05Dynamic and Static Cache SEE Results for a 49-core Processor

Brian Wie, Michael Plante, Antwong Berkley, Reza Nejad, Integrity Applications Incorporated;
Steven Guertin, NASA JPL

***Abstract* – Recent dynamic and static results are presented on a 49-core RHBD processor. Dynamic cache results show increased SEE susceptibility over static tests.**

Introduction

Selecting a space qualified processor with high throughput capability is becoming more challenging as spacecraft on-board processing demands continue to increase. Programs are now beginning to depend on upscreened COTS and for the more risk averse, relying on radiation-hardened-by-design (RHBD) processors, which utilize design enhancements at commercial foundries to mitigate the effects of Single Event Effects (SEE) and Total Ionizing Dose (TID).

The RHBD Maestro many-core processor, shown in Figure 1, consists of 49 individual processors (also referred to as "tiles") each able to run its own separate application code. Each tile contains 8kB L1 data cache, 8kB of L1 instruction cache, and 64kB of L2 cache [1]. A single tile may also use any other tiles' caches to serve as a virtual L3 cache. All cores are connected via 5 on-chip networks that pass data between each other and to XAUI, Gigabit Ethernet, DDR2, and other interfaces. The tiled architecture could also provide fault tolerant resilience and use dedicated tiles for redundancy [2, 3]. The 4x 10-Gigabit Ethernet ports offer a standard high-speed I/O interface to other payloads.

Static testing results on the Maestro cache memories and registers are summarized in [4]. Earlier test chip versions of a single Maestro tile device were tested by Boeing [5-7]. Static and dynamic testing results presented in this paper were conducted at the NASA Space Radiation Laboratory at Brookhaven.

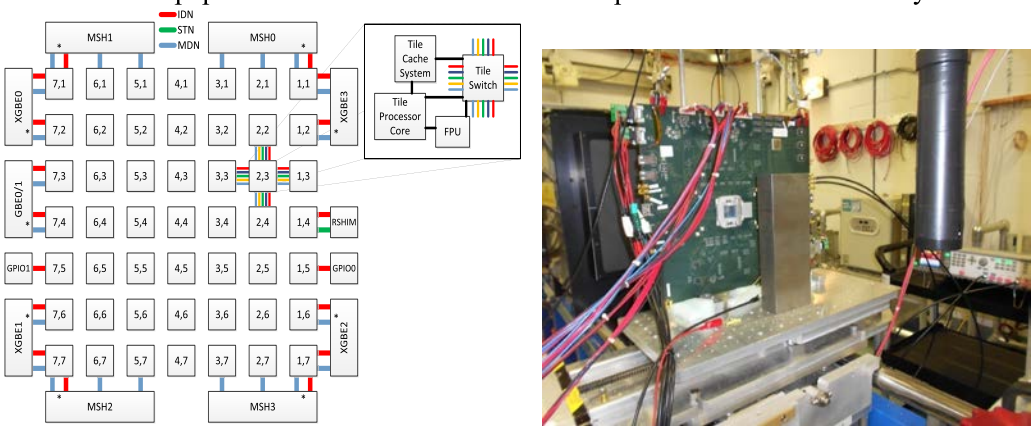

Figure 1– Maestro Device layout and NSRL test setup

Test Approach

According to [4] the combination of both RHBD and Fault Tolerance in the Maestro device adds levels of complexity not addressed by any established test methods. A multi-core processor's inherent redundant architecture lends itself well to a self-testing methodology, which is described in [4] and was utilized in the NSRL tests. The master tile (1,4) in Figure 1 administers the test to the other 48 tiles during beam exposure. The master tile and RSHIM I/O interface were always shielded and, due to mechanical constraints, the first column of 7 tiles was also shielded. This ensured the master tile and I/O interface could not be directly upset during beam exposure. A 3-inch thick tungsten block was used to shield the first column of tiles. Figure 3 shows L1D-cache static cross sections for each of the 49 cores for a 1cm and 4cm diameter beam validating the effectiveness of the tungsten shield.

The general test architecture of the register static test, register dynamic test, L1D static test, and L2 dynamic test all followed a common pattern and included an L2 static test. This test architecture

involved booting the device, preparing the test environment, then performing a roughly 2.68 second test loop repeatedly in which test tiles essentially sleep to allow upsets to accumulate during static test. Instead of sleeping/ dwelling for 2.68 sec during radiation exposure as in static testing, actions are continually performed the entire 2.68 sec for the dynamic tests. Note that in the L2 case, the code portion of the L2 is still tested statically ("NSRL L2 Static Alt" in Figure 4) as before, but what is new is that the previously unused portion of the L2 is tested dynamically. L1 cache is protected with parity, and L2 cache is protected with SECDED EDAC circuitry which protects 8 data-bits with 5 check-bits.

The L2 dynamic test consists of having a nonmaster tile write from the end of its code to the end of the same way of L2 with an address unique pattern and then reading it back. This is intended to catch problems with write enable logic that might cause multiple cache lines to receive the same data that was intended for only one line. It has a high duty cycle, so it is considered dynamic. It has the side benefit, however, of providing a test with a very high scrub rate and therefore reducing the likelihood of multiple particle effects. A distinction has been made between L2 "test area" (the dynamic L2 test region that distinguishes this test) and the L2 "code area" (the static L2 test region that all the tests must scrub to keep themselves alive longer).

Results

The Maestro test setup is shown in Figure 1. Beam ion, energy, LET, and range are shown in Table 1.

Ion	Energy (MeV/amu)	LET in Si (MeV-cm^2/mg)	Range in Si (um)
^{56}Fe	1000	1.2	146,500
^{56}Fe	150	2.9	7,350
^{131}Xe	214	10.1	7,390
^{181}Ta	182	20.0	4,750

Table 1 – NSRL Beams

Four devices were tested – three were thinned to 80µm and one was left unthinned. All measurements were performed at normal incidence. The beam spill structure at NSRL is shown in Figure 2 on two scales comparing the static and dynamic tests. A 300ms spill of particles occurs every 3.8sec. The particle flux is counted by a 1.08 cm^2 scintillator if the ion has sufficient range; otherwise, the upstream ionization chamber is used to convert absorbed dose to particle counts. If there is no EDAC, the beam structure may be neglected in analysis and only the mean flux need be considered, but, for Single Error Correction and Double Error Detection (SECDED), the expected value of the observed event count goes as the square of the fluence per scrub period, so NSRL's pulsed beam increases the average number of events over the value expected based solely on a calculation using mean flux.

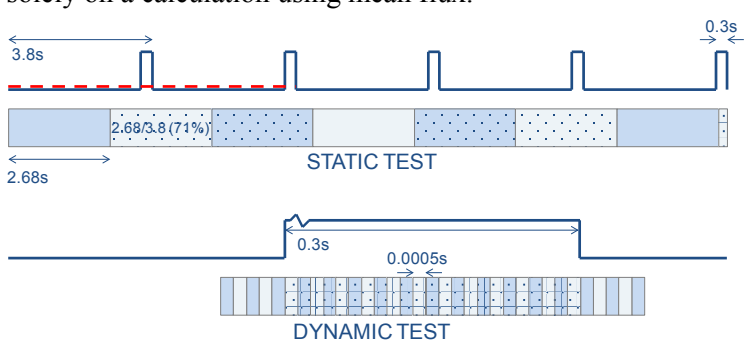

Figure 2- NSRL Beam Structure with Dynamic and Static Test Periods

Throughout all testing there was no evidence of any damage or high current events on any devices. Testing was performed in nominal operating conditions, rather than worst-case conditions for SEL. No error types were encountered that were unrecoverable; all errors conditions were recovered

from by supplying a hard asynchronous reset through an FPGA, which is under software control and did not require any power cycling.

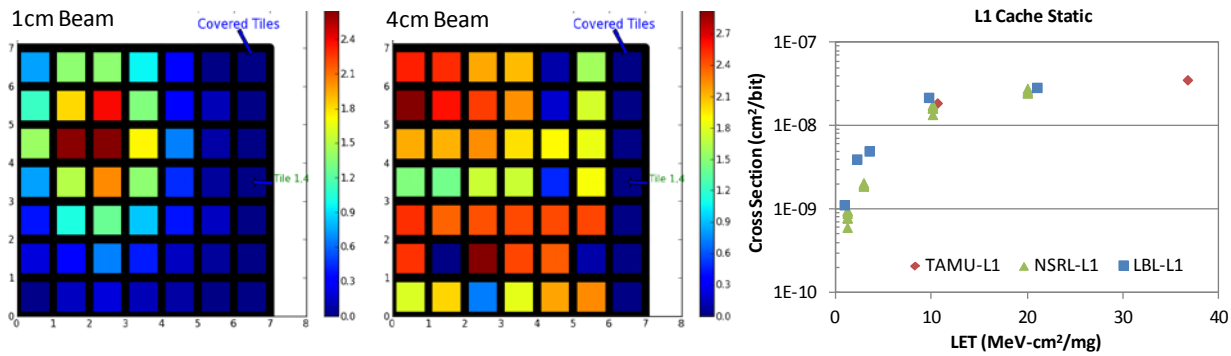

Figure 3 - Maestro L1 Cache Cross Section and Tile Profile

Master crashes/hangs were observed with particularly harsh beams. It was determined at TAMU [4] that the primary cause of nonmaster tiles crashing was the fact that L2 tag double-bit errors led to L2 misses and hangs on the MDN/MSHIMs since no external RAM is present. Tilera explained that this cannot be worked around on Tile64 or derivatives like Maestro, although it may be possible to find a solution on TilePro. A complete set of external RAM may be used to resolve the issue. [4] shows findings regarding L2 scrubbing with and without evicting data from L2. The success of this algorithm in dramatically reducing crash rates demonstrated that simply rewriting the L2 does not repair the L2 tags, even though it does repair the L2 data.

No verifiable single bit upsets were observed for static register testing. During static register testing, registers 2, 3, and 4 exhibited a data clobber (1's pattern mostly zeroed out to one of a handful of repetitive values) multiple times across different devices. Out of the 56 32-bit registers tested on each tile, the remaining registers usually held their original 1's pattern. This upset mode could be due to a program counter clobber or a sensitive reset or clock line on the device. In every case where registers 2, 3, and 4 were clobbered, the program counter for each of those tiles was also clobbered.

Both the TAMU and NSRL L1D cache static cross sections are shown in Figure 3. The different data points of the same color/type represent separate runs across multiple devices. For the NSRL and TAMU L1D data, the error bars are smaller than the data points. Cannon [6] collected data on similar 90 nm (IBM 9sf) commercial and RHBD SRAMs at Lawrence Berkeley National Laboratory (LBL). The commercial 90nm SRAM heavy ion data from the 10MeV/n cocktail at LBL is also included in Figure 3.

The dynamic versus static cross sections for L2 cache are shown in Figure 4. 1sigma error bars are shown for the NSRL L2 data. As expected, the dynamic cross section is higher than the static cross section due to Single Event Transients (SET). The dynamic access rate of the L2 cache was 2kHz, meaning each bit was rewritten or scrubbed once every 0.5ms. Both static and dynamic testing utilized a 100MHz input clock to the processor. It takes approximately 6 clock cycles to write or read a 32-bit word to/from cache.

An interesting observation is that for the dynamic L2 test, which increased the scrub rate by 5400, only single bit and double bit errors were observed. An observed single bit L2 error is actually from 2 upsets: 1 bit upset in the data byte and 1 upset in the check bits. An observed double bit error is from 2 bit upsets in a single data byte. No multi-bit upsets (MBU) were seen, unlike in the L2 static tests. This implies that the MBUs seen in the static L2 tests were not from a single particle, but from separate particles causing upsets cumulatively over the 2.68 sec scrub cycle. Taking this one step further, if the L2 is scrubbed fast enough, it may never have observable upsets, since the single bits are corrected preventing a DBU from ever occurring, and only DBUs are observable. An SET in the address decode logic may cause the same data pattern to be written to multiple addresses, which would manifest itself as

a multi-bit upset since the data pattern was address unique (data pattern=address). The reason why none of these decode logic transients were observed may be due to only running the processor at 100MHz. Further testing is planned to run dynamic tests at 300MHz to perhaps allow these decoder transients to be clocked into memory storage elements. It is assumed at the moment that clock, reset, and other combinatorial transients are responsible for part of the increase in cross section between static and dynamic testing; additional contributions may be from cache lines upsetting before being written to L2. [8] did not observe a cross section dependence on static versus dynamic testing of a IBM 90nm (9sf) RHBD SRAM test chip; however, [5] observed that the RHBD single-core Tilera Processor is susceptible to data integrity errors when performing data cache loads and stores, which frequently occurs during the dynamic L2 tests.

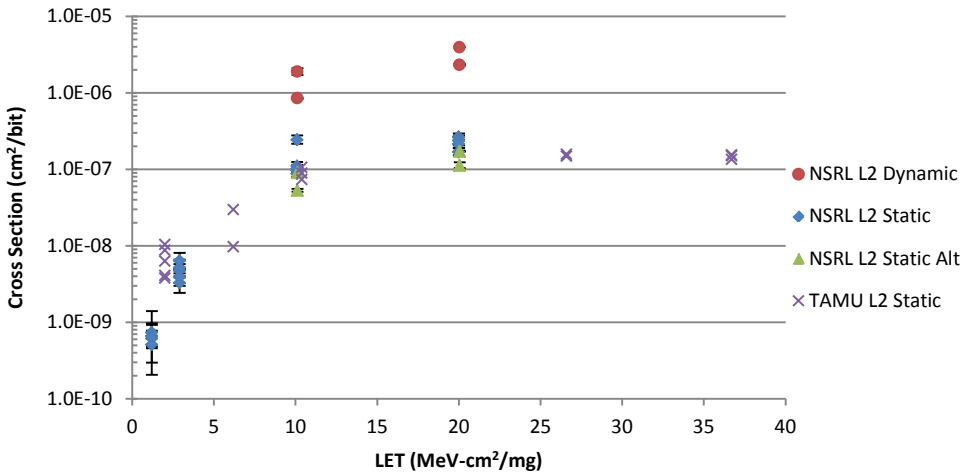

Figure 4 – Dynamic vs. Static L2 Cache

Conclusion

Results on dynamic L2 tests of the 49-core Maestro Processor show a significant increase over static test results. The increased susceptibility is likely from clock transients, reset transients, and SEE sensitivities in the cache lines, which are more easily observed during dynamic tests with a constant writing and reading of the cache. . In addition, static test results on both L1D and L2 caches utilizing the device self-testing approach were shown to be consistent with similar data sets on similar test structures collected by different methods at different heavy ion facilities.

References

[1] M. Malone, "OPERA RHBD Multi-core," MAPLD, Aug 2009.
[2] C. Villalpando, et.al., "Reliable Multicore Processors for NASA Space Missions." IEEEAC paper #1163, version 5, Nov 2010.
[3] J.P. Walters, et. al., "Software-Based Fault Tolerance for the Maestro Many-Core Processor," IEEEAC paper #1564, Jan 2011.
[4] S. Guertin, et. al., "SEE Test Results for Maestro Microprocessor," Radiation Effects on Components and Systems, Data Workshop, Sep 2012.
[5] M. Cabanas-Holmen, et. al., "Predicting the Single-Event Error Rate of a Radiation Hardened By Design Microprocessor," IEEE Trans. Nucl. Sci., vol. 58, no. 6, pp. 2726-2733, Dec. 2011.
[6] E. H. Cannon, et. al., "Heavy Ion, High-Energy, and Low-Energy Proton SEE Sensitivity of 90-nm RHBD SRAMs," IEEE Trans. Nucl. Sci., vol. 57, no. 6, pp. 3493-3499, Dec. 2010.
[7] M. Cabanas-Holmen, et. al., "Clock and Reset Transients in a 90nm RHBD Single-Core Tilera Processor," IEEE Trans. Nucl. Sci., vol. 56, no. 6, pp. 3505-3510, Dec. 2009.
[8] J. Draper, "Radiation Effects Challenges in 90nm Commercial-Density SRAMs: A Comprehensive SEE and TID Study," FTSCW, May 2008.