# Introduction to the Special Issue on Soft Errors and Data Integrity in Terrestrial Computer Systems

## I. INTRODUCTION

More than a quarter century ago, May and Woods of Intel reported on alpha-particle induced soft errors in their 2107 series 16Kb DRAMs. This paper represents the first public account of radiation-induced upsets in electronic devices at sea-level. The problem reported was caused by trace radioactivity in the memory packaging materials that emitted alpha particles. The interaction of these particles with the silicon substrate generated excess charge that was collected by the memory devices and in some cases yielded bit flips which resulted in incorrect data being stored.

A year later Ziegler and Lanford of IBM predicted that cosmic rays could result in the same phenomenon at sea-level. While it was known that cosmic rays could generate errors in electronic components in space, both papers came as a big surprise at that time, because nobody could have imagined that electronic devices, let alone whole systems, would be sensitive to ionizing particles at sea-level. Memory devices in the late 1970s stored about 1 million electrons and at the time it seemed virtually impossible that light particles such as alpha-particles could generate sufficient charge to result in failures. It is equally surprising that nowadays, almost 30 years later, with devices storing just few thousands of electrons, we are still manufacturing reliable chips and systems.

Radiation-induced soft errors (SE) have become one of the most important and challenging failure phenomena in modern electronic devices. Without a good understanding of the basic mechanisms and proper mitigation techniques, reliable electronic systems would not exist today. Mitigation schemes range from Error Correcting Codes (ECC) in memories, to radiation hardened devices and temporal or spacially redundant compute and compare schemes with overheads ranging from insignificant to very high. However, the recipes that work today might not be sufficient tomorrow. This is particularly true in the sub 65nm CMOS technology area that we are now entering. While the radiation-induced single bit Soft Error Rate (SER) of SRAMs and DRAMs is flat or even decreasing on a per bit basis, the failure rates of highly integrated chips and systems have the potential to increase dramatically over the next few process generations driven by Moore's law and its resulting increase in bit count. It could be argued, that this same statement was made soon after the May and Woods' discovery almost 30 years ago and the industry is still supplying reliable components. The reason has been that ECC schemes have almost completely eliminated the SER in memories at a minimal cost, and very high data integrity applications were able to afford the very expensive redundancy schemes. What is different now is that driven by dimensional reduction and increased latch count and functionality, the contribution of the logic (Latches, Register files, distributed small RAM arrays and combinational circuits) is not negligible anymore and has been predicted to geometrically increase with scaling. Thus, this appears to be the right time to both evaluate how well the protection schemes in use truly work and further examine the various proposed schemes to protect the logic itself.

## II. EDITORIAL PREVIEW

This special edition on soft errors induced by radiation focuses on radiation-induced soft errors and data integrity in commercial, terrestrial electronic systems. In contrast to the well known IBM Journal of Research and Development (volume 40, number 1, 1996), we are not addressing the details of the physical mechanisms behind the particle induced generation of charge but focus on its consequences, the reliability of circuits and devices from a system level perspective and techniques for mitigation. Most of the 12 papers are from industry, which is not surprising, since it is difficult for academia to keep up with the fast pace of technology development and high cost of experimental work. Academia can design their own test chips of course and have them built by one of the many fabless companies that exist nowadays, but those companies are usually not willing to share any of the detailed process information that is necessary to model and explain the observed phenomena. We propose that a more important role for academia may be in collaborating with the industry and work on discovering clever and cost effective methods to protect the logic, and in developing accurate simulation tools.

The historical perspective and the SER trends are summarized by Robert Baumann of TI in the first paper. This article briefly reviews the types of failure modes for soft errors that are responsible in terrestrial applications today and then addresses the sensitivity as a function of technology scaling for various memory and logic devices. The next section comprises three papers that provide experimental soft failure data. The first paper in this section is by Austin Lesea et al. of Xilinx. It compares soft error rates of configuration memory and block RAM in FPGAs obtained in field tests with simulation results and accelerated testing data. The failure rates observed in a large supercomputer, the Los Alamos National Laboratory's ASC Q, are reported on in the 3rd and 4th paper by Sara Michalak et al. from the Los Alamos National Laboratory and Kevin Harris et al. of HP,respectively.

The third section focuses on the modeling aspects of soft errors in integrated circuits and contains three papers. The first , paper number 5 by Paul Dodd of Sandia National Laboratories, reviews techniques for physics-based device-level simulation of single event effects. The Intel perspectives on circuit-level and chip/system-level soft error estimation techniques are presented in papers 6 and 7, by Steve Walstra et al. and Hang Nguyen et al. respectively.

The last section of this special issue on SER is devoted to soft error mitigation techniques. Process options for protecting SRAM devices and sequentials are discussed by Philippe Roche from ST Micro. Memory specific protection schemes as well as tradeoffs are discussed in detail by Charles W. Slayman of Sun Microsystems. An overview of the most popular circuit-level mitigation techniques available to designers to protect against single event upsets (SEU) and single event transients (SET) is given by Michael Nicolaidis of iRoc Technologies. Patrick Meaney et al. of IBM describe the error detection and recovery schemes within the IBM z990 series processors. The last paper of this special SER edition is

contributed by Jeffrey Wilkinson et al. of Medtronic Inc. and addresses data integrity issues and protection schemes typically used in medical electronic systems such as implantable cardioverter defibrillators and pacemakers.

The papers above, we hope, provide a cross-section of the ongoing activities and research to address data integrity issues in commercial ground-level systems. More importantly, we hope that it will stimulate the interest in the scientific community in developing novel and more effective methods in assessing the error rates and in developing cost effective protection schemes. We have purposely limited the number of Intel contributions to 2 papers only, although Intel scientists are actively engaged in all aspects of soft error modeling and mitigation. We did this in order to give the reader as broad an overview of the current state-of-the-art as possible. The selected Intel contributions summarize our understanding in a SEU field where we feel we have contributed and published significantly in the last 3-4 years.

## III. ACKNOWLEDGMENTS

It all began with the famous May and Woods paper in 1978. It is a great pleasure and honor for us to continue this tradition and as editors of this special issue, help in the communication of advances in the field and facilitate the discovery of novel methods and strategies to maintain high levels of data integrity. We however could not have done this without the commitment and dedication of the industry and academia experts that have contributed their papers and the support of the many reviewers that have unselfishly contributed their time and effort to help make this issue a success. Among them, Kerry Bernstein, Bharat Bushan, John Crawford, Richard Harper, Peter Hazucha, Hajime Kobayashi, Lloyd Massengill, Pankaj Mehra, Cecilia Metra, Subasish Mitra, Shubu Mukherjee, Sanjay Patel, Helmut Puchner, Stefan Rusu, Ronald Schrimpf, Norbert Seifert, Nelson Tam, Yoshiharu Tosaka, Kevin Zhang, James Ziegler. We would also like to acknowledge Anthony Oates, Editor-in-Chief of the IEEE Transaction on Device and Materials Reliability, Jo Ann Marsh from the IEEE EDS Publications Office and their staffs for their continuous help and support.

Jose Maiz, Guest editor
Intel Corporation
5200 NE Elam Young Pkwy
Hillsboro, OR 97229

Norbert Seifert, Guest editor
Intel Corporation
5200 NE Elam Young Pkwy
Hillsboro, OR 97229