Initial Results on the Importance of Protecting Prediction Arrays Against Hard-Faults

Veerle Desmet* Yiannakis Sazeides+ Costas Vrioni+

*Dept. of Electronics and Information Systems
Ghent University, Belgium
veerle.desmet@elis.UGent.be

+Dept. of Computer Science
University of Cyprus, Nicosia
{yanos,cs03vk1}@cs.ucy.ac.cy

Abstract

Continuous circuit and wire miniaturization increasingly exert more pressure on the computer designers to address the issue of reliable operation in the presence of hard-faults. Virtually all previous work on hard-fault reliability addresses problems that arise when a fault occurs in architectural resources, such as the register file or caches. However, hard-faults can happen in non-architectural resources, such as prediction arrays and replacement bits. Although these non-architectural hard-faults do not affect correctness they may degrade a processor performance significantly and, therefore, render them as important to deal with as architectural hard-faults.

In the past, because faults were more rare, it was acceptable for low-end systems to offer little or no protection against faults. As a result, mainly processors used in high availability systems employed advanced fault-tolerance techniques, such as using redundant and spare units [1]. With technology projections pointing to a dramatic fault increase in processors [2] a more general use of fault-tolerance techniques is emerging.

In this research we determine, using previously proposed analytical models, under what temperature conditions hard-faults in non-architectural structures are likely to occur in the same order of magnitude as hard-faults in architectural units. Furthermore, we quantify the performance implications of hard-faults in two prediction arrays: a line predictor and a return-address-stack. In particular, a simulation based analysis of a high-end processor that experiences a single stuck-at fault in one of its most frequently used cells in the return-address-stack and the line predictor, revealed a degradation up to 9% and 3%, respectively. When a single stuck-at hard-fault occurs in one of the output drivers the slowdown can be as high as 34% in the return-address-stack and 19% in the line predictor.

The above findings underline the importance to protect prediction arrays against non-architectural hard-faults. Our future work will explore the use of low-overhead detection and correction techniques for non-architectural hard-faults that are found to be more performance critical, to ensure future processors can operate with minimal degradation at the presence of non-architectural hard-faults. We will leverage existing techniques that have been proposed for error detection/correction of architectural structures, but, non-architectural resources may provide a distinct opportunity for simpler detection and correction techniques since they do not require a full repair.

References

[1] P. J. Meaney, S. B. Swaney, P. N. Sanda, and L. Spainhower. IBM z990 soft error detection and recovery. *IEEE Transactions on Device and Materials Reliability*, 5(3):419–427, Sept. 2005.

[2] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers. The impact of technology scaling on lifetime reliability. In *Proceedings of the 34th Annual International Conference on Dependable Systems and Networks*, pages 177–186, June 2004.