This article is part of the IEEE Reliability Society 2008 Annual Technology Report.

Some Faults are Worse Than Others—And How That is Useful for Low-Cost Hardening

Ilia Polian
Albert-Ludwigs-University
Email: polian@informatik.uni-freiburg.de

Traditionally, a fault handling strategy is considered effective if it covers large classes of faults, e.g., all single faults. This conventional wisdom has recently been challenged by identifying subsets of faults which are acceptable at system level. One example is a fault in an imaging microchip which does not result in a deterioration of the calculated image to the extent that a human viewer would notice the difference. In context of micro- and nanoelectronic circuits, dropping the restrictive requirement that all faults be covered enables cost-effective selective hardening solutions, where only parts of the circuit are equipped with fault protection mechanisms. In this way, the need for traditional massive redundancy schemes such as triple-modular redundancy, which are associated with massive area and energy consumption overheads, is eliminated. Handling most critical faults could be associated with overheads as low as 10%, which is practical even for cost-sensitive embedded systems with a limited energy budget.

There are two enabling technologies for selective hardening based on fault criticality. First, it must be possible to harden parts of the circuit while not spending chip area and energy budget for the parts which require no protection. Techniques recently developed are able to perform selective hardening with ultra-fine granularity: it is possible to specify individual logic gates to be hardened [1] - [4] while other gates remain unchanged and do not cause hardening costs. Second, proper methods must determine which faults may result in critical system behavior, as opposed to non-critical faults, which require no protection. These methods are addressed in the following in more detail.

Fault criticality was first studied in context of permanent chip manufacturing defects: a circuit with a defect known not to cause critical effects on system level could be sold at lower price rather than thrown away, thus increasing effective yield. A number of generic metrics such as error significance and error rate [5] as well as specific approaches for multimedia circuits [6,7] were proposed in the last few years.

More recently, the research focus turned to transient or soft errors caused by such mechanisms as electrical noise or cosmic radiation. In contrast to manufacturing defects, their impact to the system operation is limited by a very short period of time, typically one clock cycle. However, the error effect may be propagated to the memory elements and thus corrupt the system state. Consequently, one definition of non-critical errors requires that their effects be eliminated from the system state within a small number of clock cycles. In other words, errors from which the system recovers itself within a short period of time do not need to be handled. It was proven in Polian 2008 [8] that over 70% of possible error spots in an MPEG subsystem had the property that an error on one of these spots was non-critical with respect to the definition above, irrespective of the system input or system state. This concept has been enriched by probabilistic aspects in Hayes 2007 [9]. It was shown that the rate of critical errors can be reduced by several orders of magnitude by hardening 10% of the circuit or less.

A couple of application-specific criticality definitions were also proposed and studied. A communication chip with a set of formal properties describing its specification was considered in Seshia 2007 [10]. A model checker was used to formally show that errors in approximately two-

thirds of the chip's flip-flops did not lead to system behavior which violated the properties. Hence, only the remaining one third of the flip-flops required hardening. A large-scale fault injection study by May 2008 [11] demonstrated the resilience of a rather complex communication device to randomly injected soft errors up to a certain error rate. Further application-specific metrics were studied in Li 2007 [12]. In Nowroth 2008 [13], the concept of cognitive resilience was defined to denote the ability of a human user to compensate the effects of certain soft errors by her cognitive apparatus. It was shown that less than half of the flip-flops in a JPEG compressor need to be hardened according to that definition.

The future of computing is expected to be centered on a class of applications commonly known as recognition, mining and synthesis (RMS). These methods will tackle difficult problems such as natural language translation by identifying patterns in large sets of data and correspondences to data-sets already resolved (recognition); automatically learning new patterns and correspondences (mining); and derive problem solutions from principles learned (synthesis). It is obvious that these approaches must have mechanisms to deal with incorrect decisions; hence they can also be expected to be intrinsically tolerant against hardware faults, rendering only a small portion of errors critical. On the other hand, the continuing transition to nanoelectronics will result in error rates further increasing. Hence, there appears to be bright future for criticality-based hardening for applications of tomorrow.

References

[1] K. Mohanram and N. Touba, Partial error masking to reduce soft error failure rate in logic circuits. In Proc. IEEE Int'l Symp. on Defect and Fault-Tolerance, Boston, MA, USA, 2003, pp. 433–440.

[2] K. Mohanram and N. Touba, Cost-effective approach for reducing soft error failure rate in logic circuits, in Proc. IEEE Int'l Test Conf. Charlotte, NC, USA, 2003, pp. 893–901.

[3] A.K. Nieuwland, S. Jasarevic, and G. Jerin. Combinational logic soft error analysis and protection. In Proc. Int'l On-Line Test Symp., 2006.

[4] R. Garg, N. Jayakumar, S. Khatri, and G. Choi. A design approach for radiation-hard digital electronics. In Proc. IEEE/ACM Design Automation Conf., 2006, pp. 773–778.

[5] Z. Pan and M.A. Breuer. Basing acceptable error-tolerant performance on significance-based error-rate (SBER). In Proc. IEEE VLSI Test Symp., 2008.

[6] I. Chong and A. Ortega. Hardware testing for error tolerant multimedia compression based on linear transforms. In Proc. IEEE Int'l Symp. on Defect and Fault Tolerance in VLSI Systems, 2005.

[7] H. Chung and A. Ortega. Analysis and testing for error tolerant motion estimation. In Proc. IEEE Int'l Symp. on Defect and Fault Tolerance in VLSI Systems, 2005.

[8] I. Polian, B. Becker, M. Nakasato, S. Ohtake, and H. Fujiwara, Low-cost hardening of image processing applications against soft errors, In Proc. Int'l Symp. on Defect and Fault Tolerance, pp. 274-279, Arlington, VA, USA, 2006.

[9] J. Hayes, I. Polian, and B. Becker, An analysis framework for transient-error tolerance, In Proc. IEEE VLSI Test Symp., Berkeley, CA, USA, 2007.

[10] S.A. Seshia, W. Li, and S. Mitra. Verification-guided soft error resilience. In Proc. Design, Automation and Test in Europe Conf., 2007.

[11] M. May, M. Alles and N. Wehn. A Case Study in Reliability-Aware Design: A Resilient LDPC Code Decoder. In Proc. Design, Automation and Test in Europe Conf., 2008.

[12] X Li and D Yeung. Application-Level Correctness and its Impact on Fault Tolerance. In Proc. Int'l Symp. on High Performance Computer Architecture. 2007.

[13] D. Nowroth, I. Polian, and B. Becker. A study of cognitive resilience in a JPEG compressor. In Proc. IEEE/IFIP Int'l Conf. on Dependable Systems and Networks, Anchorage, AK, USA, 2008.