

Influence of Randomness during the Interpretation of Results from Single-Event Experiments on SRAMs

Francisco J. Franco, Juan Antonio Clemente, Hortensia Mecha and Raoul Velazco

Abstract

After having carried out radiation experiments on memories, the detected bitflips must be classified into single bit upsets and multiple events to calculate the cross sections of different phenomena. There are some accepted procedures to determine if two bitflips are related. However, if there are enough bitflips, it is possible that unrelated pairs of errors appear in nearby cells and they are erroneously taken as a multiple event. In this paper, radiation experiments are studied as a special case of the urn-and-balls problem in probability theory to estimate how the measured multiple-event cross sections must be corrected to remove the overestimation due to the false events.

Index Terms

Birthday Statistics, Multiple Bit Upset, Multiple Cell Upset, Radiation, Single Bit Upset, Single Event Upset, SRAMs

I. INTRODUCTION

WHEN a semiconductor memory is hit by energetic heavy ions, bitflips lead to the corruption of the stored information due to the switching of the cell content. This phenomenon was discovered in 1975 [1] but it was not considered of special interest at that time. However, it has become evident that bitflips can compromise the success of space missions. For example, it was recently discovered that even electrons can produce bitflips in SRAMs [2], [3].

Initially, such errors were detected in isolated cells and this was called “*Single Bit Upset*” (SBU). However, during the following years it was observed that, in some cases, several bitflips appeared in adjacent cells. This was immediately interpreted as particles with enough energy to flip more than one cell. This phenomenon was called “*Multiple Event*”, and classified into two subtypes: “*Multiple Bit Upset*” (MBU), if flipped cells belong to the same word, or “*Multiple Cell Upset*” (MCU) if otherwise. A common term encompassing these three concepts is “*Single Event Upset*” (SEU).

Radiation ground tests are usually carried out to determine the cross sections of those phenomena. However, results are just a list of logic SRAM addresses where the content differs from a written pattern and it is difficult to determine if a bitflip is isolated or if it belongs to a complex multiple event. Several strategies have been developed to address this problem. For example, anomalously repeated distances between addresses have been used to distinguish MCUs from the bulk of SBUs [4]–[6]. In some other cases, the internal structure of the SRAMs is known so it is possible to establish a function relating the physical address with the logical one [5]–[8]. Thus, flipped cells are attributed to the same event if the distance between them is below a threshold value.

The problem is that two cells flipped in different events can be misled with a 2-bit multiple event if they are close enough. A careful experimental setup helps to avoid this problem. For example, using a very low radiation flux and reading many times the memory [9]. However, this can be sometimes unfeasible due to the requirement of long and expensive radiation experiments. In other cases [4], [10], results from experiments were treated using statistical approaches [11]–[13] to determine the probability of false multiple events depending on the memory size and number of bitflips. Among these methods, the most accepted one is that called “*the birthday statistics*” [12].

In this paper, we propose to identify the irradiation experiments with the classical urn-and-balls problem ruled by the combinatory distribution. Thus, we will provide simple equations that allow calculating: a) the number of observed vs. actual bitflips; b) the expected number of false MBUs; c) the expected number of false MCUs. Also, the compatibility of this method with the widely-used birthday statistics is demonstrated.

II. MATHEMATICAL PREDICTIONS

In the present work, we will suppose that only SBUs occur excluding other phenomena such as multiple events, microlatchups, etc.

This work was supported in part by the Spanish MINECO projects FPA2015-69210-C6-5-R and TIN2017-87237-P, and by the UCM mobility program for young professors.

F. J. Franco is with the Departamento de Estructura de la Materia, Física Térmica y Electrónica, Facultad de Ciencias Físicas, Universidad Complutense de Madrid (UCM), Spain, e-mail: fjfranco@fis.ucm.es.

J. A. Clemente and H. Mecha are with the Computer Architecture Department, Facultad de Informática, Universidad Complutense de Madrid (UCM), Spain, e-mail: {juananc1, hortent}@ucm.es.

R. Velazco is with the Université Grenoble-Alpes & CNRS, TIMA, Grenoble (France), e-mail: raoul.velazco@univ-grenoble-alpes.fr.

A. About the number of observed bitflips vs. actual bitflips

First of all, let us postulate that an experiment is performed on a SRAM with the following characteristics:

- N_A : Address bus width.
- L_A : Total number of addresses, with $L_A = 2^{N_A}$.
- W : Data word width. Typically, 8 or 16.
- N_W , such that $2^{N_W} = W$.
- L : Memory size in bits. It is immediate that $L = W \cdot L_A = 2^{N_A+N_W}$.

The memory is exposed to radiation in such a way that m SBUs occur. One can model the generation of the m bitflips as if every time a *successful particle* hits the memory, a cell is randomly chosen from a pool of L elements and flipped. It is important to highlight that nothing excludes the same cell to be hit several times so it must return to the pool of candidates for the subsequent hits. Indeed, this is just a version of the classic urn-and-balls problem in the probability theory [14]. Assuming that all the cells have identical probability to be chosen, the probability of a specific cell to be hit is just the total probability, 1 by definition, divided by the number of elements, L :

$$p_B = \frac{1}{L} \quad (1)$$

Following the urn-and-balls problem, the probability of a specific cell to be selected k times after m trials is ruled by the binomial distribution is:

$$\begin{aligned} P(k) &= \binom{m}{k} \cdot p_B^k \cdot (1 - p_B)^{m-k} = \\ &= \binom{m}{k} \cdot L^{-k} \cdot \left(1 - \frac{1}{L}\right)^{m-k} \end{aligned} \quad (2)$$

The number of trials, m , corresponds to the total number of bitflips and the size of the pool, L , the SRAM size. The total number of balls extracted k times, which corresponds to that of cells being hit k times, is:

$$N_{FB}(k) = L \cdot P(k) = \binom{m}{k} \cdot L^{1-k} \cdot \left(1 - \frac{1}{L}\right)^{m-k} \quad (3)$$

However, only those cells hit an odd number of times will be detected once the experiment has ended. If a cell is hit an even number of times, it will return to the initial state and will be indistinguishable from unaffected cells. In conclusion, the number of observed bitflips, n , will be the sum of cells affected an odd number of times:

$$\begin{aligned} n &= \sum_{k=0}^{\infty} N_{FB}(2k+1) = \\ &= \sum_{k=0}^{\infty} \binom{m}{2k+1} \cdot L^{-2k} \cdot \left(1 - \frac{1}{L}\right)^{m-2k-1} \end{aligned} \quad (4)$$

In practice, it is very unlikely that the same cell is hit three times or more. Therefore, keeping only the first element of the series in Eq. 4, for which $k = 0$, the previous expression can be approximated as:

$$\begin{aligned} n &\simeq \binom{m}{1} \cdot \left(1 - \frac{1}{L}\right)^{m-1} \simeq m \cdot \left(1 - \frac{m-1}{L}\right) \approx \\ &\approx m \cdot \left(1 - \frac{m}{L}\right) \end{aligned} \quad (5)$$

Rearranging this expression, the following equation appears:

$$m^2 - m \cdot L + n \cdot L \approx 0$$

This identity can be easily solved following quite simple procedures. Rejecting one of the solutions due to being nonsense and using Eq. 33, shown in the Appendix, for $x = -4n/L$, we finally get:

$$m \simeq n + \frac{n^2}{L} \quad (6)$$

Thus, one can estimate the number of cells that were actually flipped from the observed value issued from a radiation experiment. In general, the correction to be done to the number of observed bitflips (n) to get the actual number of events (m) is not significant, as shown in Fig. 1, calculated from Eq. 6. For example, in an 8-Mbit memory where 2000 bitflips were observed, the difference between the actual number of bitflips and the average number of observed events would be only $m - n = 0.48$. Therefore, we can assume that $m \equiv n$ in actual experiments.

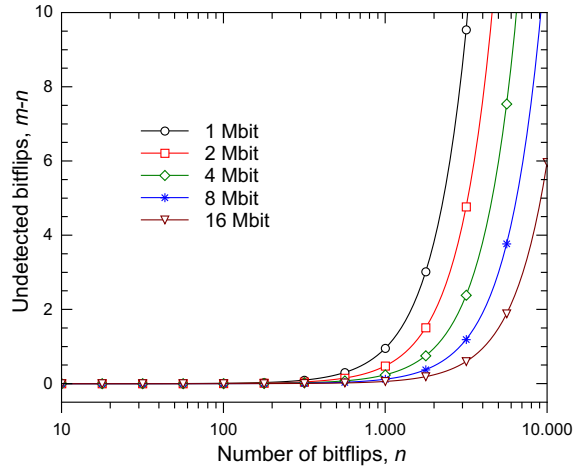


Fig. 1. Expected number of undetected bitflips, $m - n$, in function of the number of observed bitflips (n in Eq. 6) and the size of the memory, (L). Note the logarithmic scale in the X-axis.

B. About the number of false MBUs

This problem occurs when several bits in the same word are flipped due to independent events. This phenomenon can be mistaken with multiple bit upsets in such a way that the researcher may conclude that the MBU cross section is higher than what it actually is. The relevant question that should be answered is: after observing m bitflips in a L -size memory with W -width words, which is the expected number of false MBUs?

Two independent bitflips can be misled with an MBU if the involved cells belong to a single word, defined by a specific address. A simple way to calculate the number of false MBUs is to keep on using the urn-and-balls simile but selecting word addresses instead of single cells. Thus, a 2-bit MBU corresponds to word addresses chosen twice, 3-bit ones chosen three times, etc. This problem is exactly the same as the one depicted in the previous section but choosing addresses from a L_A -size pool instead of individual cells from the L -size one. Thus, the number of addresses flipped k times is equivalent to Eq. 3, but changing L by L_A :

$$\binom{m}{k} \cdot L_A^{1-k} \cdot \left(1 - \frac{1}{L_A}\right)^{m-k} \tag{7}$$

with $2 \leq k \leq m$. However, there are two important points to bear in mind. First of all, let us suppose that two independent particles not only hit the same word address, but also the same bit in it. After a double hit, the cell returns to its initial state and nothing will be detected. Therefore, it is necessary to introduce a new condition to Eq. 7: only counting those cases where cells are not hit several times.




Fig. 2. Possible combinations of two independent bitflips hitting a unique 8-bit word. Numbers ranging from 0 to 7 indicate the position of the cell inside the address. In white, combinations that cancel out the bitflips and restore the original content.

Fig. 2 illustrates the following statement. Let us suppose that an address pointing to an 8-bit word has been hit twice. This picture lists all of the possible combinations in which the false 2-bit MBU could take place according to the order of occurrence of both events. For example, if the first affected cell is #4 and the second one is #2, a false 2-bit MBU with bitflips in positions #2 & #4 is observed. The same will be observed if the positions swap. These cases are painted in grey to indicate that a 2-bit MBU is observed. However, if the same cell is hit twice, the original content is restored and the bitflips will never

be detected. These cases are shown in white in Fig. 2. Regarding the picture, one can see that there are $8^2 = 64$ possible combinations of which 8 pairs are not detectable. In short, only $64 - 8 = 8 \cdot (8 - 1) = 56$ cases out of 64 yield false 2-bit MBUs.

In general, if an address has been hit k times, the number of possible combinations is the volume of a hypercube with W units per dimension in the cartesian k -dimension hyperspace, W^k . Also, this is the number of possible permutations with allowed repetitions of k elements out of W . However, only the cases in which no elements are repeated will be observed as a k -bit MBU. According to the theory of combinatorics, only $\frac{W!}{(W-k)!}$ permutations out of W^k are without repetition. Including this correction, Eq. 7 eventually becomes:

$$N_{FMBU}(k) = \left(\frac{1}{W^k} \cdot \frac{W!}{(W-k)!} \right) \cdot \binom{m}{k} \cdot L_A^{1-k} \cdot \left(1 - \frac{1}{L_A} \right)^{m-k} \quad (8)$$

The second point is that, actually, Eq. 8 is not the observed number of false k -size MBUs since even more corrections should be included. For example, if an address is hit 4 times but twice in the same cell, it will be eventually observed as a false 2-bit MBU, not predicted by Eq. 8 for $k = 2$. Strictly, the equation should be corrected to include higher-order combinations. However, as the number of combinations is infinite but are very unlikely, it is postulated that Eq. 8 predicts the number of observed false MBUs with k -multiplicity in actual experiments.

Also, Eq. 8 fails if an address is hit more than W times. However, this is a situation extremely unlikely that will only occur if the number of bitflips is on the order of the memory size, so it will not be discussed.

Let us particularize Eq. 8 for 2- and 3-bit MBUs. As $(1 - L_A^{-1})^{m-k} \approx 1 - (m-k) \cdot L_A^{-1}$, Eq. 8 becomes:

$$N_{FMBU}(2) \approx \frac{1}{2} \cdot \frac{W-1}{W} \cdot \frac{m \cdot (m-1)}{L_A} \cdot \left(1 - \frac{m-2}{L_A} \right) \quad (9)$$

$$N_{FMBU}(3) \approx \frac{1}{6} \cdot \frac{(W-1) \cdot (W-2)}{W^2} \cdot \frac{m \cdot (m-1) \cdot (m-2)}{L_A^2} \cdot \left(1 - \frac{m-3}{L_A} \right) \quad (10)$$

Both expressions are drawn in Figs. 3(a)-(b) for illustrating purposes. Let us focus again in the case of the $1M \times 8$ -bit memory with 2000 bitflips. Eq. 9 predicts that 1.66 false 2-bit and $8 \cdot 10^{-4}$ false 3-bit MBUs can be detected. Clearly, there are serious chances of erroneously believing that 2-bit MBUs have occurred during the simulated experiment under these conditions.

C. About the number of false MCUs

Now, let us suppose that the researchers have information about the SRAM layout so they can relate the logical address of a bitflip with the XY-coordinates of the cell within the die. This information could have been obtained in several ways: perhaps the SRAM was designed by the researchers, the manufacturer provided proprietary information concerning the internal XY cell placement, or reverse-engineering was performed with a pulsed laser. This is extremely useful to determine how bitflips originated and if they are isolated or they constitute a multiple event. Thus, it is usually accepted that, if two flipped cells are closer than a reference distance, D , they belong to a multiple event.


However, in an only-SBU scenario, how many pairs of flipped cells can be erroneously classified as a multiple event? In general, cells are physically distributed within a rectangular array with the following characteristics:

- L_X is the horizontal size of the array.
- L_Y is the vertical one.

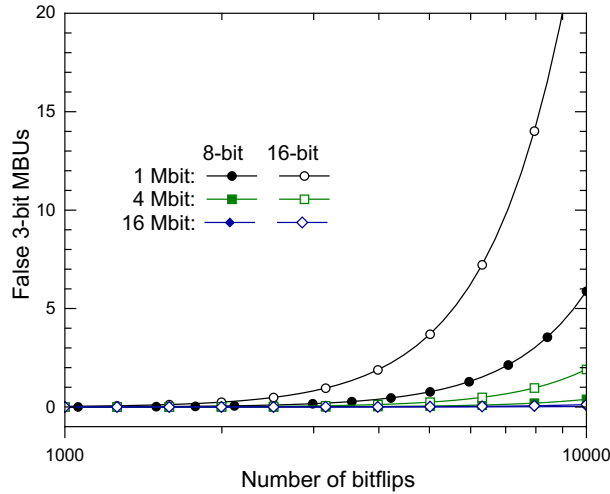
Obviously, $L_X \cdot L_Y = L$. In this rectangular array, it is frequently used the Manhattan distance instead of the Euclidean distance, more appropriate for the continuous real XY-plane. The Manhattan or Taxicab distance between two points, $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$, is defined as:

$$d(P_2, P_1) = |x_2 - x_1| + |y_2 - y_1| \quad (11)$$

Now, let us suppose that the cells in P_1 and P_2 were hit by independent particles. Their X-coordinates will be randomly chosen from the set $x_i = [0, 1, \dots, L_X - 1]$ with identical probability, $p_X = L_X^{-1}$. The first step will be to study the probability distribution of $X = x_2 - x_1$, $P_P(X)$. It is clear that X takes values from the set between $0 - (L_X - 1) = -L_X + 1$ and $(L_X - 1) - 0 = L_X - 1$. However, the shape of the probability distribution of this new variable has changed. It is possible to deduce that in the set of L_X^2 possible combinations of (x_1, x_2) , there are L_X pairs with $x_1 = x_2$ yielding $X = 0$ but only one pair, $(0, L_X - 1)$, yielding $X = L_X - 1$. In conclusion, $P_P(X = 0) \propto L_X$ but $P_P(X = L_X - 1) \propto 1$.



(a)



(b)

Fig. 3. Number of expected false MBUs according to the number of bitflips, size of the memory and word width: Eq. 9 for false 2-bit MBUs (a) and Eq. 10 for 3-bit ones (b).

Indeed, it is possible to demonstrate that $P_P(X) \propto L_X - |X|$ and, as $\sum_{X=-L_X+1}^{L_X-1} P_P(X) = 1$:

$$P_P(X) = \begin{cases} \frac{L_X - |X|}{L_X} & \text{if } X \in [-L_X + 1, L_X - 1] \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

The second step is to determine the probability distribution of the absolute value of X , $d_X = |X| = |x_2 - x_1|$, which be called hereafter P_X . The first fact is that the range of possible values is reduced to $[0, 1, \dots, L_X - 1]$. Also, $P_X(d_X \equiv |X|)$ is related to $P_P(X)$ in such a way that $P_X(d_X) = P_P(d_X) + P_P(-d_X)$ except in the case of $d_X = 0$, in which $P_X(0) = P_P(0)$. Therefore, the probability distribution of d_X is:

$$P_X(d_X) = \begin{cases} \frac{2}{L_X} \cdot (L_X - d_X) & \text{if } d_X \in [1, L_X - 1] \\ \frac{1}{L_X} & \text{if } d_X = 0 \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

A similar conclusion can be applied to the vertical distance along the Y-axis to obtain $P_Y(d_Y)$ by just changing L_X by L_Y . Now, let us use the expressions of $P_X(d_X)$ and $P_Y(d_Y)$ to determine the probability of the Manhattan distance between two

cells being $d = d_X + d_Y$. The total probability is the sum of all the possible cases where $d_X \in [0, 1, \dots, d]$ AND $d_Y = d - d_X$. Mathematically:

$$\begin{aligned} P(d) &= \sum_{k=0}^d P_X(k) \cdot P_Y(d-k) = \\ &= P_X(0) \cdot P_Y(d) + P_X(d) \cdot P_Y(0) + \sum_{k=1}^{d-1} P_X(k) \cdot P_Y(d-k) \end{aligned} \quad (14)$$

Using Eq. 13, $L = L_X \cdot L_Y$, and well-known mathematical identities shown in the Appendix at the end of the manuscript, this expression turns out to be:

$$P(d) = \frac{4}{L} \cdot \left[d - \frac{1}{2} \cdot \frac{L_X + L_Y}{L} \cdot d^2 + \frac{1}{6L} \cdot (d^3 - d) \right] \quad (15)$$

Two bitflips are supposed to be related if the distance between them is less or equal than a reference value, D , selected by the researcher so:

$$P(1 \leq d \leq D) = \sum_{d=1}^D P(d) \quad (16)$$

After some operations, it becomes:

$$\begin{aligned} P(1 \leq d \leq D) &= \frac{2 \cdot D \cdot (D+1)}{L} \cdot \\ &\cdot \left[1 - \frac{1}{6} \cdot \frac{L_X + L_Y}{L} \cdot (2D+1) + \frac{(D+2) \cdot (D-1)}{12L} \right] \end{aligned} \quad (17)$$

In the experiment, m bitflips were observed so the total number of possible pairs of addresses is:

$$N_{PAIRS} = \frac{1}{2} \cdot m \cdot (m-1). \quad (18)$$

Therefore, the total number of pairs of independent bitflips erroneously identified as related will be:

$$\begin{aligned} N_{FMCU} &= N_{PAIRS} \cdot P(1 \leq d \leq D) \approx \\ &\approx \frac{m \cdot (m-1) \cdot D \cdot (D+1)}{L} \end{aligned} \quad (19)$$

In a $1M \times 8$ -bit memory with 2000 bitflips and postulating that $D = 5$, Eq. 19 establishes that, on average, 14.3 pairs will be erroneously attributed to MCU on average.

It is also possible to use other distances instead of Manhattan. Another choice is the infinity-norm distance, defined as $d_\infty(P_1, P_2) = \max\{|x_1 - x_2|, |y_1 - y_2|\}$, which allows obtaining:

$$\begin{aligned} P_\infty(1 \leq d \leq D) &= \frac{4 \cdot D \cdot (D+1)}{L} \cdot \\ &\cdot \left[1 - \frac{1}{4} \cdot \frac{L_X + L_Y}{L} \cdot (2D+3) + \frac{D \cdot (D+2)}{4L} \right] \end{aligned} \quad (20)$$

from Eq. 13.

In both cases, it is easy to make a physical interpretation of the results. Fig. 4 shows generalized circles for different metrics. Excluding the center, it is possible to demonstrate that the number of cells at a generalized distance less or equal to D is $2 \cdot D \cdot (D+1)$ for Manhattan distance and $4 \cdot D \cdot (D+1)$ for infinity-norm. Therefore, the common factors in Eqs. 17 and 20 indicate the fraction of the SRAM cells in the vicinity of the center. The factors depending on $L_X + L_Y$ reflect border effects, either the asymmetry in the probability (Eq. 13) or the lesser contribution of cells near the SRAM borders to the total probability (the number of this group of cells is proportional to $2 \cdot (L_X + L_Y - 1)$, which is the total perimeter).




Fig. 4. Generalized circles for different metrics (radius from 1 to 3): Manhattan (a) and infinity-norm (b). A generalized d -radius circle is the set of points at a distance less or equal to d . Inside each square, the generalized distance to the center is displayed.

III. COMPARISON WITH THE BIRTHDAY STATISTICS

The birthday statistics method was initially proposed by H. J. Tausch in 2009 [12] based on a work by Abramson and Moser [15]. The latter demonstrated how likely was that, after picking p elements out of a set of n consecutive natural numbers, all of them were separated a distance *higher* than $k \in \mathbb{N}$. Initially proposed to investigate the distribution of birthday dates throughout a random set of people, Tausch extended its application to single-event experiments. Thus, n would be the size of the memory in bits, p the number of bitflips and k a threshold distance. He obtained a quite simple expression to determine the probability of collision of single bit upsets. As in the present manuscript, $n \equiv L$, $p \equiv m$, Tausch's expression becomes:

$$P_k(L, m) \approx 1 - \exp\left(-\frac{m \cdot (m-1) \cdot (2k-1)}{2L}\right). \quad (21)$$

This expression was experimentally demonstrated by Tausch himself and is nowadays widely accepted by researchers.

Similar expressions can be deduced for false MBUs & MCUs. Despite the fact of having determined the number of false multiple events from the binomial distribution, once the theoretical mean value of occurrences of a specific phenomenon, μ , is known, the probability of observing q events, with $q \in \mathbb{N}$, is determined by the Poisson distribution:

$$P(q, \mu) = \exp(-\mu) \cdot \frac{\mu^q}{q!}. \quad (22)$$

Thus, the probability of observing at least one event is:

$$P(q \geq 1, \mu) = 1 - P(0) = 1 - \exp(-\mu). \quad (23)$$

First of all, let us investigate the prediction concerning false MBUs. For the sake of simplicity, only false 2-bit MBUs will be considered. Thus, $\mu \equiv N_{FMBU}(2)$ so Eq. 23 becomes:

$$\begin{aligned} P_{FMBU} &\simeq 1 - \exp(-N_{FMBU}(2)) \simeq \\ &\simeq 1 - \exp\left(-\frac{1}{2} \cdot \frac{W-1}{W} \cdot \frac{m \cdot (m-1)}{L_A}\right). \end{aligned} \quad (24)$$

In order to simplify this expression, only the main element in Eq. 9 has been kept. What does the birthday statistics predict? In his original work, Tausch investigated false MBUs in an FPGA and postulated $k = W/2$ to validate Eq. 21. Using this value here, $2k-1 = W-1$ and, as $L = L_A \cdot W$, Eq. 21 becomes Eq. 24, which was obtained without assumptions concerning the value of k .

A more serious problem arises when dealing with false MCUs. The reason is that, in the birthday statistics, the memory is supposed to be a unidimensional set where every single cell is associated with its address, a natural number. On the contrary, in Section II-C the memory is modeled as a bidimensional grid. However, a link between the urn-and-balls model and the birthday statistics comes up regarding expressions equivalent to Eq. 19 but in the one-dimension case. This is achieved by directly using Eq. 13, appropriate for unidimensional sets, to calculate the probability of two bitflips occurring with a distance *lower* than k . This is:

$$P(0 \leq d < k) = \sum_{d=0}^{k-1} P_X(d) = P_X(0) + \sum_{d=1}^{k-1} P_X(d) \quad (25)$$

Unlike Eq. 16, the range of values goes from $d = 0$ to $d = k - 1$ since the original work by Abramson & Moser calculated the probability of “(...) that the birthdays of every day are at least k days apart (...)” [15]. This was the reason of not having initially identified Tausch’s k with D of Eq. 16. Using Eq. 13 with $L_X \equiv L$, Eq. 25 eventually becomes:

$$P(0 \leq d < k) = \frac{2k-1}{L} \cdot \left(1 - \frac{1}{L} \cdot \frac{k \cdot (k-1)}{2k-1}\right). \quad (26)$$

This expression, multiplied by the total number of possible pairs (Eq. 18) and assuming $k \ll L$, yields:

$$N_{FMCU,1D} = \frac{1}{2} \cdot m \cdot (m-1) \cdot \frac{2k-1}{L}. \quad (27)$$

Finally, the identification of $N_{FMCU,1D}$ with μ in Eq. 23 allows rededucing Eq. 21, the most important result of the birthday statistics, from a different path than that followed by Tausch in his original paper.

IV. VERIFICATION WITH MONTE CARLO TESTS

Results issued from the previous section require to be verified in experiments. However, it is difficult to perform such experimental validations in actual SRAMs since not only do SBUs occur, but also multiple events. The results issued from Section II-B, concerning false MBUs, may be tested with an SRAM with a large interleaving distance and built in old technologies and/or exposed to low energy particles. Thus, it can be guaranteed that the proportion of multiple events is low related to that of SBUs and that, at any rate, no genuine MBU can occur. Therefore, any detected candidate for MBU would be false. On the contrary, it is much more difficult to investigate false MCUs since, unlike the previous case, no multiple event is allowed. This implies using hardened memories under high radiation fluxes for a long time. Both tests however would require significant time and financial effort.

An alternative option consists in making Monte Carlo simulations. A $1M \times 8$ -bit SRAM, with a realistic size and configuration, was emulated as a 2048×4096 matrix with boolean elements initially set to “**0**” (FALSE). As every address has 20 bits, the 12 most significant bits were used to calculate the position along the Y-axis. In actual SRAMs, these bits would be used to feed the quad selector and the row decoder. The other 8 bits were used to select the X-position of the word and, finally, other 3 bits (N_W) were used to select the exact X-position of the cells.

Bitflips were emulated by just choosing two random x, y values with $x \in [1, 2, \dots, 2048]$, $y \in [1, 2, \dots, 4096]$ and then flipping the (x, y) -cell, whichever its value was. In every emulated irradiation, N_{FLIPS} matrix elements, or cells, were flipped and the same experiment was repeated N_{TRIAL} times. Thus, a set of N_{TRIAL} elements, A , was obtained for every interesting parameter (for example, the observed number of bitflips, false MBUs, etc.). The average value of the A set is the expected value and $t_{95} \cdot \sigma / \sqrt{N_{TRIAL}}$ the uncertainty, σ being the standard deviation of A , and t_{95} the Student’s t with a 95% confidence. In practice, this is 1.98 given the very high number of trials. All the calculations were done using the 0.6.1 release of the Julia programming language [16].

TABLE I
OBSERVED BITFLIPS VS. PREDICTIONS IN THE SBU-ONLY MODEL

N_{TRIAL}	N_{FLIPS}	Observed Bitflips	Predicted Bitflips
10^5	10	10	9.999989
	30	$29.99994 \pm 7 \cdot 10^{-5}$	29.99990
	60	59.99970 ± 0.00015	59.99958
	100	99.99864 ± 0.00033	99.99882
	300	299.9892 ± 0.0010	299.9893
	600	599.9581 ± 0.0018	599.9572
	1000	999.8799 ± 0.0031	999.8809
	3000	2998.932 ± 0.009	2998.927
	6000	5995.718 ± 0.018	5995.709
	10000	9988.077 ± 0.031	9988.080

Table I shows the observed number of bitflips vs. the predictions of Eq. 6. In this table, N_{FLIPS} is the total number of bitflips induced in the simulation, “*Observed Bitflips*” is the average number of bitflips observed after N_{TRIAL} experiments and “*Expected Bitflips*” is the prediction issued from Eq. 5 with $m = N_{FLIPS}$. Next, Table II compares the number of false MBUs with predictions (Eqs. 9 - 10). The probability of collision of two SBUs according to the birthday statistics model (Eq. 21 with $k = W/2$) is also included for comparison. In both cases, the predictions from the SBU model are extremely accurate. Only in Table II there is a small underestimation of the number of 2-bit MBUs. It is possible that this phenomenon is related to border effects since, in the theoretical model, it was not taken into account the XY distribution of cells.

Figs. 5a-b compare the apparent number of related pairs of cells depending on the number of bitflips and the chosen threshold distance. In both cases, results accurately fit the theoretical predictions.

TABLE II
FALSE MBUS VS. PREDICTIONS IN THE SBU-ONLY MODEL.

N_{TRIAL}	N_{FLIPS}	$P_{Birth. Stat.}$	2-bit MBUs		3-bit MBUs	
			Observed	Predicted	Observed	Predicted
100,000	10	$3.76 \cdot 10^{-5}$	$(4 \pm 4) \cdot 10^{-5}$	$3.76 \cdot 10^{-5}$	0	$7.2 \cdot 10^{-11}$
	30	$3.63 \cdot 10^{-4}$	$(3.8 \pm 1.2) \cdot 10^{-4}$	$3.63 \cdot 10^{-4}$	0	$2.4 \cdot 10^{-9}$
	60	$1.48 \cdot 10^{-3}$	$(1.4 \pm 0.2) \cdot 10^{-3}$	$1.48 \cdot 10^{-3}$	0	$2.0 \cdot 10^{-8}$
	100	$4.12 \cdot 10^{-3}$	$(4.4 \pm 0.4) \cdot 10^{-3}$	$4.13 \cdot 10^{-3}$	0	$9.7 \cdot 10^{-8}$
	300	$3.67 \cdot 10^{-2}$	$(3.77 \pm 0.12) \cdot 10^{-2}$	$3.74 \cdot 10^{-2}$	0	$2.7 \cdot 10^{-7}$
	600	0.139	0.151 ± 0.002	0.150	$(2 \pm 3) \cdot 10^{-5}$	$2.1 \cdot 10^{-5}$
	1,000	0.341	0.415 ± 0.004	0.4164	$(12 \pm 7) \cdot 10^{-5}$	$9.9 \cdot 10^{-5}$
	3,000	0.976	3.745 ± 0.012	3.743	0.0026 ± 0.0003	0.0027
	6,000	~ 1	14.963 ± 0.024	14.932	0.0215 ± 0.0009	0.0214
25,000	10,000	~ 1	41.47 ± 0.08	41.321	0.098 ± 0.004	0.098

V. DISCUSSION

Even though the model developed in Section II and verified in Section IV is quite restrictive since it only considers SBUs and discards other phenomena, its predictions can be useful because, in many situations, experiments end without microlatchups or Single Event Functional Interruptions (SEFI), so the detected bitflips are SBUs or multiple events. Even more, SBUs typically constitute the bulk of bitflips so the predictions of the model can be extrapolated without making a significant mistake. This approach is also taken in other accepted strategies such as the birthday statistics [12].

These predictions are useful to correctly interpret experimental data issued from radiation tests. First of all, Table I demonstrates that, in practice, it is not necessary to make a distinction between the actual and the observed number of bitflips (m vs. n in Eqs. 5-6). By using the properties of Poisson distribution [17], even in the cases with the highest amount of bitflips (10000 in Table I), the intrinsic statistical deviation is 1.96% with 95%-confidence whereas the correction introduced by Eq. 6 is on the order of 0.1%, twenty times smaller.

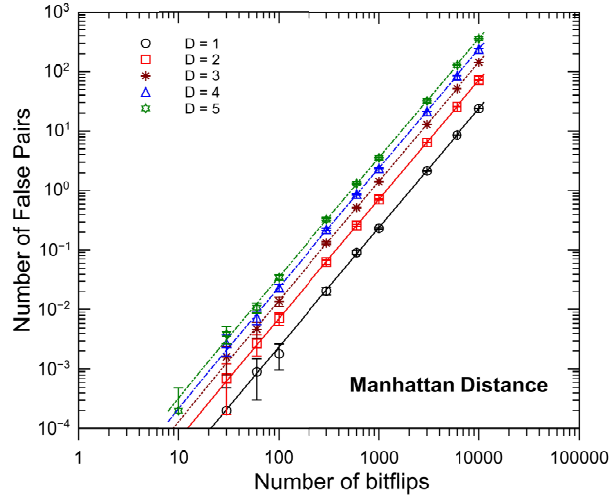
On the contrary, false MBUs are more likely to induce a misinterpretation of the experimental results. First of all, unlike Eq. 3, Eqs. 8-10 contain L_A , which is from 8 to 16 times smaller than L . Therefore, the likelihood of occurrence is higher. Second, candidates for MBUs are quite easy to detect during the experiments and draw the researchers' attention. In fact, if the internal structure of the SRAM is unknown, this is the only kind of multiple events that could be observed without carrying out a more complex analysis.

False MBUs have been detected by the authors in previous experiments. For example, in the data published in [18], 1925 bitflips were observed after irradiating a $2M \times 8$ -bit SRAM. According to Eq. 9, in that situation, the expected number of MBUs is not negligible (0.772) and, in fact, one apparent 2-bit MBU was observed in that experiment. It was eventually discarded since confidential unscrambling information from the manufacturers was available.

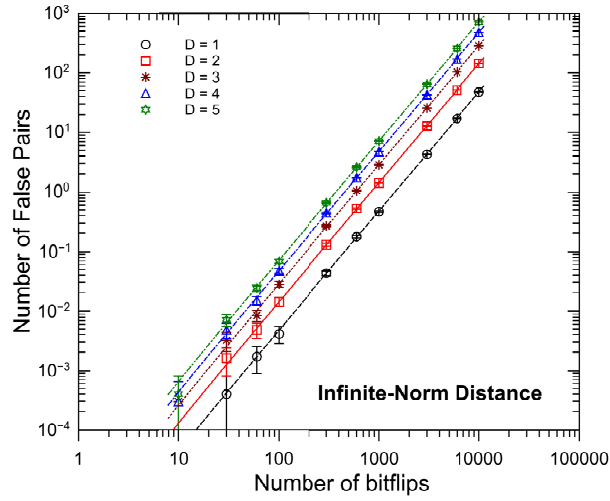
Strictly, the correct manner of using the information provided in previous sections is to calculate the mean value and the uncertainty of the experimental number of MBUs by using the χ^2 function as proposed in [17]. Later, the theoretical number of false MBUs must be subtracted and only accept the existence of MBUs beyond the shadow of a doubt if the range of possible values is strictly positive. For example, if one MBU candidate is observed, the actual value of the experimental data is between $[0.025, 5.572]$ with 95%-confidence. The subtraction of the expected number of false MBUs transforms this range into $[-0.767, 4.800]$. Therefore, it should not be concluded that MBUs have occurred during the experiment.

Finally, the most important fact is that false MCUs are not uncommon in radiation experiments. Firstly, it is important to remark that Eq. 19 does not strictly predict the number of 2-bit MCUs, but the number of unrelated neighbor pairs of bitflips. For example, two unrelated cells could be misled with a 3-bit MCU if one of them belongs to a 2-bit MCU. However, this equation can be used either to determine how trustworthy the results from a irradiation tests are; or to make a coarse correction of the number of 2-bit MCUs. Thus, if Eq. 19 casts a value comparable to the number of detected 2-bit MCUs, one can determine how correctly the experimental data have been analyzed. In the experiment depicted in [18], where 1925 bitflips were observed in a 16-Mbit SRAM, the use of proprietary information allowed classifying the events in 1645 SBUs, 96 2-bit MCUs, and other events of higher multiplicity. The criterium to relate two cells to the same event was that the Manhattan distance between them being lower than 5 ($D = 4$). According to Eq. 19 and using $m = 1925$ and $L = 2^{24}$, the expected number of apparent pairs is 4.41. Therefore, about 5% of the observed 2-bit MCUs were potentially false.

It is also interesting to investigate what happens when the internal structure of the SRAM is unknown, so the number of multiple events has to be estimated by using statistical deviations [3]–[6]. The MBU number can be corrected with Eqs. 9-10 since they do not require the knowledge of the internal memory structure. More difficult is to estimate the number of MCUs since Eqs. 17-20 require a threshold distance, which is a parameter that does not appear in the statistical approach. Nevertheless, it is possible to perform similar studies to the one depicted in Section II if these procedures are applied. However, the results strongly depend on how the statistical anomalies are detected and no general relations can be provided.



(a)



(b)

Fig. 5. Number of pairs of bit cells at a distance equal or smaller than D in proportion to the number of independent random bitflips: Manhattan Distance (a), infinite-norm distance (b). Lines are the theoretical predictions issued from Eqs. 17, 18 & 20 whereas dots and error bars represents results from Monte Carlo simulations.

VI. CONCLUSION

In this paper, the statistical properties of an SBU-only scenario have been investigated to determine the way that results issued from actual experiments on SRAMs can be erroneously interpreted. The model predicts that the number of observed events is close enough to the number of bitflips that actually occurred and both concepts are almost equivalent. Besides, it was determined that false 2-bit MBUs can be observed in realistic experiments and, finally, if the internal structure of the SRAM is known and bitflips in adjacent cells are supposed to belong to a unique multiple event, a non-negligible fraction of the postulated pairs of bitflips detected this way are, in fact, independent. In any case, some simple mathematical expressions have been proposed to estimate the correction that can be applied to provide more accurate values of soft-error cross-sections.

APPENDIX

Mathematical expressions used throughout the manuscript:

$$\sum_{k=1}^n a = a \cdot n \quad (28)$$

$$\sum_{k=1}^n k = \frac{1}{2} \cdot n \cdot (n + 1) \quad (29)$$

$$\sum_{k=1}^n k^2 = \frac{1}{6} \cdot n \cdot (n+1) \cdot (2n+1) \quad (30)$$

$$\sum_{k=1}^n k^3 = \frac{1}{4} \cdot n^2 \cdot (n+1)^2 \quad (31)$$

$$(1+x)^m \approx 1 + m \cdot x + \frac{1}{2} \cdot m \cdot (m-1) \cdot x^2 + \dots \quad (32)$$

$$\sqrt{1+x} \approx 1 + \frac{1}{2} \cdot x - \frac{1}{8} \cdot x^2 + \dots \quad (33)$$

REFERENCES

- [1] D. Binder, E. C. Smith, and A. B. Holman, "Satellite Anomalies from Galactic Cosmic Rays," *IEEE Transactions on Nuclear Science*, vol. 22, pp. 2675–2680, Dec. 1975.
- [2] M. P. King and R. A. Reed and R. A. Weller and M. H. Mendenhall and R. D. Schrimpf and B. D. Sierawski and A. L. Sternberg and B. Narasimham and J. K. Wang and E. Pitta and B. Bartz and D. Reed and C. Monzel and R. C. Baumann and X. Deng and J. A. Pellish and M. D. Berg and C. M. Seidleck and E. C. Auden and S. L. Weeden-Wright and N. J. Gaspard and C. X. Zhang and D. M. Fleetwood, "Electron-Induced Single-Event Upsets in Static Random Access Memory," *IEEE Transactions on Nuclear Science*, vol. 60, no. 6, pp. 4122–4129, 2013.
- [3] M. J. Gadlage, A. H. Roach, A. R. Duncan, A. M. Williams, D. P. Bossev, and M. J. Kay, "Multiple-Cell Upsets Induced by Single High-Energy Electrons," *IEEE Transactions on Nuclear Science*, vol. 65, pp. 211–216, Jan. 2018.
- [4] M. Wirthlin, D. Lee, G. Swift, and H. Quinn, "A Method and Case Study on Identifying Physically Adjacent Multiple-Cell Upsets Using 28-nm, Interleaved and SECEDED-Protected Arrays," *IEEE Transactions on Nuclear Science*, vol. 61, pp. 3080–3087, Dec. 2014.
- [5] J. A. Clemente, F. J. Franco, F. Villa, M. Baylac, S. Rey, H. Mecha, J. A. Agapito, H. Puchner, G. Hubert, and R. Velazco, "Statistical Anomalies of Bitflips in SRAMs to Discriminate SBUs From MCUs," *IEEE Transactions on Nuclear Science*, vol. 63, pp. 2087–2094, Aug. 2016.
- [6] F. J. Franco, J. A. Clemente, M. Baylac, S. Rey, F. Villa, H. Mecha, J. A. Agapito, H. Puchner, G. Hubert, and R. Velazco, "Statistical Deviations From the Theoretical Only-SBU Model to Estimate MCU Rates in SRAMs," *IEEE Transactions on Nuclear Science*, vol. 64, pp. 2152–2160, Aug. 2017.
- [7] A. Bossler, V. Gupta, G. Tsiligiannis, A. Javanainen, H. Kettunen, H. Puchner, F. Saigné, A. Virtanen, F. Wrobel, and L. Dilillo, "Investigation on MCU Clustering Methodologies for Cross-Section Estimation of RAMs," *IEEE Transactions on Nuclear Science*, vol. 62, pp. 2620–2626, Dec. 2015.
- [8] A. L. Bossler, V. Gupta, G. Tsiligiannis, C. D. Frost, A. Zadeh, J. Jaatinen, A. Javanainen, H. Puchner, F. Saigné, A. Virtanen, F. Wrobel, and L. Dilillo, "Methodologies for the Statistical Analysis of Memory Response to Radiation," *IEEE Transactions on Nuclear Science*, vol. 63, pp. 2122–2128, Aug. 2016.
- [9] A. Manuzato, S. Gerardin, A. Paccagnella, L. Sterpone, and M. Violante, "On the Static Cross Section of SRAM-Based FPGAs," in *2008 IEEE Radiation Effects Data Workshop*, pp. 94–97, Jul. 2008.
- [10] J. Tonfat, F. L. Kastensmidt, L. Artola, G. Hubert, N. H. Medina, N. Added, V. A. P. Aguiar, F. Aguirre, E. L. A. Macchione, and M. A. G. Silveira, "Analyzing the Influence of the Angles of Incidence and Rotation on MBU Events Induced by Low LET Heavy Ions in a 28-nm SRAM-Based FPGA," *IEEE Transactions on Nuclear Science*, vol. 64, pp. 2161–2168, Aug. 2017.
- [11] J. A. Maestro and P. Reviriego, "Study of the Effects of MBUs on the Reliability of a 150 nm SRAM Device," in *Proceedings of the 45th Annual Design Automation Conference, DAC '08*, (New York, NY, USA), pp. 930–935, ACM, 2008.
- [12] H. J. Tausch, "Simplified Birthday Statistics and Hamming EDAC," *IEEE Transactions on Nuclear Science*, vol. 56, pp. 474–478, Apr. 2009.
- [13] P. Reviriego, J. A. Maestro, and C. J. Bleakley, "Reliability Analysis of Memories Protected with BICS and a Per-word Parity Bit," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 15, no. 2, pp. 18:1–18:15, 2010.
- [14] G. Schay, *Introduction to Probability with Statistical Applications*. Springer, 2nd ed., Dic. 2016.
- [15] M. Abramson and W. O. J. Moser, "More Birthday Surprises," *The American Mathematical Monthly*, vol. 77, pp. 856–858, Oct. 1970.
- [16] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A Fresh Approach to Numerical Computing," *SIAM Review*, vol. 59, pp. 65–98, Jan. 2017.
- [17] J. L. Autran, D. Munteanu, P. Roche, and G. Gasiot, "Real-time soft-error rate measurements: A review," *Microelectronics Reliability*, vol. 54, pp. 1455–1476, Aug. 2014.
- [18] J. A. Clemente, G. Hubert, F. J. Franco, F. Villa, M. Baylac, H. Mecha, H. Puchner, and R. Velazco, "Sensitivity Characterization of a COTS 90-nm SRAM at Ultralow Bias Voltage," *IEEE Transactions on Nuclear Science*, vol. 64, pp. 2188–2195, Aug. 2017.