Selection of the Optimal Interleaving Distance for Memories Suffering MCUs

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

Abstract— **As technology shrinks, Multiple Cell Upsets (MCU) are becoming a more prominent effect with a large impact on memory reliability. To protect memories from MCUs, single error correction codes (SEC) and interleaving are commonly used. The interleaving distance (ID) is selected such that all errors in an MCU occur on different logical words. This is achieved by using interleaving distances that are larger than the largest expected MCU. However, the use of a large interleaving distance usually results in an area increase and a more complex design. In this paper, the selection of the optimal interleaving distance is explored, minimizing area and complexity without compromising memory reliability.**

Index Terms— **Interleaving distance, Soft error, MCU, Memory**

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 the memory 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 MCUs as the errors in an MCU tend to be physically close and therefore it is likely that they affect more than one bit of the same memory word [7-10]. 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 Figure 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 such 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 [11] and the MTTF [5] of memories have been proposed.

However, the use of large IDs can imply a more complex and costly memory design [11]. Therefore, if the reliability targets can be met with a smaller ID, it would be more effective to use that smaller ID value. The problem is that, to the best of the authors' knowledge, 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 will be 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.

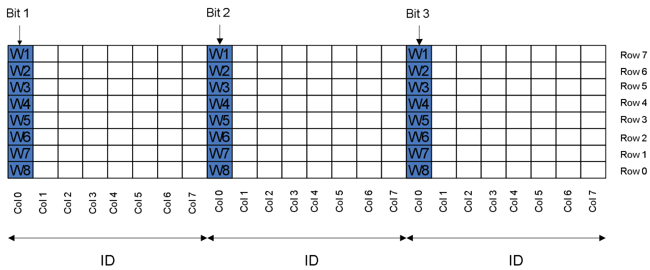

Fig. 1. Example of Interleaving in a memory.

The rest of this paper is organized as follows. Section II presents the reliability analysis of the memories for failures caused by MCUs exceeding the ID and also by failures caused by the rest of the error events. The results are then used to discuss the ID selection procedure in Section III in which 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

Pedro Reviriego and Juan Antonio Maestro are with Universidad Antonio de Nebrija, C/ Pirineos, 55 E-28040 - Madrid, Spain (phone: +34 914521100; fax: +34 914521110; email: {previrie, jmaestro}@nebrija.es).
Sanghyeon Baeg is with the School of Electrical Engineering and Computer Science, Hanyang University, 1271 Sa1-Dong Sangrok-Gu Ansan, Kyung-Gi-Do, Korea (phone/fax: +82-31-400-5237; e-mail: bau@hanyang.ac.kr).
ShiJie Wen, and Richard Wong are with the Component Engineering Group at Cisco Systems Inc. 170 W. Tasman Dr. San Jose CA 95134, U.S.A. (e-mail: shwen@cisco.com, rickwon@cisco.com).

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],[11].

Let us assume that the interleaving scheme shown in Figure 1 is used and that, as explained before, MCUs that exceed the interleaving distance always cause a failure. Let us define $e(n)$ as the probability that 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 basically adds the probabilities of all MCUs spanning more than ID columns.

Let us also 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

$$MTTF\big|_d = \frac{1}{\lambda \cdot M \cdot e^{ID}}. \quad (3)$$

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 MTTF = METF / λ, and METF = 1 / e^{ID} (being METF the Mean number of Events to Failure) [12].

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

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

The proof of this 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 [13]. 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,

[1] Not 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.

$$MTTF = \frac{1}{\dfrac{1}{MTTF_1} + \dfrac{1}{MTTF_2}}. \quad (5)$$

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 [14] for more details). Therefore, in our case, equation (5) is only an approximation for the MTTF of the memory:

$$MTTF_{memory} \cong \frac{1}{\dfrac{1}{MTTF_d} + \dfrac{1}{MTTF_a}}. \quad (6)$$

This approximation will be used in the following section to asses 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 (65nm and 45nm) for which MCUs are a major concern and large IDs are normally used to ensure that no direct failures occur. The memories were exposed to white beams up to 800 MeV at the LANCE site and neutron beams up to 180 MeV at the TSL site. Multiple devices were used and for 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 [11].

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).

Once these parameters have been determined, 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 65nm memories were added together so that a single value is shown. However, the individual α would be similar in any case.

On the other hand, 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 Figure 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.

TABLE I
AVERAGE NUMBER OF ERRORS PER EVENT

	65nmA	65nmB	45nmA	45nmB
α	2.0649	2.0649	1.9062	1.7573

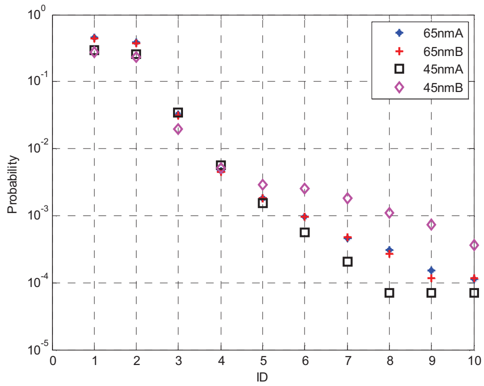

Fig. 2. Values of e^{ID} obtained in the experiments.

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 equation (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), $MTTF|_d \rightarrow \infty$, what would lead to $MTTF|_{memory} \rightarrow MTTF|_a$, per expression (6). In other words, 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 Figures 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.

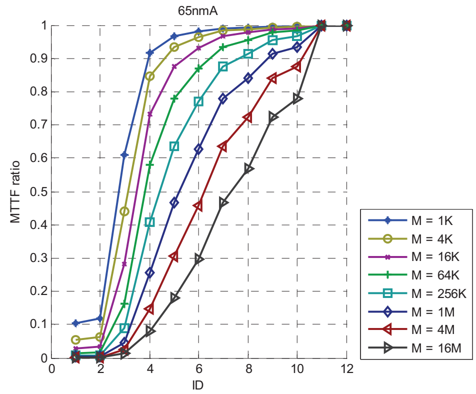

Fig. 3. $MTTF_{ratio}$ for the 65nmA memories in different configurations

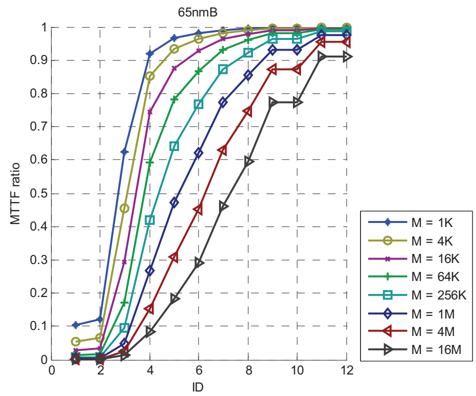

Fig. 4. $MTTF_{ratio}$ for the 65nmB memories in different configurations

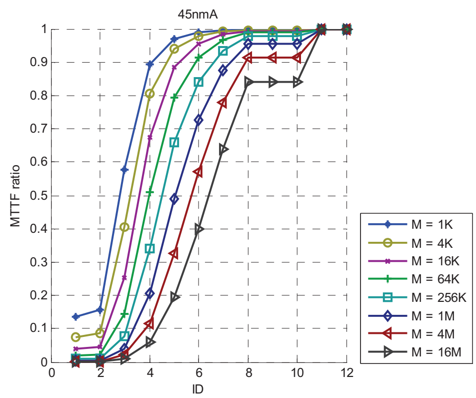

Fig. 5. $MTTF_{ratio}$ for the 45nmA memories in different configurations

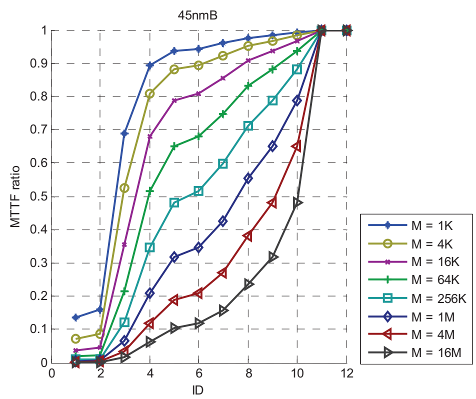

Fig. 6. $MTTF_{ratio}$ for the 45nmB memories in different configurations

In order to make a more detailed analysis of the ID vs reliability trade-off, let us now focus on the design of a 256-Kword memory. Let us also consider a reliability goal such that direct failures can only have an effect of 10% or less on the MTTF. Then, using the plots in figures 3 to 6, the minimal ID distances can be obtained in order to meet that goal. These results are depicted in Table II where a conservative value for the ID is also proposed. The conservative ID is defined in such a way that direct failures have a negligible impact, thus having no effect on the MTTF. The results show that ID values smaller than the maximum MCU size can be used in some cases. This will reduce the area and power of the memory making the design more competitive. To illustrate the benefits of the proposed approach, the cost of a memory with the previously selected ID values has been compared. The cost calculation is based on data from [11], summarized in Table III where the relative area/power overhead versus an ID of four are shown. The ID values available are only powers of two in this particular memory design, a situation that is common unless a full-custom design is made. It can be seen that both the area and power increase significantly with the ID.

In Table IV, the relative area overhead has been described for the ID values determined in Table II. For each case presented in Table II the closest power of two that is equal or larger than the required ID in each case (ID_{min}, $ID_{conservative}$) is selected from Table III. Those values are also shown in parenthesis in Table II. 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 (65nmA, 65nmB, 45nmA) 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 II
ID VALUES

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

TABLE III
AREA OVERHEAD FOR DIFFERENT ID VALUES

ID	Area increment	Power Increment
4	1	1
8	1,027	1,072
16	1,316	1,247
32	2,031	2,016

TABLE IV
AREA INCREMENT FOR THE TWO ID CONFIGURATIONS

	65nmA	65nmB	45nmA	45nmB
ID_{min}	1,027	1,027	1,027	1,316
$ID_{conservative}$	1,316	1,316	1,316	1,316

TABLE V
POWER INCREMENT FOR THE TWO ID CONFIGURATIONS

	65nmA	65nmB	45nmA	45nmB
ID_{min}	1,072	1,072	1,072	1,247
$ID_{conservative}$	1,247	1,247	1,247	1,247

IV. CONCLUSIONS

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,” IEEE Custom Integrated Circuit Conference, pp. 437-444, 2006.

[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,” 2004 IEDM, pp.945-948, Dec. 2004.

[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,” *IEDM*, pp. 941-944, Dec. 2004.

[5] P. Reviriego, J. A. Maestro, and C. Cervantes, “Reliability analysis of Memories suffering multiple bit upsets,” IEEE Trans. On Device and Materials Reliability, vol. 7, no 4, pp 592-601, Dec. 2007.

[6] G.C. Yang, “Reliability of semiconductor RAMs with soft-error scrubbing techniques”, IEE Proceedings Computers and Digital Techniques, Volume 142, Issue 5, Sept. 1995 Pages: 337 – 344.

[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, pp310-312, June 2000.

[8] A. D. Tipton, J. A. Pellish, 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 Transactions on Nuclear Science, Volume 53, Issue 6, Part 1, Dec. 2006 Pages: 3259 – 3264.

[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 Transactions on Nuclear Science, Volume 53, Issue 6, Part 1, Dec. 2006 Pages: 3139 – 3144.

[10] J. Maiz, S. Hareland, K. Zhang and P. Armstrong, “Characterization of multi-bit soft error events in advanced SRAMs”, IEEE International Electron Devices Meeting, 2003. IEDM’03 Technical Digest, Dec. 2003, Pages: 21.4.1 - 21.4.4.

[11] S. Baeg, S. Wen and R. Wong, “SRAM Interleaving Distance Selection with a Soft Error Failure Model”, IEEE Transactions on Nuclear Science Vol 56, Issue 4, Part 2, pp.2111 – 2118, Aug. 2009.

[12] W. Feller, “An Introduction to Probability Theory and its Applications” Hoboken, NJ: Wiley, 1968.

[13] I. Koren, C. M. Krishna, ”Fault-Tolerant Systems” Morgan Kaufmann, 2007. (ISBN: 0120885255).

[14] S. Baeg, P. Reviriego, J.A. Maestro, S. Wen, and R. Wong, "Analysis of a Multiple Cell Upset Failure Model for Memories", Proc. of the 2009 IEEE Workshop on Silicon Errors in Logic - System Effects (SELSE'09), Stanford University (USA), March 2009.