Evaluation of a 32-bit microprocessor with built-in concurrent error-detection

Jiri Gaisler
European Space Research and Technology Centre
2200AG Noordwijk, The Netherlands
email: jgais@ws.estec.esa.nl

Abstract

This paper describes the test results from heavy ion testing of ERC32, a 32-bit processing core with on-chip concurrent error-detection. The parity based error-detection mechanisms succeeded in detecting more than 97.5% of all injected errors, significantly reducing the MTBF for undetected SEU errors. Most errors occurred in registers, but some errors in combinational logic could also be observed. The cross-section for errors in combinational logic is however to small to have an influence on the overall error rate. The conclusion is therefore that parity based error-detection is well suited to detect SEU errors in VLSI devices for space applications.

1 Introduction

The ERC32 is a 32-bit SPARC compatible processor for space applications. It is implemented in three VLSI devices and manufactured on the Temic/MHS 0.8 µm SCMOS technology. This technology is radiation-hard from a total dose point-of-view but is sensitive to SEU errors. The ERC32 chipset therefore implements on-chip error-detection mechanisms to detect and isolate SEU errors [1]. This paper describes the results of the first system level SEU test campaign performed at Brookhaven National Laboratories (BNL).

2 ERC32 architecture

The ERC32 consists of three devices; a SPARC V7 integer unit (IU), a floating point unit (FPU) and a memory controller (MEC). Concurrent error-detection mechanisms are implemented in all three devices. The detection scheme is targeted to mainly detect transient errors, originating from single-event upsets (SEU). More than 99% of all on-chip registers are provided with parity bits which are checked each time the register content is used. External address, data and control buses are provided with parity bits to detect interconnection errors and SEU errors in output latches. Program flow control is implemented in the IU; a signature is generated from each executed instruction and can be checked at an arbitrary point in the program using a special instruction. To reach 100% device error-detection coverage, the IU and FPU are provided with on-chip comparators and can be operated in duplex (master/checker) mode.

The error-detection mechanisms in the IU and FPU are mapped on the trap generation logic. When an error is detected, an error trap is generated and the error can be handled and reported via the trap handler. The error traps are of two major types; restartable and non-restartable. Errors occurring in registers belonging to pipeline stages before the write stage are not affecting the state of the IU or FPU and are considered restartable. The trap handler for a restartable error simply returns to the interrupted instruction. The pipeline is refilled and the error is effectively removed.

The MEC does not have a pipeline and all MEC errors are therefore non-restartable. The error is reported to the error manager inside the MEC which will then either generate an interrupt to the IU, reset the system or halt.

3 Test procedure

Although SEU testing of VLSI circuits for space applications are commonplace, no experience exists of testing a complex VLSI device with as extensive parity-based SEU error-detection mechanisms as incorporated in the ERC32. The main test objective was therefore to determine the coverage and characteristics of the ERC32 error-detection mechanisms during typical operation, including measurement of the relative frequency of the various error types. To be able to calculate absolute error rates for typical orbits, the threshold LET (Linear Energy Transfer) and cross-section also has to be determined.

A single board computer with the ERC32 chipset, 2 Mbyte static RAM and 512 Kbyte flash PROM was used as test system (figure 1). The IU and FPU under test were connected to a “golden chip” [5] using the on-chip comparators, to guarantee detection of all manifested errors. The test software used was a mix of typical applications and special test programs exercising as many on-chip resources as possible. Each SEU event was reported to a hosts system that monitored the tests. The test software reported the error-type and if possible error type before re-starting the board. The test software was able to report more that 100 events/s in order to cope with a high ion flux. All error-detection mechanisms were enabled except the program flow control which could not be used due to lack of support in the used compiler (gcc).

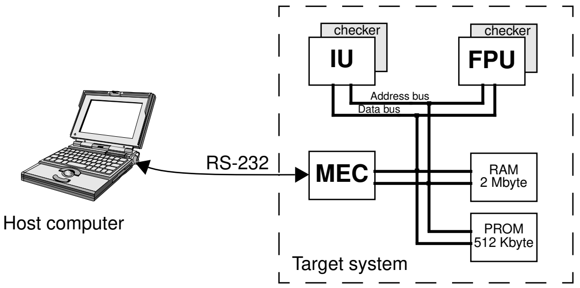

Figure 1: SEU test system architecture

4 Test results

During testing, the three devices were individually radiated using ions with the effective LET ranging from 20 to 82 MeV. Figure 2 shows the measured cross-section versus the LET.

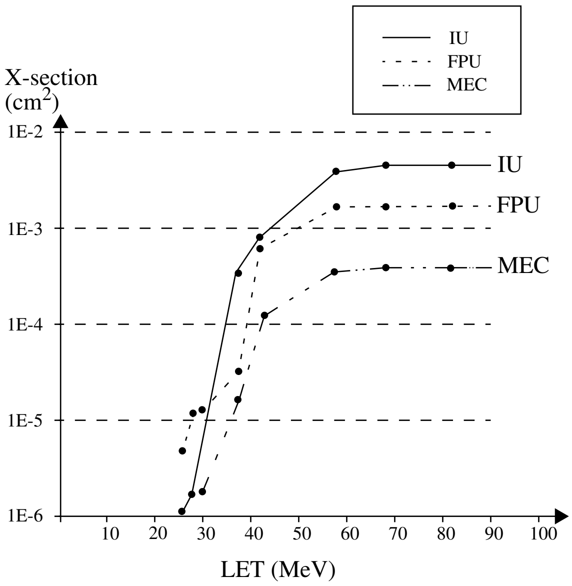

Figure 2: SEU sensitivity for ERC32

Tables 1 to 6 show the obtained error distribution in both absolute and relative terms for the three devices. The errors can be categorized into three groups:

Device level detected. The error is detected on-chip by parity checkers or other mechanisms (e.g. illegal instruction trap). These errors include *recoverable errors*, *register file errors* and *other errors* in the tables.

System level detected. The error is detected at system level through external bus parity errors, watchdog time-out or memory access errors.

Undetected. The error is not detected by any error-detection mechanism, but signalled through a checker device interrupt (IU and FPU) or program failure (MEC).

The error rate during testing varied with the ion flux and LET and was between 0.001 to 1Hz. Limiting the maximum error rate to 1 Hz guaranteed that if an error ever would become effective it would also be reported, since the used test programs had a cycle time of less than 100 ms.

4.1 IU test results

Table 1 and 2 show the obtained results for the IU using the *iutest* program. This program exercises most of the IU functionality including 94% of the register file.

Test	Restart. errors	Regfile errors	Other errors	System errors	Undetec errors	Total errors	LET	X-sect (cm^2)
Ni-58, 0°	9	6	2	8	4	29	26.6	1.2E-5
Ni-58, 15°	4	1	0	0	0	5	28.5	2.1E-5
Ni-58, 30°	3	5	3	4	4	19	30.7	-
Br-79, 0°	3	402	14	2	4	425	37.2	4.3E-4
N-58, 45°	1	96	2	1	1	101	37.6	-
Br-79, 30°	2	343	15	5	2	367	42.9	9.9E-4
I-127, 0°	36	1,294	183	18	10	1,541	58.8	3.7E-3
I-127, 30°	21	833	140	10	2	1,006	67.9	4.2E-3
Au-197, 0°	9	435	61	8	1	514	82.3	4.3E-3
TOTAL	**88**	**3,415**	**420**	**56**	**28**	**4,007**		

Table 1: IU error distribution

Test	Restart. errors	Regfile errors	Other errors	System errors	Detected errors (IU)	Detected err (total)	Ratio undet
Ni-58, 0°	31%	21%	7%	28%	59%	86%	14%
Ni-58, 15°	80%	20%	0	0	100%	100%	0%
Ni-58, 30°	16%	26%	16%	21%	58%	79%	21%
Br-79, 0°	0.7%	94.6%	3.3%	0.5%	98.6%	99.1%	0.9%
N-58, 45°	1%	95%	2%	1%	98%	99%	1%
Br-79, 30°	0.5%	93.5%	4.1%	1.4%	98.1%	99.5%	0.5%
I-127, 0°	2.3%	84%	11.9%	1.2%	98.2%	99.4%	0.6%
I-127, 30°	2.1%	82.8%	13.9%	1%	98.8%	99.8%	0.2%
Au-197, 0°	1.8%	84.6%	11.9	1.6%	98.2%	99.8%	0.2%
AVERAGE	**2.1%**	**85.2%**	**10.5%**	**1.4%**	**97.9%**	**99.3%**	**0.7%**

Table 2: IU relative error distribution

The results from the IU tests show that the error-detection mechanisms are as efficient as projected; 97.9% of all errors are detected directly by the IU and 99.3% are detected in the system. It can further be noted that the register file errors account for more than 85% of all errors, and only 2% of the errors are restartable without software intervention.

4.2 FPU test results

Tables 3 and 4 show the obtained results for the FPU using the *fputest* program. This program exercises most of the FPU functionality including 100% of the register file.

Test	Restart. errors	Regfile errors	Data bus errors	System errors	Undetec. errors	Total errors	Ratio undetec.
Ni-58, 0°	9	7	0	0	1	17	6.2%
Ni-58, 15°	5	7	2	1	1	16	6.2%
Ni-58, 30°	27	6	0	1	7	41	17%
Br-79, 0°	20	148	1	1	10	180	5.6%
Br-79, 30°	26	211	1	0	12	249	4.8%
I-127, 0°	190	866	5	0	64	1,125	5.7%
i-127, 30°	58	410	1	0	30	499	6.0%
Au-197, 0°	37	178	3	0	10	228	4.4%
TOTAL	**372**	**1,833**	**13**	**3**	**135**	**2,355**	**5.7%**

Table 3: FPU absolute error distribution

Test	Resta. errors	Regfile errors	Data bus errors	System errors	Detected err. (FPU)	Detected err. (total)	Ratio undetec.	LET	X-sect (cm^2)
Ni-58, 0°	53%	41%	0	0	94%	94%	6.2%	26.6	5.9E-6
Ni-58, 15°	31%	44%	12%	6%	87%	93.8%	6.2%	28.5	1.3E-5
Ni-58, 30°	66%	15%	0	2.4%	80%	83%	17%	30.7	1.4E-5
Br-79, 0°	11%	82%	0.6%	0.6%	93.9%	94.4%	5.6%	37.2	3.8E-4
Br-79, 30°	10%	85%	0.4%	0	95.5%	95.1%	4.8%	42.9	7.1E-4
I-127, 0°	17%	77%	0.4%	0	94.3%	94.3%	5.7%	58.8	1.9E-3
i-127, 30°	12%	82%	0.2%	0	94%	94%	6.0%	67.9	1.8E-3
Au-197, 0°	16%	78%	1.3%	0	95.6%	95.6%	4.4%	82.3	1.9E-3
TOTAL	**15.8%**	**77.8%**	**0.5%**	**0.1%**	**94.1%**	**94.2%**	**5.8%**		

Table 4: FPU relative error distribution

Table 4 shows that the FPU error-detection coverage is about 94%, somewhat lower than expected. Register file errors account for 77.8% of all errors while 15.8% of all errors are restartable without software intervention. The reason for the low error-detection coverage is further discussed below.

4.3 MEC test results

Tables 5 and 6 show the MEC test results using the well known *paranoia* program. The *error manager* column correspond to MEC register parity errors detected by the MEC error manager. The *system error* column reflects watchdog time-out, unexpected traps and spurious resets. *Undetected errors* indicate program failure, i.e. *paranoia* error.

Test	Error manager	System errors	Detected errors	Undetec. errors	Total errors	Ratio undetec.
Ni-58, 0°	1	6	7	1	8	12%
Ni-58, 30°	0	8	8	0	8	0%
Br-79, 0°	6	1	7	0	7	0%
Br-79,3 0°	19	3	22	0	22	0%
I-127, 0°	92	15	107	0	107	0%
I-127,3 0°	86	11	97	2	99	2%
Au-197, 0°	72	11	83	0	83	0%
TOTAL	**276**	**55**	**331**	**3**	**334**	**0.9%**

Table 5: MEC absolute error distribution

Test	Error manager	System errors	Detected errors	Undetected errors	LET	X-sect (cm^2)
Ni-58, 0°	12.5%	75%	87.5%	12.5%	26.6	7.1E-7
Ni-58, 30°	0	100%	100%	0%	30.7	2.5E-6
Br-79, 0°	85.7%	14.3%	100%	0%	37.2	1.8E-5
Br-79,3 0°	86.3%	13.6%	100%	0%	42.9	1.1E-4
I-127, 0°	86.0%	14.0%	100%	0%	58.8	4.2E-4
I-127,3 0°	86.9%	11.1%	98.0%	2.0%	67.9	4.9E-4
Au-197, 0°	86.7%	13.2%	100%	0%	82.3	4.7E-4
TOTAL	**82.6%**	**16.5%**	**99.1%**	**0.9%**		

Table 6: MEC relative error distribution

Table 6 shows that the MEC overall error-detection coverage is 99.1%. About 82% of the errors were detected directly by the error manager and originated from register parity errors, while 16% were detected through secondary effects. The MEC registers are continuously checked for parity errors by hardware, and the test coverage is therefore independent of the application. The registration of undetected errors rely on the self-checking paranoia program - if the program reports an error without prior detection of a hardware error then an undetected error is assumed. The detection method is obviously not ideal since *paranoia* is not 100% self-checking, and the MEC results must therefore be taken as indicative.

4.4 Discussion of test results

The overall test results are considered to be very promising; 6,530 of 6,696 injected errors were detected corresponding to a system level error-detection coverage of 97.5%. By studying the error distribution at different LET levels, it is apparent that for LET < 37 MeV, register parity errors are less frequent than at LET levels above 37 MeV. For the IU, only 23% of the errors below 37 Mev occur in the register file while 85% of all errors above 37 MeV are register file errors. The ratio of restartable errors, i.e. errors in the first three pipeline stages, is 31% below 37 MeV but only 2% above. The ratio of undetected errors and errors detected at system level, i.e. not on-chip, is also significantly higher below 37 MeV. The FPU and MEC show similar behaviour; the ratio for FPU register file errors is 27% below 37 MeV but 78% above. The restartable error ratio in the FPU is 55% below 37 MeV but only 14% above. For the MEC, only 6% (!) of the errors below 37 MeV are detected through register parity checking, compared to 82% above. Figure 3 shows errors in registers compared to the effective LET.

The FPU results also show a higher ratio of undetected errors even at high LET levels (5.8%) than the IU and MEC.

One way to explain the above behaviour is assume that the LET threshold for SEU errors in combinational logic is lower than the LET threshold for registers. At low LET levels (< 37 MeV), the dominant SEU effect will be upsets in combinational logic, for which the parity-based detection scheme is less effective. At higher LET levels (> 37 MeV), the SEU sensitive cross-section for registers will become significant and SEU errors in registers will become the dominant error source. This assumption is supported by following observations:

- The majority of combinational logic in the IU is in the three first pipeline stages, where detected errors are restartable.
- The FPU has significantly more combinational logic than the other two devices, explaining the higher ratio of undetected errors at high LET levels.
- The majority of the combinational logic in the FPU is in the datapath where detected errors are restartable.
- The low on-chip error-detection coverage below 37 MeV in the MEC.
- The register cell used in all three devices is hardened against SEU errors by design and has shown a threshold LET of about 35 MeV in earlier tests and simulations [4].

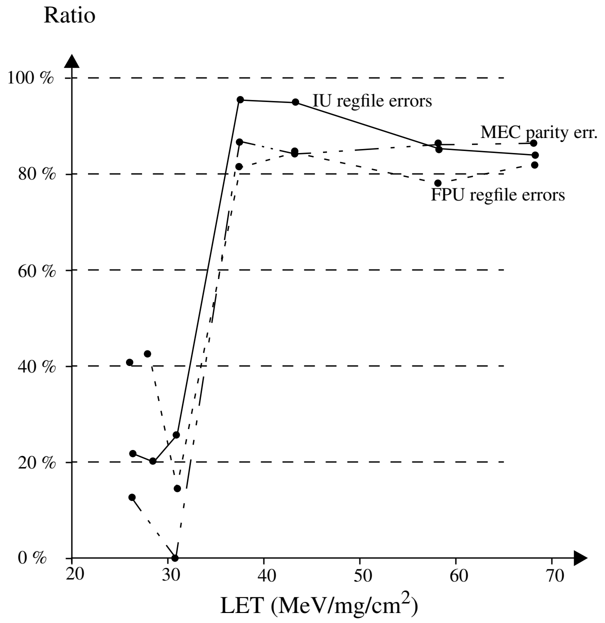

Figure 3: Ratio register errors of all errors

Whether an SEU error originating from combinational logic is detected by register parity depends on the length of the induced pulse and the delay through the parity generator. Typically, parity for a certain data is generated by a parity tree and latched on the same clock edge as the corresponding data (fig. 4a). If the delay through the parity generator is longer then the pulse length of the SEU error, the error will be detected since the latched parity bit will never reflect the erroneous data at the input of the latch (fig. 4b). If the delay in the parity generator is shorter then the SEU pulse length, there is a chance that the erroneous data and erroneous parity bit are simultaneously latched and the error will not be detected (fig. 4c). This could explain the low error-detection coverage in the MEC below 37 MeV - the parity in the MEC is mostly built over a few bits only, resulting in a short delay through the parity generator and therefore a higher probability for undetected parity errors. Increasing the delay by adding logic is not desirable since the cycle time would increase, thereby limiting the performance.

There is also the possibility that one SEU pulse splits into several pulses creating a "shower" of errors at the input of a register - this can happen if a complex structure such as a barrel shifter or multiplier is hit. In this case, it is impossible to predict whether the error will be detected or not.

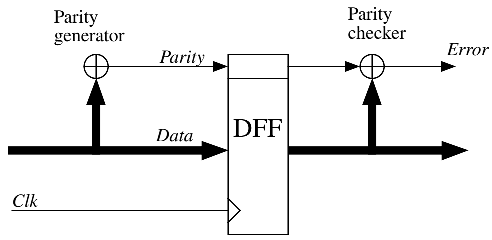

4a. General parity protected register

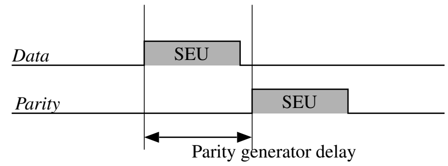

4b. Parity generation delay longer than SEU pulse

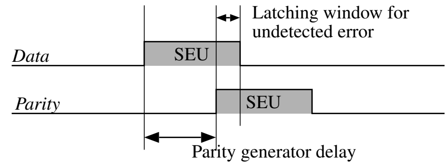

4c. Parity generation delay shorter than SEU pulse

5 Absolute error rates

Error calculations for the three devices has been done using CREME [3] software from Naval Research Laboratory. The results are worst-case results assuming solar minimum conditions and a minimum shielding of 1 g/cm^2, but exclude protons and solar flares. The upset rate has been calculated for two orbits:

- Geostationary orbit (GEO), altitude 35,786 km. No geomagnetic shielding - also applies to eccentric orbits.
- Low earth orbit (LEO), altitude 400 km, inclination 53^o (space station). Geomagnetic shielding.

Other orbits like Highly Elliptical Orbit (HEO) or Geostationary Transfer Orbit (GTO) can be expected to give results very similar (but slightly lower) to GEO. The table 7 and 8 below show the upset rates and corresponding FITS (failures in 10^9 hours) figures. The results show that in LEO, the SEU error-rate for the ERC32 core is approximately one error in 17 years. Taking into account the error-detection mechanisms, the rate for undetected errors is one in 732 years. For GEO, the figures are 2.1 and 87 years respectively.

Device	LEO (upset/day)	FITS (all err.)	Error-det. coverage	FITS (undet. err.)	MTBF (all err.) (years)	MTBF (undet. err.) (years)
IU	9.9E-5	4,125	99.3%	29	27.7	3,953
FPU	5.1E-5	2,125	94.2%	123	53.7	926
MEC	9.3E-6	388	99.1%	4	295	32,373
ALL	**1.6E-4**	**6,638**		**156**	**17.1**	**732**

Table 7: Error rates in LEO

Device	GEO (upset/day)	FITS (all err.)	Error-det. coverage	FITS (undet. err.)	MTBF (all err.) (years)	MTBF (undet. err.) (years)
IU	8.3E-4	34,583	99.3%	242	3.3	471
FPU	4.3E-4	17,917	94.2%	1,039	6.4	110
MEC	8.1E-5	3,375	99.1%	30	33.8	3,758
ALL	**1.3E-3**			**1,311**	**2.1**	**87**

Table 8: Error rates in GEO

It should be noted that the above error-rates are absolutely worst-case figures which will not occur in real applications. The test programs used to obtain the error-rates continuously uses 94% of all IU registers and 100% of all FPU registers. Analysis from existing programs have shown that only 5 of the 8 IU register windows are used on average, and only half of the registers in each window contains 'live' data. This means that the error-rate coming from register-file errors would drop with 69% since errors are only detected when the register value is used. From table 2, it can be seen that register-file errors contribute to 82-95% of all IU errors. It is then a valid assumption that the error-rate for detected errors would be about 50% of the worst-case error-rate. The error-rate for undetected IU errors is not likely to change, since the undetected errors do not originate from hits in the register-file, but rather in unprotected state registers and combinational logic. Some reduction in the error rate would probably also come from the fact that the utilisation of the processor is seldom 100%, and the idle task is likely to put the processor in power-down mode, i.e. shut down the clock. During this time, hits in combinational logic would not lead to any errors.

Similar analysis show that even rather floating-point intensive applications only use the FPU 10-20% of the time, with an average of 8 'live' registers. From table 7, it can be seen that FPU register-file errors contribute to about 80% of the error-rate and that 15% of all errors are recoverable without software intervention. With these assumptions the actual FPU rate for detected errors will be less than 5% of the worst-case figure. The rate for undetected errors in not affected by the register usage, but by the FPU utilisation. The rate for undetected errors is the proportional to the FPU utilisation, here assumed to 15% of the worst-case figure. The error rate in the MEC is affected along similar lines as the FPU, but since the error-rate is so small, this effect can be neglected. The tables 9 and 10 show the predicted error-rates adjusted to an average application as described above. As reference, the failure rate for permanent failure for the whole chipset according MIL-STD-217E is 250 FITS (space environment).

Device	LEO (upset/day)	FITS (all err.)	Error-det. coverage	FITS (undet. err.)	MTBF (all err.) (years)	MTBF (undet. err.) (years)
IU	4.4-E-5	2,062	99.3%	11	55.4	7.906
FPU	2.5 E-6	106	94.2%	6	1,074	18,520
MEC	9.3E-6	388	99.1%	4	295	32,373
ALL	**1.6E-4**	**2,556**		**21**	**44.7**	**5,436**

Table 9: Error rates in LEO (compensated)

Device	GEO (upset/day)	FITS (all err.)	Error-det. coverage	FITS (undet. err.)	MTBF (all err.) (years)	MTBF (undet. err.) (years)
IU	4.1E-4	17,292	99.3%	121	6.6	942
FPU	2.2E-5	896	94.2%	156	127	732
MEC	8.1E-5	3,375	99.1%	30	33.8	3,758
ALL	**1.3E-3**	**21,563**		**307**	**5.3**	**372**

Table 10: Error rates in GEO (compensated)

6 Conclusions

The initial tests of the ERC32 have shown that parity-based error-detection mechanisms are effective to detect SEU errors in VLSI devices. The error-detection coverage ranged from 94% (FPU) to 99% (IU and MEC). The majority of the errors occurred in the register files (~80%). The error-rates for various orbits were calculated, and the usage of parity based error-detection resulted in suppression of the MTBF for undetected SEU errors from 45 to 5,436 year in low-earth orbit and from 5 to 372 years in geostationary orbit.

It is planned to further characterize the effects of SEUs in combinational logic with new tests in 1997, where variation in clock frequency will be used to differentiate between combinational and sequential SEUs. An attempt to use control flow checking will also be done to investigate whether increased detection coverage can be obtained.

References

[1] "Concurrent error-detection and modular fault-tolerance in a 32-bit Processing Core for Embedded Space Flight Applications", Jiri Gaisler ESA/ESTEC, FTCS-24

[2] "On latching probability of particles induced transients in combinatorial networks", Peter Linden et al., FTCS-24

[3] "Cosmic Ray Effects on Microelectronics", Part IV, J.H.Adams, NRL Memorandum Report 5901, 1986

[4] "Detailed design of ERC32 integer and floating-point unit", MHS internal document AMS/LET/0001/JMSP, July 24, 1994

[5] "Techniques of microprocessor testing and SEU rate prediction", R.Koga et al., IEEE Trans Nucl. Sci., NS-32, 1985