

Accelerated Testing of a 90nm SPARC64 V Microprocessor for Neutron SER

Hisashige Ando¹, Ken Seki¹, Satoru Sakashita¹, Masatoshi Aihara¹, Ryuji Kan¹, Kenji Imada¹, Masaru Itoh¹, Masamichi Nagai¹, Yoshiharu Tosaka², Keiji Takahisa³, Kichiji Hatanaka³

¹ Fujitsu, Kawasaki, Japan

² Fujitsu Laboratories, Akiruno, Japan

³ Research Center for Nuclear Physics, Osaka University, Ibaraki, Japan

Abstract—Soft-error tolerance of SPARC64 V microprocessor is measured with a series of accelerated tests using a white neutron beam facility of Osaka University.

It is found that a derating factor (before ECC) of level-2 cache SRAM is more than 50% for most of test programs measured. But, the derating factors of level-1 caches for the same set of programs are much lower than those of level-2 cache. All of these detected level-1 and level-2 SRAM errors were corrected by ECC and some other recovery means.

93.6% of estimated occurrences of latch errors were harmlessly vanished and the remaining 6.4% were detected or left some trace of an error. With the logic error detection and recovery mechanisms incorporated in the SPARC64 V microprocessor, 76% of the detected logic errors were corrected. Based on these measurements, the system down failure rate of the SPARC64 V microprocessor due to terrestrial neutron hits is confirmed to be well below 10fit.

Index Terms—soft-error, neutron beam, error detection, error correction, microprocessors, reliability, SPARC.

I. INTRODUCTION

The occurrence of soft errors due to terrestrial neutron hit is an increasing concern as the integration level of an LSI is continuously increasing with the Moore's law. SPARC64 V microprocessor[1][2] is designed for mission critical unix servers which require high resilience to the soft errors for un-interrupted operation. Also, high data integrity is especially important because a silent data corruption in mission critical system can cause serious damage. To meet with these requirements, the processor was designed not only to correct SRAM errors, but also to detect errors in logic and to recover from those detected errors as much as practically possible. But, the actual benefit of these mechanisms has not been proven by measurement.

We have carried out an accelerated soft error testing of the Fujitsu's PRIMEPOWER 650 server using the white neutron beam facility of Research Center for Nuclear Physics of Osaka University.

II. SPARC64 V MICROPROCESSOR

This processor is built with a 90nm bulk CMOS technology with 10 layers of Cu metallization. It integrates 128KB level-1 instruction cache and the same 128KB level-1

data cache. It also integrates 4MB level-2 unified cache. Major specifications of the processor are summarized in Table. I .

TABLE I . 90NM SPARC64 V MICROPROCESSOR SPECIFICATIONS

Technology	90nm bulk CMOS, 10 layer Cu metallization
Chip size	18.46mm x 15.94mm
Clock Frq.	2.16GHz
Power dissipation	Max 65W @Vdd=-1V, 2.16GHz
CPU Core	Single Core, 4 inst. decode O-o-O, ~240K Latches
Caches	L1\$:128KB+128KB, L2\$: 4MB

Level-1 data cache and level-2 cache are protected with SECDED code. Level-1 instruction cache is parity protected since its original contents are in the main memory or level-2 cache with ECC protection. Other smaller SRAM structures, i.e., cache tags, TLB and branch history table are also designed to be able to recover from any single bit error.

The processor logic is protected by parity check to detect single bit logic errors. 83% of latches (excluding boundary scan and other latches not related to normal execution of instructions) are covered by parity check. Major data and address buses are also covered by parity check.

The error detection methods and the error recovery methods of the processor are summarized in Table. II .

TABLE II . ERROR DETECTION AND RECOVERY METHODS OF SPARC64 V MICROPROCESSOR

	Error Detection	Recovery
L1\$ Data, TLB	Parity	Invalidate & Miss
BRHIS	Parity	Branch Miss-prediction Recovery
L1\$ & L1D\$ Tag	Parity+Duplication	Use correct data and rewrite
L1D\$ Data, L2\$ Data & Tag	SECDED	Hardware ECC
Registers	Parity	Instruction Retry
ALU, Shifter, VIS	Parity Prediction	
Mult/Div	Residue check +Parity Prediction	

III. ACCELERATED TEST SET-UP

PRIMEPOWER 650 is an 8 CPU server with 4 CPU modules in one side of the server enclosure and the other 4 on the other side. 4 CPU modules on one side are placed in the course of the neutron beam.

The flux density of the original RCNP neutron beam is 2.8

$\times 10^9$ neutron/cm²/s. But, as we piggy-backed our experiment with the semiconductor test chip soft error measurement with those DUT placed in front of the server, the beam is attenuated. The beam is also attenuated by the server enclosure and CPU modules themselves as it travels downstream. So, we have estimated the beam flux at each CPU chip by measuring the number of level-2 cache SRAM bit flips. Fig.1 and Fig.2 shows the test setup. Fig. 3 shows the CPU module placement in the server.




Fig.1 Accelerated test setup. The diameter of the neutron beam at the output of the collimator is



Fig.2 Photograph of the test setup. 4 CPU modules are placed behind a shiny metal grid.




Fig.3 Photograph of the 4 CPU modules in PW650 server. Beam is irradiated from left to right.

IV. SOFT ERROR MEASUREMENT

A. Static Test

The contents of the 4MB level-2 cache are initialized and the neutron beam is turned on for 5 minutes. Then, the contents were read back and compared with the initialized values to detect neutron induced bit flips. These 5 minutes measurements were repeated many times.

Although some of the neutron hits resulted in multi-bit errors in an array, with the use of the 8:1 column multiplex, those multi-bit errors are distributed one bit each into multiple words. So, these errors are correctable with SECDED ECC.

Several double-bit errors were detected in the SRAM static

test. But, the failed bit positions are far apart and it is judged to be caused by the accumulation of errors from two random hits instead of a multi-bit errors caused by a single energetic neutron hit.

The latch bit flip was measured by initializing latches through a scan chain. The length of the beam exposure is much longer than the SRAM experiment as the number of latch bits is about 240K that is far less than the 32M SRAM bits.

These measurements set the basis of the estimated number of bit flips for the dynamic measurements described below. From these measurements, it is found that our latch is 1.3 times more sensitive than the level-2 cache SRAM cell.

B. Beam Attenuation

The incident neutron beam is attenuated as it passes through CPU modules. The CPU module is composed of aluminum heat sink, copper heat spreader, CPU chip itself, ceramic package, printed circuit board and steel bolster plate which totals to about 10mm thickness as shown in Fig.4.




Fig. 4 Side view of the SPARC64 V CPU module

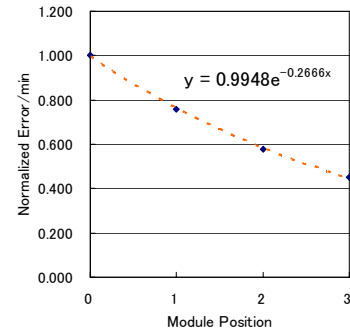


Fig. 5 Module position dependency of Normalize number of L2 cache SRAM errors from the dynamic tests.

Fig.5 shows a module position dependency of the normalized L2 SRAM bit errors. The frequency of the error is reduced to 76.6% as the module position shifts one position to downstream of the beam. The beam is attenuated by 1.8% due to the beam divergence alone. With this correction, beam attenuation due to the module passage is calculated to be 22%.

JEDEC standard JESD89-3[3] gives the following equation for attenuation.

According to this equation, the total attenuation from single CPU module should be 4% or less. It is believed that the scattering of the beam (and not hitting the downstream CPU chip) is a part of the reason of this big discrepancy, but

$$\text{Total Attenuation} = \prod_{i=1}^{i=n} \{1 - \exp[-t_i(\text{cm}) \times \rho_i(\text{g/cm}^3)/[216(\text{g/cm}^2)]]\}$$

the entire mechanism of the observed large attenuation is not well understood.

All 4 points in Fig.5 closely match with the exponential curve. This indicates that the variation in soft error rate among the CPU chips is minimal.

C. Dynamic SRAM Soft-error Measurement

The occurrence of the soft error was recorded while running 6 different programs, SVOP, LINPACK, gcc, bzip2, galgel and sixtrack. SVOP is an in-house hardware verification program which is designed to catch hardware failures. LINPACK is a standard HPC benchmark of linear equation solver. gcc and bzip2 are selected from the SPEC CPU2000 integer benchmark suite. galgel and sixtrack are selected from the SPEC CPU2000 floating point benchmark suite. All these programs check the result for the correctness of the program execution.

Table III summarizes the level-2 cache SRAM single bit error derating factors of each program before ECC.

TABLE III. DERATING FACTOR OF LEVEL-2 CACHE SRAM SINGLE BIT ERROR

SVOP	LINPACK	gcc	bzip2	galgel	sixtrack
0.301	0.524	0.554	0.730	0.539	0.651

Low derating factor of SVOP compared to other programs is due to the design of the hardware verification program. It runs without OS and touches very small portion of level-2 cache for performing tests except the level-2 cache test itself. For other application programs running under the unix operating system, derating factors range from 0.524 to 0.73.

Measured derating factors for the level-1 instruction cache and data cache SRAM single bit error are 24.3% and 14.9% respectively. For the level-1 caches, we could not record enough bit errors to identify the difference in derating factor between the programs. But, it is noticeable that the level-1 caches have much smaller derating factors than the level-2 cache. This is due to the shorter latency of ACE (architecturally correct execution) bits in level-1 cache as analyzed in [4].

With the implementation of the SECCED error correction mechanism and other recovery mechanisms, all single bit errors in these SRAM arrays were corrected. Hence the SRAM contribution to the chip level FIT is zero.

We removed multi-bit SRAM errors from the count as we believe those errors were caused by the accumulation of multiple random single bit errors. The frequency of accumulated double bit errors is proportional to the square of the acceleration and those multi-bit errors are an artifact of the accelerated testing.

D. Logic Soft-error Measurement with Programs

The occurrence of the logic errors was also recorded while running above mentioned 6 programs.

As the PRIMEPOWER 650 is an 8 CPU SMP server, innocent error free CPU can fail with watchdog timeout when a logic error occurred in bus control logic (including cache coherency mechanism between the CPU chips) of another CPU and the bus protocol hang. These induced failures are removed from the counts.

The total logic error count estimate is based on the latch

flip data from the static measurements only. The frequency of the logic error in combinatorial logic is neglected. Hence the following logic derating factors are conservative estimates.

For the overall measurement, 93.6% of estimated occurrences of latch flips were vanished, i.e., did not cause any failure nor leave traces. Remaining 6.4% of latch flips were detected by parity checkers and other error detection mechanisms. Then, error recovery mechanisms implemented try to recover from the error as summarized in Table II.

TABLE IV. DERATING FACTOR FOR LATCH ERRORS

		Frequency
Vanished Error		93.6%
Noticed Error	Recovered	4.9%
	Fatal Error	1.5%

Fig. 6 shows the logic derating factor for the noticed errors from the 6 tested programs. Paper [5] reports 3%~4% logic derating obtained from statistical fault injection logic simulations. Our result of logic derating for noticed error is in the same ball park with the reported result.

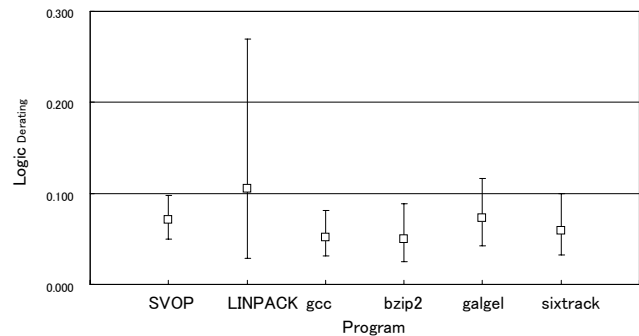


Fig. 6 Logic Derating for Noticed Errors

With the implementation of error detection and recovery mechanisms summarized in Table II, 76% of the noticed errors were recovered. Only 1.5% of estimated logic errors resulted in fatal (system down or SDC) errors. The measured system down error rate of 90nm SPARC64 V CPU is well below 10fit assuming the system is operated at NYC sea level. We have observed only one SDC. Although this is not a statistically significant data, from this measurement, its error rate is less than 1/10 of the system down failure rate.

V. SUMMARY

The accelerated soft error measurement was carried out for the SPARC64 V CPU manufactured with the 90nm bulk CMOS process. The white neutron beam facility of RCNP of Osaka University was used.

The error rates of the SRAM cell and the latch were measured while those elements were statically holding its data. Then, the error rate of the cache SRAMs and the logic were measured while various programs were running on the PRIMEPOWER 650 server.

Derating factor of level-2 cache single-bit error ranges from 30% to 73% among the programs measured. Derating factor of the level-1 cache SRAM single-bit error is 24.3%

for instruction cache and 14.9% for data cache. With the ECC and other correction mechanisms, all those single-bit SRAM errors were corrected.

For the overall measurement, it is estimated that 93.6% of the latch flips were vanished. With the implementation of the extensive error detection and recovery mechanisms, only 1/4 of the remaining latch errors resulted in system down failure. These measurements confirmed the effectiveness of the error detection and recovery mechanisms implemented in SPARC64 V microprocessor. From this accelerated testing, its system down failure rate due to the terrestrial neutron hit is confirmed to be well below 10fit.

ACKNOWLEDGMENT

The authors acknowledge the support from Tsuyoshi Motokurumada, Kunita Morita and Hideo Yamashita for the interpretation of logic error logs, Ryoza Takasu for beam flux measurement, Taiki Uemura for useful discussions. The authors would like to thank for the support offered by the people in Server systems group, Product business support group and Electron device group of Fujitsu and Fujitsu Laboratories.

These measurements became possible with the support of higher level managers, Akira Yamanaka, Noriyuki Toyoki and Masayoshi Kimoto.

REFERENCES

- [1] A.Inoue "Fujitsu's New SPARC64 V for Mission-Critical Servers", Microprocessor Forum 2002, Oct. 2002.
- [2] H.Ando, Y.Yoshida et. al "A 1.3GHz Fifth Generation SPARC64 Microprocessor", IEEE J.SSC, Vol.38, Issue 11, pp.1896-1905, Nov. 2003.
- [3] JEDEC Standard, "Test Method for Beam Accelerated Soft Error Rate JESD89-3", Sept. 2005
- [4] S.Mukherjee, C.Weaver, J.Emer, S.Reinhardt and T.Austin, "A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-Performance Microprocessor", MICRO-36, Proc. 36th Int. Symp. on. Microarchitecture, 2003.
- [5] H.Nguyen, Y.Yagil, N.Seifert and M.Reitsma, "Chip-Level Soft Error Estimation Method", IEEE Trans. Dev. Mat. Rel., Vol.5, No.3, pp. 365-381, Sept. 2005.