monitor platform, in such a way that periodical read-out/correct processes are performed in the memory. After the experiment starts, the first read-out process should happen at a pre-determined instant. In this process, the whole memory map is read and possible errors detected. If no errors, the experiment would continue until the next read-out process. If in one of these processes errors are detected, the topology of the errors is annotated and cleaned immediately afterwards. Then, when the experiment continues, the whole memory is correct again. If we can guarantee that the read-out processes are so frequent that only one event happens between every pair of them, then n errors observed at a certain instant would mean an n-error MBU, since events do not accumulate in the memory.

However, there may be situations in which the previous experimental methodology cannot be used. One possible scenario are environments where a real-time monitoring system is not feasible. For example, experiments to characterize radiation in remote areas, in-flight tests, etc. In this case, only the final memory map would be available, making the error-per-event calculation unfeasible. Another possible scenario are environments where a real-time monitoring system is feasible, but where it is not possible to guarantee the minimum read-out frequency so that events do not accumulate in the memory. This problem is also addressed in the literature [26], and it can be due to a very high event arrival rate (very intense radiation), which would make the read-out frequency unfeasible.

In those situations, some degree of error accumulation will occur and techniques to predict and correct its effect on the per-event distribution are needed. Two such techniques are proposed in the next section.

As a final remark, it should be noted that in all cases the total number of error events that the memory suffers during the testing will limit the accuracy in the estimation of $p(n)$. For example, if only 10 events are recorded in the experiments, small values of $p(n)$ like for example 0.001 would not be correctly estimated. For a given number of events N_e, the $p(n)$ distribution would be estimated in a reliable way when $N_e \gg 1/p(n)$. A large number of events can be obtained by repeating the test multiple times or by performing the test on many devices at the same time. This limitation also applies when estimating the effects of error accumulation on $p(n)$, as it will be discussed later.

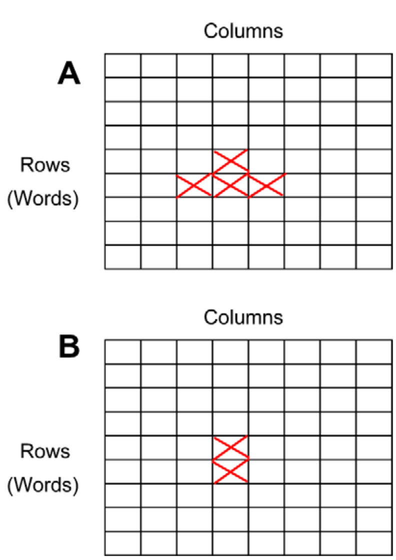

Fig. 1. A, false MBU example and B, false MBU example.

3. Techniques to deal with error accumulation

In this section, two approaches to deal with the error accumulation problem in the calculation of the error-per-event distribution are presented. In both the effects of error accumulation on the error-per-event distribution are estimated. In the first one, the estimation is used to select the radiation time so that those effects are kept below a given level, while in the second a more elaborate estimation is used to correct the effects of error accumulation.

In both cases, the discussion concentrates on the estimation of $p(2)$ and $p(3)$ as those are the most common types of MBUs [15–17]. Addressing only those values also helps keeping the focus on the proposed ideas avoiding unnecessary complexity. The ideas can then easily be extended for larger values of n following similar derivations.

The assumptions used in the derivations are the following:

- Error events are assumed to arrive following a Poisson process and they are equally distributed across all memory cells.
- It is assumed that a large number of experiments have been done. This is needed in any case to perform a good memory characterization.
- It is also assumed that there is a limited amount of error accumulation. More details on this assumption are given at the end of this section.
- MBUs are assumed to be formed by a number of events all of which are adjacent. This is consistent with the spatial distribution of MBUs expected in most cases [21].
- Error patterns for MBUs are assumed to be equally likely. When this is not the case, the approximations presented in

the following can be easily adapted to account for those differences.

- The memory is assumed to have a large number of columns and rows so that it is acceptable to assume that all cells are surrounded by other cells. Therefore, the effects due to cells on or close to the edges are neglected. It is straightforward to modify the proposed approximations to account for those effects.
- The effect of an error falling on a cell previously affected by another error and, as a result, both errors are cancelled, is neglected. This effect is small compared to the others included in the approximations and in any case it can be easily added to the proposed approximations.

The last four assumptions are introduced to keep the equations and the reasoning simple, as the main goal of the paper is to present an initial approach to deal with error accumulation. The proposed approximations can be easily adapted to account for these effects.

3.1. Selection of the optimum radiation time

To deal with error accumulation, a methodology to deduce how the different errors came grouped into MBUs is required. Let us consider the memory map depicted in Fig. 1B, where the positions marked with an "x" denote erroneous positions affected by radiation. The first thought would be to consider that the memory has been affected by a double MBU. However, it could also be a "false" MBU, due to two independent SEUs. Depending on the case, we would have a different MBU count for the error-per-event distribution: in other words, the observed MBU distribution does not have to be the real error-per-event distribution.

Having defined in the previous section the $p(n)$ distribution (the probabilities of the different types of events), let us now define $p_m(n)$ as the observed value of $p(n)$ in a given experiment with k events. $p_m(\mathrm{n})$ will be $\hat{p}(n)$ plus the effects of event accumulation, where $\hat{p}(n)$ is the estimate of $p(n)$ in that particular experiment. We now define $\{p^*(2)|^k, p^*(3)|^k, \ldots, p^*(\mathrm{n})|^k\}$, being $p^*(n)|^k$ the probability of obtaining a false n-bit MBU (formed in reality by n independent SEUs) after k events in the memory. Then if we are able to predict $p^*(n)|^k$, it would be possible to remove those false events from $p_m(n)$ to obtain $\hat{p}(n)$.This is the basic idea behind the proposed technique described in the following. However, as we will see shortly, there are other effects of event accumulation for which prediction and correction is also needed.

Let us study, from a probabilistic point of view, how likely the occurrence of a false MBU is. In the simple case of two adjacent errors (Fig. 1B would be one of the possibilities of this case), the second SEU needs to have hit the adjacent area of the first one. If the memory has M words of L bits, this probability would be $8/(M{\cdot}L)$, since there are eight neighbors per bit position (for positions that are along the edges, this number is lower, but these cases are negligible if the memory is large enough, as assumed at the beginning of the section). This would determine the probability that a second SEU creates a false MBU with a previous one. In the case that there were two isolated errors, the probability that a third SEU creates a false 2-bit MBU with any of the two previous would be $(8{\cdot}2)/(M{\cdot}L)$, since now there are already two candidate erroneous positions to group with. Generalizing, the probability that the i-th SEU hits a position adjacent to any of the $i-1$ previous events (to form the false 2-bit MBU), can be bounded by (assuming $i \ll M \cdot L$):

$$p_{adj}(2)|^i \leqslant \frac{8 \cdot (i-1)}{M \cdot L} \tag{1}$$

Then, if an experiment with k events is designed, the average number of false 2-bit MBUs obtained in the system, $f(2)|^k$ would be:

$$f(2)|^k \leqslant \sum_{i=2}^{k} p_{adj}(2)|^i = \frac{4 \cdot k \cdot (k-1)}{M \cdot L} \tag{2}$$

And therefore, the probability of false 2-bit MBUs after k events (that is to say, $p^*(2)|^k$), will be:

$$p^*(2)|^k \leqslant \frac{f(2)|^k}{k} = \frac{4 \cdot (k-1)}{M \cdot L} \tag{3}$$

Analyzing expression (3), if k is much smaller than $M{\cdot}L$ (memory size), then $p^*(2)|^k$ will tend to 0. In other words, the effect of false MBUs is negligible, being the observed MBU distribution the actual error-per-event function.

Let us define the error margin of the experiment, e, as the fraction of MBUs that we tolerate to be false. Then, the maximum number of events that we can generate in the memory, according to (3), is:

$$p^*(2)|^k \leqslant e \Rightarrow e \geqslant \frac{4 \cdot (k-1)}{M \cdot L} \Rightarrow k \leqslant \frac{e \cdot (M \cdot L)}{4} + 1 \tag{4}$$

In this way, k is the threshold for the number of events: below that threshold the amount of false 2-bit MBUs is smaller than e. Therefore, the observed MBU distribution in the final memory map would be in fact the actual error-per-event distribution (with the error tolerance e). The same derivation could be done to characterize the rest of $p^*(n)|^k$ for $n > 2$. However, if k is such that the effect of false 2-bit MBUs is negligible, the effect of false MBUs with more bits will be even smaller.

From the previous discussion, if the number of events is kept below the value of k given by (4) then the effects of error accumulation on the per-event distribution will be smaller than e. If the event arrival rate per word, λ, is known for the radiation source, that would imply that the experiment should last less than $k/(\lambda{\cdot}M)$ units of time. Larger radiation times would cause a larger number of events and the weight of false MBUs would be noticeable, distorting the observed distribution. In summary for a given error tolerance and arrival rate, expression (4) can be used to select an optimum radiation time.

Before moving on to a more sophisticated approach to deal with error accumulation, it should be noted that in predicting the effects of error accumulation a number of assumptions have been used. For example, equally likely MBU error patterns have been assumed and also that groups of errors are interpreted as an MBU only when they are adjacent. For other memory configurations (for example isolated columns so that only vertical MBUs are possible) or other MBU grouping criteria (for example errors at a distance of two) the same ideas can be applied and similar formulas would be obtained.

3.2. Correction of error accumulation effects

Another approach would be to use larger radiation times and then correct the effect of error accumulation by predicting the values of $p^*(n)|^k$ and subtracting them from the observed distribution, in order to get the right values. As the predicted values of $p^*(n)|^k$ are in fact the average values, a large number of experiments are needed to ensure that the observed $p^*(n)|^k$ are close to the theoretically predicted average value. However, a large number of experiments is also needed to accurately estimate the values of $p(n)$ as discussed before and therefore this is assumed in the following.

As a first step, expression (3) can be used to estimate $p^*(2)|^k$ and subtract it from the observed value. However as the purpose of this second approach is to enable larger number of errors per measurement (k), other effects may start to be noticeable and should be accounted for.

3.2.1. Calculation of the false 3-bit MBU probability

One such effect is the occurrence of false triple MBUs, $p^*(3)|^k$. These can be estimated by noting that in order to produce a triple false MBU, two errors that are adjacent or separated by just one cell are needed before the third error occurs. For a given error there are 24 positions adjacent or at a distance of two, as shown in Fig. 2. So, for event j, the number of candidates, c, would be on average:

$$c = \frac{24 \cdot \sum_{i=2}^{j}(i-1)}{M \cdot L} = \frac{12 \cdot j \cdot (j-1)}{M \cdot L} \tag{5}$$

This would indicate the probability that a second error happens in the proximity of a previous one, leading to a situation in which a third error could cause a false triple MBU.

If the two previous errors are adjacent, the third error will have 12 or 10 (11 on average) candidate positions to create the false triple MBU. On the other hand, if the two previous errors are at a distance of two, the third one will have 1, 2 or 3 candidate positions (2 on average) that would create a false triple MBU (see Fig. 3).

There are eight adjacent error configurations and 16 error configurations at a distance of one, and therefore the probability of failure upon event j can be approximated as

$$P_f(j) \leqslant 12 \cdot \frac{j \cdot (j-1)}{M \cdot L} \cdot \frac{11 \cdot \frac{1}{3} + 2 \cdot \frac{2}{3}}{M \cdot L} = 60 \cdot \frac{j \cdot (j-1)}{(M \cdot L)^2} \tag{6}$$

The first product term is the average number of 2-error candidate configurations given by (5). The last product term is the average probability that a third error falls near any of these 2-error candidates, forming a false 3-error MBU. In this way, the average number of false triple MBUs up to event k is given by

$$\begin{aligned} f(3)|^k &= \sum_{j=3}^{k} P_f(j) \leqslant \sum_{j=3}^{k} 60 \cdot \frac{j \cdot (j-1)}{(M \cdot L)^2} \\ &= \frac{60}{(M \cdot L)^2} \cdot \left(\frac{1}{3} \cdot k^3 - \frac{1}{3} \cdot k - 2\right) \end{aligned} \tag{7}$$

Which can be approximated for $k \gg 1$ as:

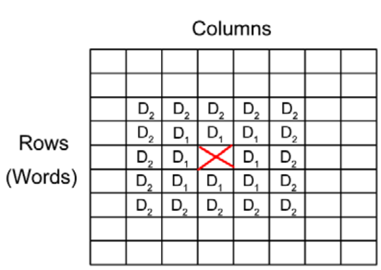

Fig. 2. Positions at distance of one or two of a cell affected by an SEU.

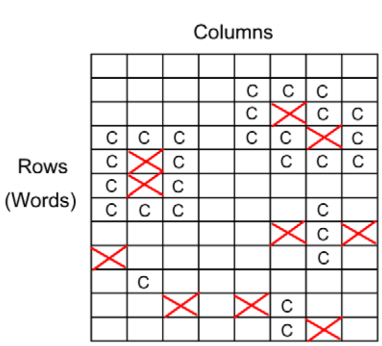

Fig. 3. Candidate positions to create a false triple MBU for different error configurations.

$$f(3)|^k \cong \frac{20 \cdot k^3}{(M \cdot L)^2} \tag{8}$$

And finally,

$$p^*(3)|^k = \frac{f(3)|^k}{k} \leqslant \frac{20 \cdot k^2}{(M \cdot L)^2} \tag{9}$$

With this, an estimate of the number of false triple MBUs has been obtained. The issue is that of those triple MBUs, some would have originated from false double MBUs (see Fig. 3), and therefore they should not be counted also as false double MBUs. Otherwise, we would be counting the same scenario twice. In the following, a technique to avoid this effect and properly determine the right number of false events will be presented.

3.2.2. Calculation of the correct number of events avoiding event accumulation

As a first approach, let us define $m(n)$ as the number of measured events of size n during the experiment. We will only consider double and triple MBUs since they are the most common type of events, as discussed before. An approximation of the total number of error events (N_e) suffered during the testing would be (ignoring events of more than three errors):

$$N_e = \sum_{i=1}^{3} m(i) \tag{10}$$

Using the previous equations, $f(2)|^{N_e}$ and $f(3)|^{N_e}$ can be expressed as

$$f(2)|^{N_e} \cong N_e \cdot p^*(2)|^{N_e} = \frac{4 \cdot (N_e - 1)}{M \cdot L} \cdot N_e \tag{11}$$

$$f(3)|^{N_e} \cong N_e \cdot p^*(3)|^{N_e} = \frac{20 \cdot (N_e)^2}{(M \cdot L)^2} \cdot N_e \tag{12}$$

With these values the approximation for the total number of error events can be refined as

$$N'_e = \sum_{i=1}^{3} m(i) + f(2)|^{N_e} + 2 \cdot f(3)|^{N_e} \tag{13}$$

This means that all the measured events ($m(i)$) are associated to at least one real event. However, each false 2-error MBU is formed by two independent SEUs, which is why an extra event needs to be added ($f(2)|^{N_e}$). In the same way, each false 3-error MBU is formed by three independent SEUs, implying two extra events ($2 \cdot f(3)|^{N_e}$).

The computation of the estimates of $p(2)$ and $p(3)$ is now performed subtracting the number of false MBUs from the number of measured events:

$$\hat{p}(2) = \frac{m(2) - f(2)|^{N'_e} + 2 \cdot f(3)|^{N'_e}}{N'_e} \tag{14}$$

and

$$\hat{p}(3) = \frac{m(3) - f(3)|^{N'_e}}{N'_e} \tag{15}$$

Note that the number of triple false MBUs is added twice to the estimate $p(2)$ as they are included as two events in $f(2)|^{N'_e}$ but are counted in $m(3)$ and not in $m(2)$.

Now that the estimates of $p(2)$ and $p(3)$ have been obtained, and ignoring false MBUs with more than 3 errors, then:

$$\hat{p}(1) \cong 1 - \hat{p}(2) - \hat{p}(3) \tag{16}$$

With the estimate of $p(1)$, another refinement can be done to $f(2)|^{N'_e}$ and $f(3)|^{N'_e}$:

$$f'(2)|^{N'_e} = \hat{p}(1)^2 \cdot f(2)|^{N'_e}$$
$$f'(3)|^{N'_e} = \hat{p}(1)^3 \cdot f(3)|^{N'_e} \quad (17)$$

With this, it is guaranteed that all the estimated events are formed by 2 or 3 independent SEUs (determined by $p(1)$).

3.2.3. Introduction of false 3-bit MBUs formed by the accumulation of a 2-bit MBU and an SEU

Up to now, only false MBUs formed by the accumulation of SEUs have been considered. However, the estimates of $p(2)$ and $p(3)$ in (14) and (15) can be further refined to account for the case in which a double MBU falls on a position adjacent to an SEU creating a false triple MBU. Let us define $f_{ds}(3)|^{N'_e}$ as the number of triple MBUs that are a combination of a real double MBU and an SEU. Following a reasoning similar to (11), this can be calculated as

$$f_{ds}(3)|^{N'_e} \cong \frac{11.5 \cdot (N'_e - 1)}{M \cdot L} \cdot N'_e \cdot \hat{p}(2) \quad (18)$$

The expression is weighed by the estimation of $p(2)$, in order to force the existence of a double MBU in the pattern. The value 11.5 is the average number of candidates to form a false MBU, as explained in previous cases. Also, it has been assumed that all real MBUs patterns of size two are equally distributed. That may not be the case in some memories depending on their physical structure. In those situations, (18) can be modified to account for the specific probabilities of the different patterns. Those probabilities can be predicted to some degree from the memory design or can be estimated in a first pass of the estimation of $p(2)$ and $p(3)$ in expressions (14) and (15), and used in the second phase.

With the introduction of $f_{ds}(3)|^{N'_e}$, a further refinement of the total number of events calculated in (13) can be proposed:

$$N''_e = \sum_{i=1}^{3} m(i) + f'(2)|^{N'_e} + 2 \cdot f'(3)|^{N'_e} + f_{ds}(3)|^{N'_e} \quad (19)$$

So that the refined estimates are finally:

$$\hat{p}'(2) = \frac{m(2) - f'(2)|^{N''_e} + 2 \cdot f'(3)|^{N''_e} + f_{ds}(3)|^{N''_e}}{N''_e} \quad (20)$$

and

$$\hat{p}'(3) = \frac{m(3) - f'(3)|^{N''_e} - f_{ds}(3)|^{N''_e}}{N''_e} \quad (21)$$

Since the proposed estimates are an approximation, it is interesting to explore the conditions on which that approximation is valid. The first condition can be derived from Eq. (1) in which it is assumed that each error is distant enough to the others, in order to obtain the number of candidates to form a false double MBU. For this to be true, there should be few error events in the memory. Eq. (3) can be used to that end by limiting the applicability of the approximation to, for example, values of $p^*(2) < 0.04$, which would mean that at most one error event should occur per 100 memory cells:

$$p^*(2)|^k < \frac{4 \cdot (k-1)}{M \cdot L} = \frac{4}{100} = 0.04 \Rightarrow k - 1 = \frac{M \cdot L}{100}.$$

With this consideration, Eq. (1) would be a reasonable approximation. The second condition is due to effects that are not modeled in the first phases of the estimation, like for example triple false MBUs caused by a double MBU overlapping with an SEU. This effect is proportional to the number of real double MBUs and one way to ensure that it is negligible is to apply the method only when the proportion of MBUs is low. For example, values of $p(2) < 0.05$ could be used. In summary, the approximation would be valid when both the proportion of cells with error and the percentage of MBUs ($p(n)$, $n \geqslant 2$) are low.

In order to characterize the memory, a large number of events are needed as previously discussed. Therefore, in many cases multiple experiments will be done. In that case the average of the estimates (20) and (21) of all experiments can be used as the final estimate. Another option is to add all the estimates of $f(2)|^{N_e}, f(3)|^{N_e}, N_e, f_{ds}(3)|^{N_e}$, and then use the additions to compute the estimates of $p(2)$ and $p(3)$. This is the approach used in this paper.

It should be noted that estimates of some of the variables presented in this section are mutually dependent. For example, the total number of error events depends on $f(2)|^{N_e}$ but also $f(2)|^{N_e}$ depends on the total number of events. The approach used in the derivations has been to use initial estimates of some variables to estimate others, and then refine the initial estimates with the new values. This is done for example with the total number of events N_e, N'_e, N''_e. The results presented in the next section show that this strategy seems to be sufficient in most cases. A more sophisticated approach can be used by formulating the solution as an iterative approximation problem in which the initial estimates are used as the starting point to recalculate refined estimates of $p(2)$ and $p(3)$, which are then used to iterate again until convergence to the final values is achieved. This approach may enable the correction of memories with higher degrees of error accumulation, what is left for future work.

To summarize this section, two approximations that correct the effects of error accumulation have been proposed: (14), (15) and (20), (21). The first one is simpler at the cost of reduced accuracy compared to the second. In the next section some simulation experiments are presented to illustrate the validity of the approximations.

4. Simulation results

In this section, the simulation work to put in perspective the proposed techniques will be described. A simulation script that reproduces the expected error events and patterns in the memory after a given radiation time has been developed to obtain the results presented in the following. In the simulation, the events are assumed to occur following a Poisson process which has been a commonly used assumption in memory reliability analysis [14]. The script generates random exponential inter-arrival times. The events that arrive before the radiation time specified in the simulation experiment are then inserted in a matrix that is used to model the memory. To achieve this, a row and a column are randomly selected (using a uniform probability distribution function) to determine the cell on which the error is inserted. Then, the type of event is randomly selected using the $p(n)$ specified in the simulation experiment. Finally, if the event is an MBU, its topology is also randomly selected. After all the errors are inserted, the number and types of errors in the memory are analyzed and used as the $m(n)$ for the proposed approximations.

The behavior of a memory with M = 512 K words and L = 12 bits has been simulated for a radiation associated to a constant event arrival rate per word λ = 1/400. In this way, the average number of events is a linear function of the radiation time (and vice versa). Alternatively, for a constant radiation time, the number of events is a linear function of the fluence, and therefore the results can also be interpreted in these terms. Notice that the value of λ would depend both on the memory technology and the radiation nature, and therefore it is not necessary to explicitly characterize them.

The purpose of the first simulation experiment is to determine how the number of false MBUs evolves with radiation time, and whether they are kept under the error tolerance e when the maximum number of events (radiation time) k is not reached. To perform the experiment, a value of e = 0.001 has been chosen, which

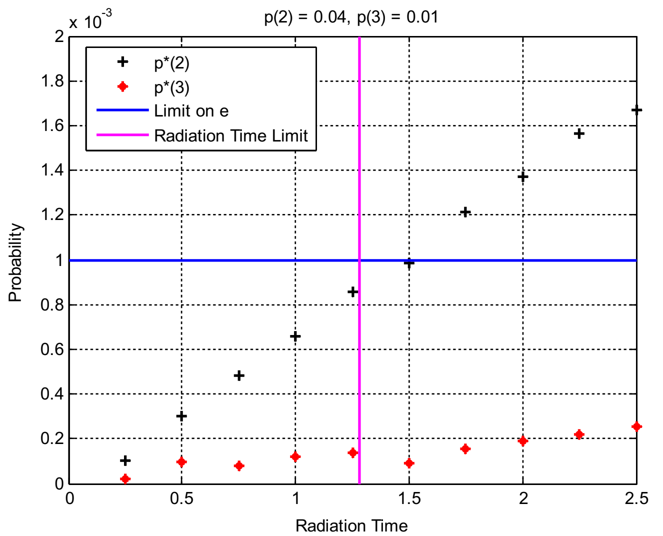

Fig. 4. Probability of false MBUs in the presence of real MBUs.

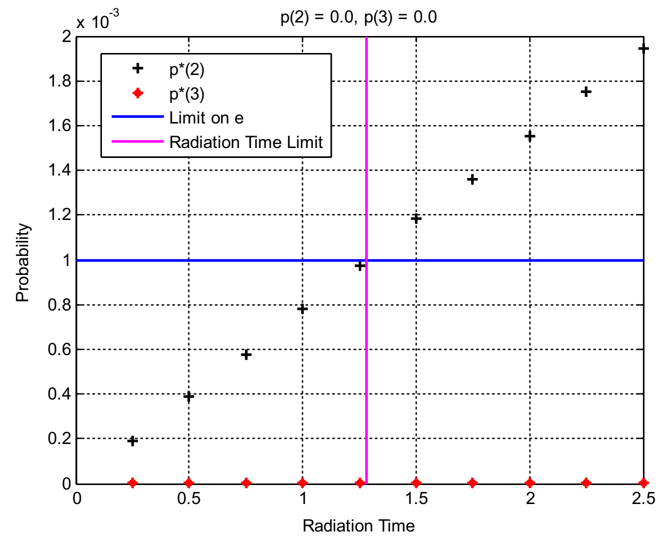

Fig. 5. Probability of false MBUs in the absence of real MBUs.

corresponds to a radiation time $t_r = k/(\lambda \cdot M) = 1.3$ units of time (approximately). It can be seen in Fig. 4 that while the radiation time is below that threshold, the probability of false 2-bit MBU is always below the selected e level. This means that since e is small, these false MBUs will have a negligible weight in the overall count of observed MBUs, and therefore these observed MBUs will mostly correspond to real MBUs (which would lead to the calculation of the per-event distribution $p(n)$). On the other hand, when k is exceeded, the false MBUs are higher than e, meaning that they will have a noticeable impact on the $p(n)$ calculation. Another thing to remark is when the probabilities of 3-bit MBUs and higher ($p(n)$ for $n \geqslant 3$) are smaller than $p(2)$, then the probabilities of false MBUs, $p^*(n)$ for $n \geqslant 3$ are also below e, as it can be seen in Fig. 4.

The results of a second simulation experiment are shown in Fig. 5. The same configuration of the first simulation experiment is used, but now in the case that no real MBUs are present (only SEUs). It can be seen that the result trend is similar to the one in Fig. 4. However, the value of the calculated false MBUs ($p^*(2)$ and $p^*(3)$) are more accurate than in the previous case (Fig. 5). This is because due to the absence of real MBUs, there are no side effects for complex combinations of them (e.g. a 2-bit MBU combined with an SEU to form a false 3-bit MBU). This can be directly seen in Fig. 5: the probability $p^*(3)$ is nearly 0 as expected (since there are no MBUs, the combination of 3 SEUs is needed to produce a false 3-bit MBU, which is unlikely). In the same way, the probability of $p^*(2)$ is also more accurate: the break point (where $p^*(2)$ meets the value of e) is exactly reached at the time threshold marked by k. However, although the first simulation experiment is less accurate, the model is robust enough to keep all $p^*(n)$ under the e level when the radiation time is not exceeded (Fig. 4).

In the third simulation experiment, the quality of the calculated per-event distribution, $p(n)$, has been measured for a scenario with $p(2) > 0$ and $p(3) = 0$. Different memory sizes have been tested in order to verify that the method is size-independent for $k \gg 1$. The observed values of $p(2)$ and $p(3)$ have been compared with the real ones in order to check if the radiation time limit effectively removes the presence of false MBUs. In Fig. 6, the results are offered when the radiation time is below the threshold defined by k. It can be seen that the observed MBUs coincide with the real ones for all memory sizes, what implies that no false MBUs were accumulated in the memory. On the other hand, the same results are depicted in Fig. 7, but now exceeding k. It can be seen that the real and observed $p(2)$ differ due to the false MBUs that have started to accumulate in the memory.

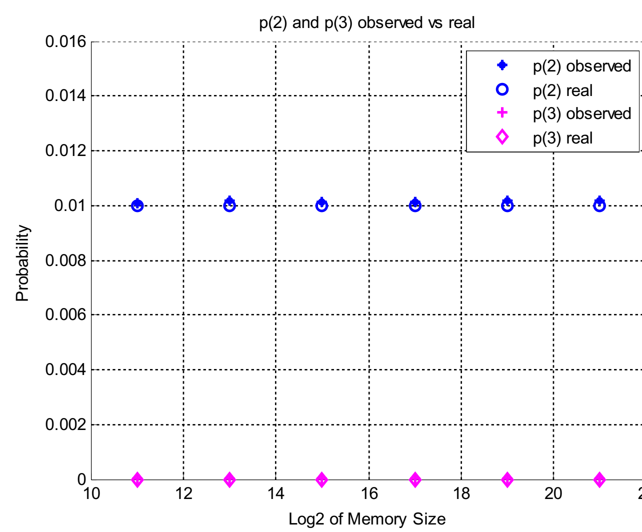

Fig. 6. MBU probability for a number of events below the threshold.

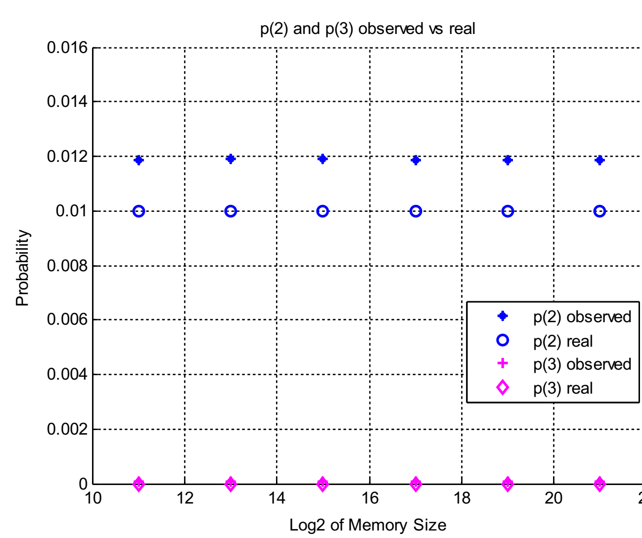

Fig. 7. MBU probability for a number of events over the threshold.

The third simulation experiment has then been used to check the proposed techniques that correct the effects of error accumulation (Section 3.2). The starting point is the result of the third experiment shown in Fig. 7 where error accumulation produced noticeable effects. In Figs. 8 and 9 the initial estimates (14) and (15) and the refined ones (20) and (21) are shown respectively and compared with the real values of $p(2)$ and $p(3)$. In this case, the simple approximation to correct the effects of error accumulation seems to be enough (both Figs. 8 and 9 offer similar results).

In the fourth simulation experiment, a higher value of the radiation time and different values of $p(2)$ and $p(3)$ are used and the results are shown in Figs. 10–12. In this case, the simple approximation (Fig. 11) reduces the error but it is still noticeable, while the refined approximation (Fig. 12) is able to correct most of the effects of error accumulation.

In the last simulation experiment, the validity of the approximations is tested. This is done using a wide range of radiation times to check if the approximation is valid when the percentage

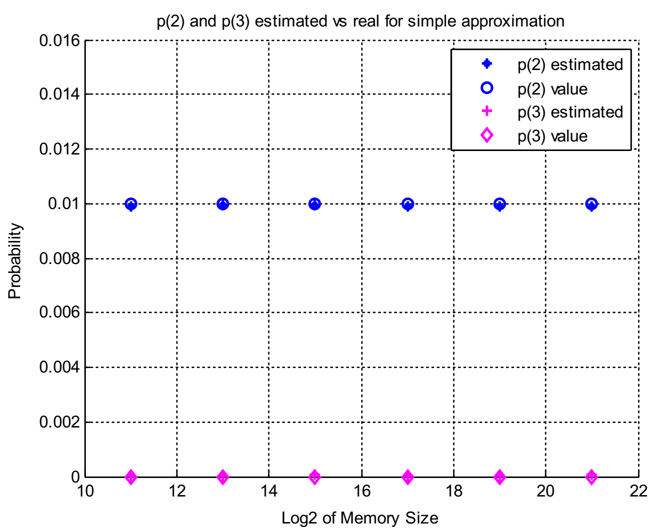

Fig. 8. Correction of the event accumulation using the initial estimates in expressions (14) and (15).

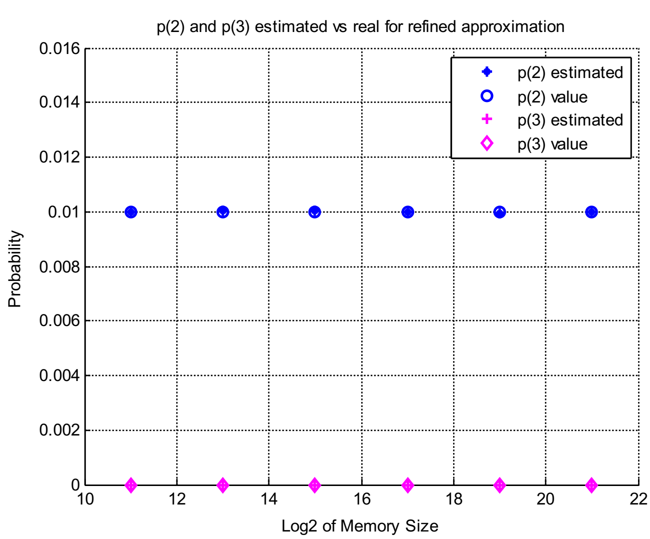

Fig. 9. Correction of the event accumulation using the refined estimates in expressions (20) and (21).

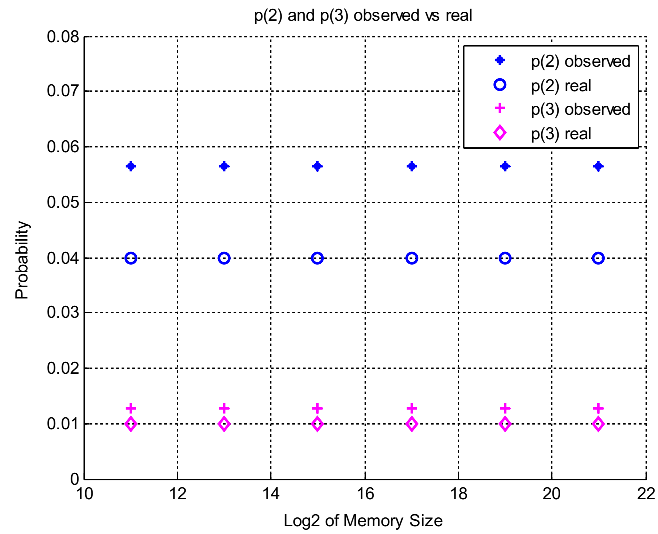

Fig. 10. MBU probability for a higher number of events over the threshold.

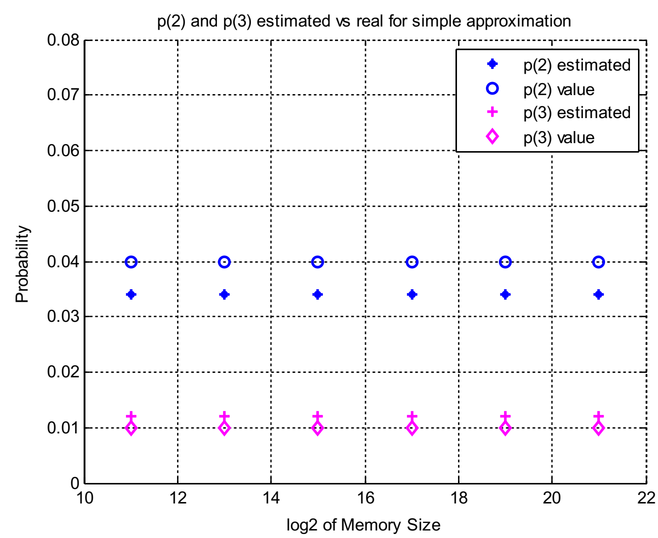

Fig. 11. Correction of the event accumulation using the initial estimates in expressions (14) and (15).

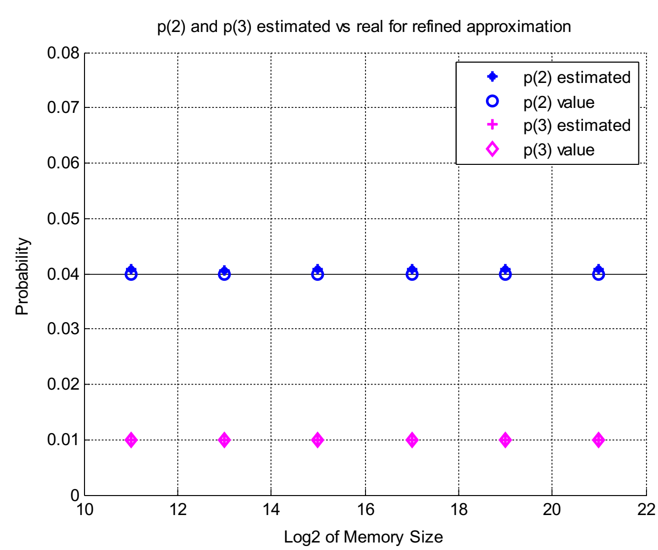

Fig. 12. Correction of the event accumulation using the refined estimates in expressions (20) and (21).

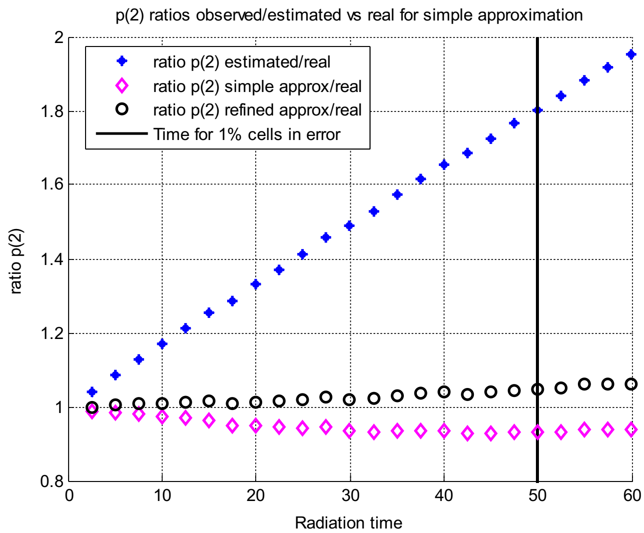

Fig. 13. Ratios of values of $p(2)$: observed/estimated versus real.

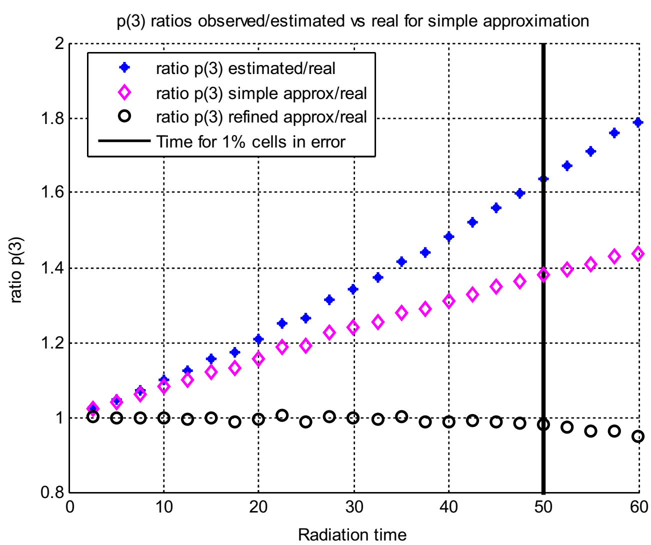

Fig. 14. Ratios of values of $p(3)$: observed/estimated versus real.

of cells with error is below 1%. The results are shown in Figs. 13 and 14 where the ratio of the observed or estimated versus the real value of $p(2)$ and $p(3)$ are plotted for different values of the radiation time. There is a direct relation between radiation time and the percentage of cells in error (expression (3)) and the limit assumed for the applicability of the approximation (1%) is also shown as a vertical line in the figures. The results show how, without the approximation, the ratio of $p(2)$ diverges from one, meaning that the estimated number of double MBUs differs greatly from the real values (Fig. 13). Then, the two approximations are shown to reduce the error significantly when the radiation time is such that the percentage of cells in error is low. For the ratio of $p(3)$, it can be seen that the simple approximation is only able to slightly reduce the error, while the refined one reduces the error significantly (Fig. 14).

In a general case, as discussed before, the accuracy of the first approximation would be better if the radiation time is not very large as in that case the initial error caused by error accumulation is small and the reduction of the error provided by the first approximation is enough. The accuracy of the initial approximation is also better when $p(2)$ and $p(3)$ are small, as in that case the refinements introduced in the second approximation would have little effect on the estimates of $p(2)$ and $p(3)$.

Besides from the presented results, additional simulation experiments have been performed using much larger radiation times, what implies a higher number of events arriving at the memory. It has been observed that when a certain number of events is reached, the proposed techniques to derive the correct error information are not able to eliminate the effects of event accumulation. This is due to the fact that for a large number of events, some of the approximations used in the derivations are no longer valid, since other phenomena start appearing in the system, e.g. errors overlapping with other previous existing errors or the appearance of false MBUs formed by a number of errors greater than three. These scenarios have not been deeply explored, since they are out of the scope of this paper.

5. Conclusions

In this paper, techniques to derive the error-per-event distribution (MBU distribution) on a memory radiated with a certain source when error accumulation is present have been studied.

The proposed techniques allow experiments with some degree of event accumulation, what enables larger radiation times (number of events). Two alternatives have been proposed. The first one estimates the effects of event accumulation and uses the estimation to select values of radiation time that ensure that those effects are below a given threshold. The second one uses more refined estimates of those effects to correct them allowing a reliable estimation of the per-event distribution when a larger amount of error accumulation exists. Both methods have been validated with a number of simulation experiments that show their applicability to different memory sizes and MBU distributions.

Given that the objective of the proposed approaches is to optimize the experimental setup for radiation testing of memories, the next step would be to apply the proposed methods to real radiation experiments in order to evaluate their advantages compared to the traditional approach that forces a negligible level of error accumulation. Additional work to extend the proposed approach to cover more general $p(n)$ distributions that include MBUs with larger multiplicity and the use of iterative approximations will also be considered for future work.

References

[1] Ziegler JF, Lanford WA. The effect of sea level cosmic rays on electronic devices. J Appl Phys 1981;52:4305–18.

[2] Gossett CA, Hughlock BW, Katoozi M, LaRue GS, Wendler SA. Single event phenomena in atmospheric neutron environments. IEEE Trans Nucl Sci 1993;40:1845–56.

[3] Tosaka Y, Kanata H, Itakura T, Satoh S. Simulation technologies for cosmic ray neutron-induced soft errors: models and simulation systems. IEEE Trans Nucl Sci 1999;46(3):774–9.

[4] Schrimpf RD, Fleetwood DM. Radiation effects and soft errors in integrated circuits and electronic devices. World Scientific Publishing; 2004. ISBN: 981-238-940-7.

[5] Dodd PE, Massengill LL. Basic mechanisms and modeling of single-event upset in digital microelectronics. IEEE Trans Nucl Sci 2003;50(3):583–602.

[6] Nicolaidis M. Design for soft error mitigation. IEEE Trans Dev Mater Reliab 2005;5(3).

[7] May TC, Wood MH. A new physical mechanism for soft errors in dynamic memories. In: Proceedings of the 16th annual international reliability physics Symposium; 1978. p. 33–40.

[8] Normand E. Single event upset at ground level. IEEE Trans Nucl Sci 1996;43:2742–50.

[9] Mavi DG, Eaton PH. Soft error rate mitigation techniques for modern microcircuits. In: Proceedings of the 40th annual reliability physics Symposium; 2002. p. 216–25.

[10] Baumann RC. Soft errors in advanced computer systems. IEEE Des Test Comput 2005;22(3):258–66.

[11] Schiano L, Ottavi M, Lombardi F. Markov models of fault-tolerant memory systems under SEU. In: Records of the 2004 international workshop on memory technology, design and testing. Issue 9–10; 2004. p. 38–43.
[12] Cardarilli GC, Leandri A, Marinucci P, Ottavi M, Pontarelli S, Re M, et al. Design of a fault tolerant solid state mass memory. IEEE Trans Reliab 2003;52(4): 476–91.
[13] Sasada T, Ichikawa S, Kanai T. Measurement of single-event effects on a large number of commercial DRAMs. IEEE Trans Nucl Sci 2006;53(4):1806–12 [Part 1].
[14] Goodman RM, Sayano M. The reliability of semiconductor RAM memories with on-chip error-correction coding. IEEE Trans Inform Theory 1991;37(3):884–96 [Part 2].
[15] Tipton AD et al. Multiple-bit upset in 130 nm CMOS technology. IEEE Trans Nucl Sci 2006;53(6):3259–64 [Part 1].
[16] Chugg AM, Moutrie MJ, Jones R. Broadening of the variance of the number of upsets in a read-cycle by MBUs. IEEE Trans Nucl Sci 2004;51(6):3701–7 [Part 2].
[17] Maiz J, Hareland S, Zhang K, Armstrong P. Characterization of multi-bit soft error events in advanced SRAMs. In: IEEE international electron devices meeting. IEDM'03 technical digest; December 2003. p. 21.4.1–21.4.4.
[18] Chugg AM, Moutrie MJ, Burnell AJ, Jones R. A Statistical technique to measure the proportion of MBU's in SEE testing. IEEE Trans Nucl Sci 2006;53(6): 3139–44 [Part 1].
[19] Maestro JA, Reviriego P. Study of the effects of MBUs on the reliability of a 150 nm SRAM device. In: Proceedings of the design automation conference (DAC'08), Anaheim (USA); June 2008. p. 930–35.
[20] Radaelli D, Puchner H, Wong S, Daniel S. Investigation of multi-bit upsets in a 150 nm technology SRAM device. IEEE Trans Nucl Sci 2005;52(6):2433–7.
[21] Satoh S, Tosaka Y, Wender SA. Geometric effect of multiple-bit soft errors induced by cosmic ray neutrons on DRAM's. IEEE Electron Dev Lett 2000;21(6):310–2.
[22] Tosaka Y, Ehara H, Igeta M, Uemura T, Oka H, Matsuoka N et al. Comprehensive study of soft errors in advanced CMOS circuits with 90/130 nm technology. In: IEEE international electron devices meeting. IEDM'04 technical digest; 2004. p. 941–44.
[23] Reviriego P, Maestro JA, Cervantes C. Reliability analysis of memories suffering multiple bit upsets. IEEE Trans Dev Mater Reliab 2007;7(4):592–601.
[24] Baeg S, Wen S, Wong R. SRAM Interleaving Distance Selection with a Soft Error Failure Model. In: Proceedings of the European workshop on radiation effects on components and systems (RADECS), Jyvaskyla (Finland); September 2008. p. 306–10.
[25] Johansson K, Ohlsson M, Olsson N, Blomgren J, Renberg P-U. Neutron induced single-word multiple-bit upset in SRAM. IEEE Trans Nucl Sci 1999;46(6): 1427–33.
[26] Bajura MA, Boulghassoul Y, Naseer R, et al. Models and algorithmic limits for an ECC-based approach to hardening sub-100-nm SRAMs. IEEE Trans Nucl Sci 2007;54(4):935–45 [Part 2].