Design and Simulation of an EM-Fault-Tolerant Processor with Micro-Rollback, Control-Flow Checking and ECC

Franco Trovo, Shantanu Dutt* and Hasan Arslan
Dept. of Electrical & Computer Engr., University of Illinois at Chicago

Electromagnetic (EM) radiation can cause significant disruption to computer systems as demonstrated in previous works. Previous research dealing with fault detection and tolerance in digital and computer systems have typically considered one to two random faults occurring at a time, arising mainly from operational problems like electromigration, transistor malfunction due to, say, heating, and random noise events. For such disruptions, single or double random faults are indeed a reasonable fault assumption and also makes its analysis and simulation tractable. While data on the fault-extent and fault-pattern caused by EM disruption is not completely available, we postulate that most EM disruptions will result in multiple clustered faults, and further that external (i.e., off-chip) wiring, like, memory address and data buses, and input/output buses are most vulnerable to such disruptions. With these assumptions, we have developed and simulated a processor system that can withstand such fault patterns to a reasonable degree.

The core of our error detection mechanisms are: 1) Hamming error-correcting coding (ECC) of data/address on the memory data/address buses; 2) Control-flow checking (CFC) of the program under execution by a simple watchdog (WD) processor; 3) Reasonableness checks of addresses on the address bus by the WD; 4) Inherent processor detection of invalid instructions and other exceptions (e.g., divide-by-zero). The fault tolerance (FT) mechanism we have used is *microrollback*. In this technique, the ``commitment" of the processor state (internal special registers, register file, cache and main memory contents) after each instruction execution is delayed via FIFO buffers interfacing with these storage entities by a maximum of d instruction cycles. Two of the above four detection techniques (ECC and CFC) are coupled to the microrollback , and when either detects an error, the microrollback unit rolls back the processor states to the one existing i instructions earlier, where i is the maximum error detection latency of the particular scheme that detected the error (the maximum rollback distance d is chosen to be the maximum of the maximum error latencies of the two detection mechanisms). The processor thus resumes execution from a point just prior to the fault event that caused the error. We believe this can be an effective FT technique for EM disruptions, because ``malignant" EM radiation is characterized by short-duration repeating pulse trains. Each such pulse train can cause transient errors in the data/address on the memory buses which can be recovered from via the microrollback mechanism if the errors are detected by the above methods.

We implemented the above fault detection and tolerance mechanisms in VHDL for the Motorola 68040 processor. We simulated a variety of fault patterns and numbers on both memory buses ranging from two random faults to 4-bit cluster faults that we believe could be typical of fault patterns caused by EM radiation. We also simulated various frequencies of EM pulse-train repetition ranging from the most disruptive (one pulse train every bus clock cycle) to a low disruptive one (one pulse train every 100 bus clock cycle). Our current simulation results (we are in the process of tuning and improving our techniques) reveal that the program we ran (matrix multiplication) completed correctly under these fault environments and the above fault detection and tolerance methods between 67-100% of the time. The program finishes incorrectly only 0-16% of the time, while for the rest of the runs it does not finish; most probably this is due to the program getting into an infinite loop, and this can be detected by a timeout mechanism and the program can be stopped in a ``fail-safe" manner. These results are promising and demonstrate that our techniques could be effective in protecting computer systems from actual EM disruptions.