

Abstract

Relentless CMOS scaling coupled with lower design tolerances is making ICs increasingly susceptible to transient faults, wear-out related permanent faults and process variations. Decreasing CMOS reliability implies that high-availability systems which were previously restricted to the domain of mainframe computers or specially designed fault-tolerant systems may become important for the commodity market as well. In this thesis we tackle the problem of enabling efficient, low cost and configurable fault-tolerance using Chip Multiprocessors (CMPs).

Our work studies architectural fault detection methods based on redundant execution, specifically focusing on “leader-follower” architectures. In such architectures redundant execution is performed on two cores/threads of a CMP. One thread acts as the leading thread while the other acts as the trailing thread. The leading thread assists the execution of the trailing thread by forwarding the results of its execution. These forwarded results are used as predictions in the trailing thread and help improve its performance. In this thesis, we introduce a new form of execution assistance called critical value forwarding. Critical value forwarding uses heuristics to identify instructions on the critical path of execution and forwards the results of these instructions to the trailing core. The advantage of critical value forwarding is that it provides much of the speedup obtained by forwarding all values at a fraction of the bandwidth cost.

We propose two architectures to exploit the idea of critical value forwarding. The first of these operates the trailing core at lower voltage/frequency levels in order to provide energy-efficient redundant execution. In this context, we also introduce algorithms to dynamically adapt the voltage/frequency level of the trailing core based on program

behavior. Our experimental evaluation shows that this proposal consumes only 1.26 times the energy of a non-fault-tolerant baseline and has a mean performance overhead of about 1%. We compare our proposal to two previous energy-efficient fault-tolerant CMP proposals and find that our proposal delivers higher energy-efficiency and lower performance degradation than both while providing a similar level of fault coverage.

Our second proposal uses the idea of critical value forwarding to improve fault-tolerant CMP throughput. This is done by using coarse-grained multithreading to multiplex trailing threads on a single core. Our evaluation shows that this architecture delivers 9–13% higher throughput than previous proposals, including one configuration that uses simultaneous multithreading (SMT) to multiplex trailing threads. Since this proposal increases fault-tolerant CMP throughput by executing multiple threads on a single core, it comes at a modest cost in single-threaded performance, a mean slowdown between 11–14%.