

Protection of Memories Suffering MCUs Through the Selection of the Optimal Interleaving Distance

Pedro Reviriego, Juan Antonio Maestro, Sanghyeon Baeg, ShiJie Wen, and Richard Wong

Abstract—Interleaving, together with single error correction codes (SEC), are common techniques to protect memories against multiple cell upsets (MCUs). This kind of errors is increasingly important as technology scales, becoming a prominent effect, and therefore greatly affecting the reliability of memories. Ideally, the interleaving distance (ID) should be chosen as the maximum expected MCU size. In this way, all errors in an MCU would occur in different logical words, thus being correctable by the SEC codes. However, the use of large interleaving distances usually results in an area increase and a more complex design of memories. In this paper, the selection of the optimal interleaving distance is explored, keeping the area overhead and complexity as low as possible, without compromising memory reliability.

Index Terms—Interleaving distance, memory, multiple cell upsets (MCUs), soft error.

I. INTRODUCTION

COMPUTER memories are sensitive to soft errors which can affect system reliability. Memory cells can be disturbed by high-energy neutron particles from terrestrial atmosphere or alpha particles resulted from IC package material. Previous studies showed that the soft error rate is closely related to critical charge [1], [2] and process [3], [4]. Therefore, a natural way to mitigate soft error issues is to increase critical charge at state nodes, or to use process-related immunity techniques such as well and substrate engineering. Another option is to include error correction capabilities on memories so that some of the errors can be corrected. This is normally done by using a single error correction (SEC) code on each memory word to deal with isolated errors [5]. Scrubbing can be combined with single error correction to further increase reliability by periodically reading the memory and correcting the single errors, so that they do not accumulate over time [6]. The combination of SEC and scrubbing is effective against single event upsets but not against multiple cell upsets (MCUs), as the errors in this kind of event tend

to be physically close, and therefore it is likely that they affect more than one bit of the same memory word [7]–[12]. To deal with MCUs, interleaving is commonly used [5]. Interleaving ensures that cells that belong to the same word are physically apart so that only one can be affected by errors in the same MCU. This is illustrated in Fig. 1 where an ID of eight is used, so that an MCU should affect columns at a distance larger than eight to upset two bits of the same word. The words are selected by a combination of row and column and only three bits are shown.

It is commonly assumed that the ID is selected large enough to assure that no MCU causes errors on two or more bits of the same logical word. Based on that assumption, reliability models to calculate the failure probability versus time [13] and the mean time to failure (MTTF) [5] of memories have been proposed.

However, the use of large IDs can imply a more complex and costly memory design [13]. Therefore, if the reliability targets can be met with a smaller ID, it would be more effective to use that smaller ID value. There is no systematic methodology to determine the optimal ID for a memory configuration, thus usually leading to higher values that overprotect the system.

In this paper, an analysis of the impact of the ID on the memory reliability is presented. The goal is to quantify the effect of reducing that ID, helping the designer choose the optimal value and to bound the probability of error due to MCUs.

This paper is organized as follows. Section II presents the reliability analysis of the memories for failures caused by MCUs exceeding the ID and failures caused by the rest of the error events. The results are then used to discuss the ID selection procedure in Section III where a case study is used to illustrate the ID selection process. Finally Section IV concludes the paper.

II. RELIABILITY ANALYSIS

In this section, the reliability of a memory is studied considering two types of failures.

- Direct failures caused by an MCU exceeding the ID. If this happens, it is possible that two or more errors hit the same logical word,¹ therefore producing a failure.
- Accumulation failures caused by two independent events, producing two or more errors on the same word. This second type of failures is independent of the ID, and has been previously studied in [5], [13], [14].

Assume that the interleaving scheme shown in Fig. 1 is used and that, as explained before, MCUs that exceed the interleaving distance always cause a failure. Define $e(n)$ as the probability that

¹Not all the MCUs exceeding the ID will cause a failure, because the MCU pattern (physical distribution) also affects. In this paper, this effect will be disregarded, and therefore all such MCUs are modeled as causing failures.

Manuscript received September 07, 2009; revised December 17, 2009, January 19, 2010; accepted February 02, 2010. Date of current version August 18, 2010. This work was supported in part by the Spanish Ministry of Science and Innovation under Grant AYA2009-13300-C03-01, by the Regional Government of Madrid, and in part the European Union FEDER programme.

P. Reviriego and J. A. Maestro are with the Universidad Antonio de Nebrija, C/Pirineos, 55 E-28040 Madrid, Spain (e-mail: previrie@nebrija.es; jmaestro@nebrija.es).

S. Baeg is with the School of Electrical Engineering and Computer Science, Hanyang University, Kyung-Gi-Do, Korea (e-mail: bau@hanyang.ac.kr).

S.J. Wen and R. Wong are with the Component Engineering Group, Cisco Systems Inc., San Jose, CA 95134 USA (e-mail: shwen@cisco.com; rickwon@cisco.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNS.2010.2042818




Fig. 1. Example of interleaving in a memory.

a given error event spans n columns and e_{ID} as the probability that an error event causes a direct failure when the interleaving distance is ID . Then, for a given ID value, the probability that an error event causes a direct failure is given by how likely MCUs span more than ID

$$e_{ID} = \sum_{n=ID+1}^{\infty} e(n). \quad (1)$$

This adds the probabilities of all MCUs spanning more than ID columns.

Define $p(n)$ as the probability that an event causes n cell errors (therefore being $p(1)$ the probability that a given event is an SEU, and $p(n)$, for $n > 1$, the probability of an n -bit MCU). If we denote by α the average number of errors per event, then α can be computed as follows:

$$\alpha = \sum_{n=1}^{\infty} n \cdot p(n). \quad (2)$$

Under these assumptions, the probability of failure due to the two mentioned mechanisms can be studied: direct failure when an event provokes errors that exceed the ID , and accumulation failure caused by two independent events causing errors on the same word.

To study the memory reliability, the mean time to failure (MTTF) will be used as a figure of merit. For direct failures, the MTTF is given by

$$\text{MTTF}|_d = \frac{1}{\lambda_{\text{word}} \cdot M \cdot e_{ID}}. \quad (3)$$

where λ is the per-word error event arrival rate and M is the memory size in words. This is a direct conclusion if we consider that events arrive following a Poisson distribution, since $\text{MTTF} = \text{METF}/\lambda$, and $\text{METF} = 1/e_{ID}$ (METF being the mean number of events to failure) [15].

For accumulation failures, the MTTF can be approximated when M is large by

$$\text{MTTF}|_a \cong \frac{1}{\lambda_{\text{word}} \cdot \alpha} \cdot \sqrt{\frac{\pi}{2 \cdot M}}. \quad (4)$$

Which is an extension of the traditional MTTF approximation for SEUs only [14]. The proof of this equation can be found in

[5], where the scenario in which MCUs accumulate in memories is modeled.

The total MTTF of the memory will be determined by both effects. This is equivalent to the traditional model of two elements connected in series such that the system fails when one of them fails [16]. For those systems, when the probability of failure is uniformly distributed with time, the total MTTF can be expressed as a function of the partial MTTFs as

$$\text{MTTF} = \frac{1}{\frac{1}{\text{MTTF}_1} + \frac{1}{\text{MTTF}_2}}. \quad (5)$$

In the memory case, the direct failures have a uniformly distributed probability of failure with time (all the direct failures have the same probability of occurrence), but the accumulation failures do not. This is due to the fact that as errors accumulate, a new error is more likely to affect a word that already contains a previous error causing a failure (see for example [17] for more details). Therefore, in our case, (5) is only an approximation for the MTTF of the memory

$$\text{MTTF}_{\text{memory}} \cong \frac{1}{\frac{1}{\text{MTTF}_d} + \frac{1}{\text{MTTF}_a}}. \quad (6)$$

This approximation will be used in the following section to assess the impact of the ID selection on the MTTF. Note that the ID affects the probability of direct failure, e_{ID} , per expression (1), and this probability is related to the MTTF per expressions (3) and (6).

III. SELECTION OF THE INTERLEAVING DISTANCE

In this section, and based on the previous analysis, the selection process of the ID is presented now using a real case study. Four different memory technologies have been studied which have been previously characterized with real radiation experiments. They correspond to advanced geometries (65 nm and 45 nm) for which MCUs are a major concern and large ID s are normally used to ensure that no direct failures occur. The memories were exposed to white beams (which contain neutrons with a broad energy spectrum) up to 800 MeV at the LANCE site and neutron beams up to 180 MeV at the TSL site. The neutron flux that has been used is 10^6 neutron/cm² · second. These white beams have been chosen to replicate the neutron spectrum that exists at ground level. Multiple devices were used and for

TABLE I
AVERAGE NUMBER OF ERRORS PER EVENT

α	65 nm		45 nm	
	A	B	A	B
	2.0649	2.0649	1.9062	1.7573




Fig. 2. Values of e_{ID} obtained in the experiments.

each one multiple tests were performed. For all tests, the mean time between upsets was much larger than the mean time of an SRAM read cycle for the entire memory. Such configuration was achieved by adjustment of the flux intensity. Once an error was detected, a checking procedure was launched to check the error types. More details on the experiments are given on [13].

The purpose of the characterization process has been to determine the two parameters described in the previous section: $e(n)$ (probability that an MCU spans n columns) and $p(n)$ (probability of an n -error event). Note that the conducted experiments have focused on neutron upsets, typical of terrestrial applications. Therefore, the obtained ID would only be valid for this environment.

Once these parameters have been measured, the value of α (average number of errors per event) has been calculated through (2) using $p(n)$. The results for the different memories are shown in Table I. The values for the two types of 65 nm memories were added together so that a single value is shown. This single value is used to illustrate how the proposed technique works. In a real scenario, it might be convenient to use independent values for each memory type to ensure an optimal reliability level. Nevertheless, this decision does not affect the technique procedure explained in the following.

It should also be noted that the experiences in 65 nm were used to enhance the soft error immunity in the 45 nm process. This explains why the average number of errors per event is smaller in the more advanced technology.

The values of $e(n)$ for the four types of memories were used to compute the probability of an event causing a direct failure for each ID value (e_{ID}), using (1). The results are shown in Fig. 2, where it can be observed that these values decrease as the ID increases. This is obvious, as the probability of a direct failure lowers with higher values of IDs. But, those high IDs, although safer, introduce an unnecessary complexity in the memory.

Therefore, the objective is to find the minimal ID that produces a reasonable MTTF in the memory.

With the previously described parameters, expressions (3) and (4) can be used to estimate the MTTF for direct and accumulation failures. Then using (6), the MTTF of the memory can be approximated. From the memory designer perspective, the main concern is choosing the minimal ID, but with a negligible impact of direct failures on the MTTF.

This negligible impact would imply that $e_{ID} \rightarrow 0$ (no direct failures due to MCUs). Therefore, according to (2), $MTTF|_d \rightarrow \infty$, that would lead to $MTTF|_{memory} \rightarrow MTTF|_a$, per expression (6). The mean time to failure of the memory would only be affected by the accumulation of several independent events. Therefore, the closer the ratio $MTTF|_{ratio} = MTTF|_{memory}/MTTF|_a$ is to 1, the less impact of direct failures. As this ratio decreases from 1, that would represent a decrement of the MTTF due to those direct failures. For example, given a value of ID, a ratio of 0.8 would mean that the reliability of the memory is 80% of its optimal value due to the direct failures caused by large MCUs that cannot be handled by the interleaving. In this case, a higher ID would be advisable (what would lead to a higher MTTF ratio).





In this way, the effect of the interleaving distance is quantified, helping the designer with the selection of an optimal value.

Considering the case under study, the value of the ratio has been computed using expression (6) for the different geometries, implementing various IDs and memory sizes. The results are presented in Figs. 3 to 6. Analyzing the plots the following observations can be made. First, as the memory size increases, larger ID values are needed to ensure a small impact of direct failures (high MTTF ratios). This can be explained as for larger memories, more errors are needed to cause a failure by error accumulation and therefore even a small percentage of errors causing direct failures will affect the reliability. The conclusion is that the optimal ID tends to grow with the memory size. The second observation is that the four memories follow a similar trend, and therefore similar IDs would produce a similar impact on all of them.

In order to make a more detailed analysis of the ID versus reliability tradeoff, let us now focus on the design of a 256-Kword memory. Traditionally, the ID would be selected as the maximum MCU size expected in the system, in order to avoid direct failures. In the experiments, for 65 nm, this maximum size has been detected to be 11 (for type A) and 12 (for type B). For 45 nm, the maximum size is 11 (for both types A and B). These values are represented in Table II as $ID_{conservative}$. These results have been rounded up to the closest power of 2 (in parenthesis in Table II), since these are the values usually chosen when designing a memory.

Theoretically, with this ID, the effect of direct failures would have a negligible impact on the MTTF on the memories. This can be seen in Figs. 3 to 6, where the MTTF ratio is 1 for these values.

However, let us now consider a reliability goal such that direct failures may have an effect of up to 10% on the MTTF. Obviously, the previous ID values would meet this requirement, but they could be reduced in order to decrease the interleaving complexity as much as possible. In order to find out the minimal ID


 Fig. 3. MTTF_{ratio} for the 65 nmA memories in different configurations.

 Fig. 5. MTTF_{ratio} for the 45 nmA memories in different configurations.

 Fig. 4. MTTF_{ratio} for the 65 nmB memories in different configurations.

 Fig. 6. MTTF_{ratio} for the 45 nmB memories in different configurations.

that is able to meet the reliability goal, the plots in Figs. 3 to 6 have to be used again. The results are shown in Table II as ID_{\min} , together with the closest power of 2. These ID values will reduce the area and power of the memory making the design more competitive.

To illustrate the benefits of the proposed approach, the costs of a memory with the different ID values have been compared. The cost calculation is based on data from [13], summarized in Table III, where the relative area/power overhead versus an ID of four are shown. The table shows that both the area and power increase significantly with the ID.

In Table IV, the relative area overheads have been depicted for the ID values (powers of 2) determined in Table II, as well as the percentages of area savings that are achieved by using the optimal ID versus the conservative one. The area estimates for each ID value are based on the data in Table III. The results for the power consumption overhead are shown in Table V.

The results show that the area and power can be significantly reduced in this case for three of the memory types (65 nmA,

 TABLE II
 MINIMUM ID VALUES

	65nmA	65nmB	45nmA	45nmB
$ID_{\text{conservative}}$	11→(16)	12→(16)	11→(16)	11→(16)
ID_{\min}	8→(8)	8→(8)	7→(8)	11→(16)

 TABLE III
 AREA OVERHEAD FOR DIFFERENT ID VALUES

ID	Area increment	Power Increment
4	1	1
8	1,03	1,07
16	1,32	1,25
32	2,03	2,02

65 nmB, 45 nmA) with a negligible impact on reliability, using ID_{\min} as described before, versus the conservative ID values (which will be the natural choice if this methodology is not applied). Therefore, the proposed ID selection process achieves the goal of choosing the ID that minimizes the cost without impacting reliability.

TABLE IV
AREA INCREMENT FOR THE TWO ID CONFIGURATIONS

	65nmA	65nmB	45nmA	45nmB
ID _{conservative}	1.32	1.32	1.32	1.32
ID _{min}	1.03	1.03	1.03	1.32
% savings	22%	22%	22%	0%

TABLE V
POWER INCREMENT FOR THE TWO ID CONFIGURATIONS

	65nmA	65nmB	45nmA	45nmB
ID _{conservative}	1.25	1.25	1.25	1.25
ID _{min}	1.07	1.07	1.07	1.25
% savings	14%	14%	14%	0%

As a final comment, it is worth to mention that the applicability of the presented technique relies on the knowledge of several technology parameters ($e(n)$, $p(n)$, α). This implies that in order to get these values, initial radiation tests should be conducted on new memories in order to characterize them. This somewhat limits the applicability of the method in earlier designs but, as explained in this section, produces clear benefits once the technology is well characterized.

IV. CONCLUSION

In this paper, the reliability of memories that use SEC and interleaving has been analyzed. A procedure to ensure that failures caused by MCUs exceeding the ID have a negligible impact on reliability has been presented. The procedure helps memory designers choose the minimal ID (thus reducing area and complexity), but assuring an appropriate reliability level. A case study has also been presented showing the potential benefits of the proposed approach using real radiation data. The results show that significant area and power savings can be obtained in some cases.

Another interesting observation from the analysis is that larger memories are more likely to need larger ID values, as they tolerate less percentage of MCUs exceeding the ID. As technology shrinks, MCUs tend to affect more cells and memories tend to be larger. Those two factors will reinforce the need for larger ID in future memory designs. This in turn will result in a larger area and power overhead due to the ID.

REFERENCES

- [1] E. Ibe, S. Chung, S. Wen, H. Yamaguchi, Y. Yahagi, H. Kameyama, S. Yamamoto, and T. Akioka, "Spreading diversity in multi-cell neutron-induced upsets with device scaling," in *Proc. IEEE Custom Integrated Circuit Conf.*, 2006, pp. 437–444.
- [2] Y. Kawakami, M. Hane, H. Nakamura, T. Yamada, and K. Kumagai, "Investigation of soft error rate including multi-bit upsets in advanced SRAM using neutron irradiation test and 3D mixed-mode device simulation," in *Proc. IEDM*, Dec. 2004, pp. 945–948.
- [3] D. Radaelli, H. Puchner, S. Wong, and S. Daniel, "Investigation of multi-bit upsets in a 150 nm technology SRAM device," *IEEE Trans. Nucl. Sci.*, vol. 52, no. 6, pp. 2433–2437, Dec. 2005.
- [4] Y. Tosaka, H. Ehara, M. Igeta, T. Uemura, H. Oka, N. Matsuoka, and K. Hatanaka, "Comprehensive study of soft errors in advanced CMOS circuits with 90/130 nm technology," in *Proc. IEDM*, Dec. 2004, pp. 941–944.
- [5] P. Reviriego, J. A. Maestro, and C. Cervantes, "Reliability analysis of memories suffering multiple bit upsets," *IEEE Trans. Device Mater. Rel.*, vol. 7, no. 4, pp. 592–601, Dec. 2007.
- [6] G. C. Yang, "Reliability of semiconductor RAMs with soft-error scrubbing techniques," *IEE Proc. Computers and Digital Techniques*, vol. 142, no. 5, pp. 337–344, Sep. 1995.
- [7] S. Satoh and Y. Tosaka, "Geometric effect on multiple-bit soft errors induced by cosmic ray neutrons on DRAM's," *IEEE Electron Dev. Lett.*, vol. 21, no. 6, pp. 310–312, Jun. 2000.
- [8] A. D. Tipton, J. A. Pellig, R. A. Reed, R. D. Schrimpf, R. A. Weller, M. H. Mendenhall, B. Sierawski, A. K. Sutton, R. M. Diestelhorst, G. Espinel, J. D. Cressler, P. W. Marshall, and G. Vizkelethy, "Multiple-bit upset in 130 nm CMOS technology," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 6, pt. 1, pp. 3259–3264, Dec. 2006.
- [9] A. M. Chugg, M. J. Moutrie, A. J. Burnell, and R. Jones, "A statistical technique to measure the proportion of MBU's in SEE testing," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 6, pt. 1, pp. 3139–3144, Dec. 2006.
- [10] J. Maiz, S. Hareland, K. Zhang, and P. Armstrong, "Characterization of multi-bit soft error events in advanced SRAMs," in *IEEE Int. Electron Devices Meeting IEDM'03 Digest*, Dec. 2003, pp. 21.4.1–21.4.4.
- [11] D. Giot, P. Roche, G. Gasiot, J. L. Autran, and R. Harboe-Sorensen, "Heavy ion testing and 3-D simulations of multiple cell upset in 65 nm standard SRAMs," *IEEE Trans. Nucl. Sci.*, vol. 55, no. 4, pt. 1, pp. 2048–2054, Aug. 2008.
- [12] R. Lawrence and A. Kelly, "Single event effect induced multiple-cell upsets in a commercial 90 nm CMOS digital technology," *IEEE Trans. Nucl. Sci.*, vol. 55, no. 6, pp. 3367–3374, Dec. 2008.
- [13] S. Baeg, S. Wen, and R. Wong, "SRAM interleaving distance selection with a soft error failure model," *IEEE Trans. Nucl. Sci.*, vol. 56, no. 4, pt. 2, pp. 2111–2118, Aug. 2009.
- [14] A. M. Saleh, J. J. Serrano, and J. H. Patel, "Reliability of scrubbing recovery-techniques for memory systems," *IEEE Trans. Rel.*, vol. 39, no. 1, pp. 114–122, Apr. 1990.
- [15] W. Feller, *An Introduction to Probability Theory and Its Applications*. New York: Wiley, 1968.
- [16] I. Koren and C. M. Krishna, *Fault-Tolerant Systems*. New York: Morgan Kaufmann, 2007.
- [17] S. Baeg, P. Reviriego, J. A. Maestro, S. Wen, and R. Wong, "Analysis of a multiple cell upset failure model for memories," in *IEEE Workshop on Silicon Errors in Logic—System Effects (SELSE'09)*, Stanford, CA, Mar. 2009.