

SEE Characterization of the New RTAX-DSP (RTAX-D) Antifuse-Based FPGA

Sana Rezgui, *Member, IEEE*, Paul Louris, and Rezwana Sharmin

Abstract—A comprehensive SEE characterization at high-frequencies (up to 120 MHz) of the new space-flight RTAX-D antifuse-based FPGA family is presented. SEE hardening-by-design techniques in the main FPGA programmable architectures have been implemented. It is evaluated in-beam to show their efficacy in mitigating SETs with little area and time penalty. In particular, the FPGA core and the new embedded DSP blocks were mitigated and tested in heavy-ion beam. Comparing to its predecessor, RTAX-S, these mitigations reduce the overall orbital error-rates by an order of magnitude.

Index Terms—SET Characterization and Mitigation, Hardening-By-Design, antifuse-based FPGAs, non-volatile FPGAs, DSPs, radiation tests, heavy-ions.

I. INTRODUCTION

Based on the existing ACTEL RTAX-S 0.15- μ m Space-Flight Field Programmable Gate Array (FPGA) family [1], the new RTAX-DSP FPGA product family [2] has been developed, using the same processing technology, to meet the increasing performance demand of data processing. Based on the design and SEE data of the RTAX-S, DSP blocks have been added with minimum design changes, and mitigations have been implemented to enhance the SEE tolerance of the RTAX-DSP (called also RTAX-D).

Indeed, the RTAX-S had been hardened to tolerate Single Event Upsets (SEU) by a hard-wired, asynchronous Triple Module Redundancy (TMR) scheme on its register-cells (R-Cell); but there is no embedded hardwired SET-mitigation solution for its combinational-cells (C-Cell). Based on this knowledge, the first published data, with the clock-frequency ranging up to 150 MHz [3], attributed the frequency-dependent SET-induced errors to ion-strikes on C-Cells. However, further beam tests by Actel radiation group [4] and confirmed later by The Aerospace Corporation radiation group [5] showed that most of these errors are actually originating from ion-strikes on the R-Cells. Detailed analyses indicated the sensitive element to be the output buffer of the R-Cell [4].

These findings prompted adding several new mitigation solutions to the RTAX-D: the R-Cells are enhanced with additional SET mitigation on their output-buffer; the DSP blocks are SEU-hardened by TMR of all of its registers and SET-hardened by filtering technique [2]. The implemented

mitigation solutions are derived from techniques such as the guard-gate cells with temporal delay [4, 6-8] and triple-drive of the sensitive nodes.

This paper presents a comprehensive SEE characterization of the new RTAX4000D device, focusing on the DSP Blocks as well as the C-Cells and the R-Cells in the core. The SET cross-sections of the C-Cells were measured by using the previously published [4, 6] techniques for the RTAX-S and the Radiation Tolerant Flash-based FPGA (RT3P). The mitigation approaches as well as their efficacy in mitigating SET effects are demonstrated by beam tests and presented in this paper. The heavy-ion beam tests were performed in Lawrence Berkeley National Laboratories (LBNL) with a wide heavy-ion (HI) cocktail (Neon, Argon, Copper, Krypton and Xenon) at normal incidences. The projected orbital error-rates in geosynchronous orbits (GEO) are also compared with those of the RTAX-S to show the efficacy of the new employed mitigation techniques.

The paper is organized as follows. In section II, we will present a brief description of the RTAX-S and the RTAX-D features targeting mostly the FPGA core and the new embedded DSP blocks as well the details of the adopted mitigation techniques. In section III and IV, we will show and discuss the beam test results (SEE cross-sections and orbital error-rates) to evaluate the mitigation solutions.

II. Device Under-Test: RTAX-D

Actel RTAX-D architecture (shown in Fig. 1), derived from the RTAX-S architecture has been designed for high performance and total logic module utilization. Unlike traditional FPGAs, the entire floor of the RTAX-D, similarly to the RTAX-S device, is covered with a grid of logic modules. The RTAX-D/RTAX-S FPGA families use a metal-to-metal antifuse programmable interconnect element that resides between the upper two metal-layers. The antifuses are normally open circuit and, when programmed, form a permanent, passive, low impedance connection, leading to fast signal propagation [1-2].

A. FPGA Core: Functionality and Mitigation

Actel's RTAX-S/RTAX-D families provide two types of logic modules in the FPGA core: the R-Cell and the C-Cell. The C-Cell can implement more than 4,000 combinational functions of up to five inputs (Fig. 2) with carry logic for the implementation of arithmetic functions. While each SEU-hardened R-Cell appears as a single D-Type flip-flop to the user, each is implemented in silicon using triple redundancy.

Manuscript received July 16, 2010. Sana Rezgui, *Member, IEEE*, and Paul Louris are with Actel Corporation, Mountain View, CA 94043 USA (650-318-4928; fax: 650-318-2571; e-mail: sana.rezgui@actel.com).

Rezwana Sharmin was with Actel Corporation, Mountain View, CA 94043 USA. She is now with the University of Toronto, Toronto, ON Canada.




Fig. 1: Block Diagrams of the RTAX-S and the RTAX-D FPGAs

Each TMR R-Cell consists of three master-slave latch pairs (Fig. 3), each with asynchronous self-correcting feedback paths. The output of each latch on the master or slave side votes with the outputs of the other two latches on that side. If one of the three latches is struck by an ion and starts to change state, the voting with the other two latches prevents that change from feeding back and latching. Care was also taken in the layout to ensure that a single ion strike could not affect more than one latch.




Figure 2 – C & R-Cell Block Diagram Architecture

The R-Cell contains a flip-flop featuring asynchronous clear, asynchronous preset, and active-low enable control signals (Fig. 2). The R-Cell in both FPGAs is identical except that the one in the RTAX-D is enhanced with additional SET mitigation of the last output-buffer. For RTAX-D R-Cell, this output buffer is tripled to three buffers with their outputs connected to a single node. If one of the tripled buffers suffers an SET event, the other two will subdue this SET effect and maintain the correct signal. This

mitigation solution is called triple-drive. With it, the R-Cell is expected to be fully SEE mitigated.




Figure 3: R-Cell Implementation Using Voter Gate Logic

The clock source for the R-Cell can be chosen from the hardwired clocks, routed clocks, or internal logic. Two C-Cells, a single R-Cell, and two Transmit (TX) and two Receive (RX) routing buffers form a Cluster, while two Clusters comprise a SuperCluster (Fig. 1).

Each SuperCluster also contains an independent Buffer (B) module, which supports buffer insertion on high-fanout nets by the place-and-route tool, minimizing system delays while improving logic utilization. The logic modules within the SuperCluster are arranged so that two combinational modules are side-by-side, giving a C-C-R – C-C-R pattern to the SuperCluster. This C-C-R pattern enables efficient implementation (minimum delay) of two-bit carry logic for improved arithmetic performance.

B. Embedded DSP Blocks: Functionality and Mitigation

As shown in Fig. 4, the flexible elements of the RTAX-DSP Math-blocks enable easy integration into many signal-processing topologies, such as Fast Fourier Transforms (FFT), Inverse FFT, Finite/Infinite Impulse Response Filters, and Discrete Cosine Transforms [2]. They also enable the acceleration of high precision single/double floating point multiplications. To achieve uniform performance of DSP functions, the five DSP columns are evenly distributed across the device, each one of them adjacent to an SRAM/FIFO column. Each DSP-column includes 24 blocks of 18x18 bits DSP with a total of 120 DSP blocks per FPGA. Each block can be implemented as a simple multiplier, an adder/subtractor or an accumulator.

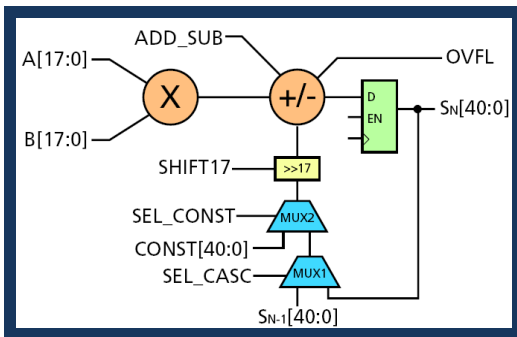


Fig. 4: Internal Architecture of the Embedded DSP Block

For SET mitigation of the DSP blocks, a low-pass SET filter has been inserted between each bit of the multiplier output and its following output-register (as shown in Fig. 5). This type of SET-mitigation is derived from a previously proposed SET filtering technique [6-8] and has the advantage of an easy implementation with minimum hardware overhead and time penalty if the delay time is short. The added SET filter first splits the original signal path to two with an additional path having an inverter string to delay the original signal. The latter and the delayed signal are both feeding to a guard-gate (GG). Fig. 5 shows this guard-gate of 4 transistors; it functions as an AND gate when the 2 input-signals agree, or as a latch of the previous state when the input signals differ. It will then pass only those transients with widths exceeding the inverter-string delay. Each GG is tripled to avoid SETs in it. In normal operation, any SET having a pulse width wider than the delay time will propagate to the R-Cell. More details about this SET mitigation solution are given in [6-8].




Fig. 5: Block Diagram of the DSP Mitigation: The output of each math-block is connected to an SET filter, whose filtering strength is controlled by a delay (an inverter-string). The guard-gate (GG) is tripled to avoid mitigate SETs in it.

III. SEE Characterization of the RTAX-D FPGA Core

The test targets three configurable architectures in the RTAX-D FPGA, the FPGA core C-Cells and R-Cells, as well as the embedded DSP blocks in three different configurations. Table 1 shows the features of the RTAX4000D and the RTAX2000S, to facilitate the comparison between the beam test results in the following.

Table 1: Features of the RTAX2000S and RTAX4000D

Part	RTAX2000S	RTAX4000D
System Gates	2M	4M
R-Cells	10,752	18,480
C-Cells	20,504	36,960
RAM Kbits	288	540
DSP (18x18) Blocks	-	120
Secure (AES) ISP	Yes	Yes
Clocks (Hardwired & Routed)	4 H and 4 R	4 H and 4 R
I/O Banks	8	8
User IOs	684	840

A. Experimental Test Setup

An updated version of the test setup built previously for the RTAX-S [5] was used. It includes three boards: 1) a motherboard housing a master FPGA that does the in-beam monitoring and controlling of the DUT operation, 2) a new RTAX-D daughterboard housing the DUT that provides the interfacing to the motherboard and 3) an interface board housing a slave FPGA for the communication between a central PC and the motherboard through two USB ports.

The master FPGA placed in the motherboard is an A3P1500-FG484, while the slave FPGA in the interface board is an A3P1000-PQ208. The daughterboard includes a socket for the RTAX4000D-CQ352 DUT. Considering the high number of IOs (94) connecting the DUT to the master FPGA, many sub-design versions were implemented and exercised simultaneously on the same DUT with no interferences between them. This feature is very important for testing different designs in the same conditions. Fig. 6 shows the RTAX-D experimental radiation test setup.




Fig. 6: RTAX-D Radiation Test-Setup

For communication with the host PC, a generic user interface was designed to communicate with the slave board. The communication protocol between the slave board and the host PC always remains the same for easy and fast implementation of any new SEE test experiment, always with a maximum of 64 display counters, whose functions are adjustable according to the running experiments. These counters are usually used for displaying the number of SEE events among other indicators of the operation of the DUT design. More details about this user interface are given in [4]. For beam test experiments, two designs have been built for the testing of: 1) the FPGA core and 2) the DSP blocks.

B. FPGA Core SEE Characterization (nC-R-Cells)

The purpose of this test is to determine the SEE cross-section of the combination of a number of C-Cells-string whose output is connected to an R-Cell. This type of implementation is named in this paper an nC-R-Cell. In the following, its cross-section in an RTAX-D FPGA will be measured and compared to that in an RTAX-S FPGA.

As mentioned above, the R-Cell should be relatively SEU immune because of the TMR implementation in it. However, for the RTAX-S, since the output-buffer of each R-Cell was not TMR'd, each R-Cell in the RTAX-S FPGA was sensitive to SETs and its cross-section increases with the clock frequency, as previously shown in [4-5].

To measure the SEE cross-section of an nC-R-Cell, a basic test design including two shift registers (SR) was implemented on the RTAX4000D-5Columns. As shown in Fig. 7, SRL4 is a shift register of 302 R-Cells configured as D-Flip-Flops (DFFs) with one global clock signal and where four C-Cells configured as inverters (C-Cell-Inverter) are inserted between adjacent R-Cells. SRL10 is identical to SRL4 except that ten C-Cell-Inverters are inserted this time between each two R-Cells.




Fig. 7: Block Diagram of the nC-R-Cells Test Design Implemented on an RTAX4000D

The purpose of this implementation is to show the dependence of the clock frequency and the number of logic C-Cells inserted between two R-Cells on the resulting SEE cross-sections. SRL4 and SRL10 were tested at 15 and 30 MHz, although their maximum frequencies are 120MHz for SRL4 and 60 MHz for SRL10. Fig. 8 shows the SEE cross-sections obtained for these sub-designs and their dependence on the design's frequency as well as the number of inserted C-Cells between each two R-Cells.

Few upsets were observed for the SRL4 and SRL10 running the clock frequency of 15 or 30 MHz. At 30 MHz, their SEE cross-sections are twice the SEE cross-sections at 15 MHz. From the linear dependency on the clock frequency of these cross-section data, we concluded that most of the errors are due to SET events. Also, at the same frequency (15 or 30 MHz), the data shows that the SEE cross-section of a 10C-R-Cell is twice the SEE cross-section of a 4C-R-Cell, suggesting that it is due to the increased number of C-Cells in the SRL10 test design compared to the SRL4.

To estimate the worst-case for SRL10, whose highest frequency is 60MHz, we multiplied its SET cross-section obtained at 30 MHz by two. Similarly, the SEE cross-section of a 4C-R-Cell, whose maximum frequency is 120 MHz, was multiplied by four. Both worse-case SEE cross-sections are shown in Fig. 8, which coincide but do not exceed a C-Cell SET cross-section. Both show a saturation SET cross-section of $4.61 \times 10^{-8} \text{ cm}^2/(\text{nC-R-Cells})$, and a threshold LET (LET_{th}) of $28 \text{ MeV.cm}^2/\text{mg}$. Indeed, this obtained SET cross-section coincides at the saturation level with an RTAX-S C-Cell SET cross-section but not at the knee, specifically at the threshold LET. At this point, we conjecture that the difference in the threshold LETs is due to the added capacitance in the RTAX-D FPGAs routing increasing hence the SET filtering effects.

This also means that the R-Cell is SEE-mitigated and is not contributing anymore or very little to an nC-R-Cells cross-section and that the main component contributing to the overall SEE cross-section is one C-Cell. Therefore, once the R-Cells are fully SEE-mitigated, the highest SEE cross-section of a combinational shift-register equals any other one operated at its highest frequency. Furthermore, the highest SEE cross-section (worse-case) for a given design is the product of the number of used R-Cells in a given design and a C-Cell SET cross-section as demonstrated previously in Ref. 6. Also, based on this beam data, reducing the propagation delay between two R-Cells by half (by reducing the number of C-Cells) or running the design at half of its maximum allowed frequency, will lead to the same results.

It is then possible to estimate the SEE cross-sections of a given design running at its maximum frequency by knowing its critical data-paths and the frequencies of their clocks and their data-inputs. If a design's frequency or the input data-rate is reduced by two then the overall SET cross-section will be reduced by two. This can be very useful for the

estimation of the data-rate masking effects in a given design. Note that other calculations are needed to take in account the logic masking effects such as the case of an AND gate when one of its inputs is grounded at the time of an SET occurrence on the other input.




Fig. 8: SEE Cross-sections of the RTAX4000D nC-R-Cell

More importantly, at all frequencies, the SEE cross-section of an nC-R-Cell in the RTAX-D are much lower than that in the RTAX-S; and the LET_{th} in the RTAX4000D increases to 28 $MeV.cm^2/mg$, from 10 $MeV.cm^2/mg$ in the RTAX-S (Fig. 9). These results show the efficacy of the SET mitigation solution (triple-drive) implemented on the R-Cells output buffers. Note that we only observed soft errors and no power cycle were needed during the beam test.

The resulting WEIBULL parameters along with the predicted geosynchronous orbital error-rates for an nC-R-Cell implemented in an RTAX2000S or an RTAX4000D are given in Table 2. The data show a reduction of an order of magnitude of the orbital error-rate of an nC-R-Cell in an RTAX4000D compared to that in an RTAX2000S. This proves the efficacy of the added mitigation scheme for the R-Cell.

Table 2: WEIBULL Parameters, Threshold LETs and SET Cross-Sections of a Single nC-R-Cell and Estimated GEO Orbital-Error Rates in an RTAX4000D and the RTAX2000S Solar Minimum, Aluminum Thickness= 100mil; Z=2; no Funneling

Unit	Frequency [MHz]	S	W	Limit [$cm^2/Unit$]	Onset	LET_{th} [$MeV.cm^2/mg$]	Saturation Cross-Section [$cm^2/C-Cell$]	GEO Orbital Error Rate [Error/Cell/Day]
nC-R-Cell (RTAX2000S)	120	1.3	100	1E-7	10	$9.74 < LET_{th} < 21.17$	6.39E-8	5.21E-9
nC-R-Cell (RTAX4000D)	120	2	40	6E-8	23	$21.17 < LET_{th} < 30.86$	5.21E-8	4.19 E-10




Fig. 9: Efficacy of the Mitigation Solution in the RTAX-D

IV. SEE Characterization of the DSP Blocks

To select the appropriate filtering strength or delay for a DSP-block that optimizes the trade-off between the low orbital error-rates and the performance, a proto-type named RTAX4000D-5Columns in a CQ352 package was used for beam-test experiments. The DUT has five different DSP

columns; each column has a SET filter with a unique delay (from 0 to 1 ns), as shown in Fig. 10. This delay increments with 250ps from column 0 to column 4. For instance, column 0 doesn't include SET filtering, column 1 uses an SET filter of 250ps, while column 5 uses an SET filter of 1 ns. The main purpose of the beam experiments was to determine the appropriate delay.

For the beam test experiments, a DSP test design was implemented to exercise different implementations of the DSP blocks and measure their cross-sections. It uses five identical sub-designs; each one of them is implemented on a single DSP-column and uses 24 DSP blocks. These 24 DSP blocks are shared in three sets of eight DSPs each. As shown in Fig. 10, the first DSP set is configured as multipliers (MULT), the second set as adders/subtractors (ADD/SUB), and the third set as accumulators (ACC). These three different DSP-configurations use the math-block in different manners, as shown in Fig. 11. For instance, both of the adder and the accumulator use a feedback path in the DSP blocks, but not the multiplier.




Fig. 10: Block Diagram of the RTAX4000D-5Columns varying the levels of SET Mitigation with the Delay




Fig. 11: Block Diagram of the Three Tested DSP implementations: a) Multiplier (MULT), b) Adder/Subtractor (ADD/SUB), c) Accumulator (ACC)

A. SET Characterization of the DSP Blocks

Because of the high number of DSP blocks and their inputs/outputs, the math-blocks were tested in a built-in self test (BIST) manner. Indeed, supplied only by an external clock (HCLK), a counter cycles through the inputs A and B of two separate DSP-multipliers, whose outputs are compared at each clock-cycle. In-beam, each of these two DSP-blocks acts as a DSP-Under-Test or a Golden-DSP-block. Since both DSP blocks are placed far apart (more than 10 μm), the probability of simultaneous upsets in them within one clock cycle (8.5ns for 120MHz) is extremely low.

To achieve a reliable error-detection of SEEs on the DSP blocks, the comparison circuit (XOR gates, R-Cells, etc.) is tripled. This same sub-circuit (the set of two DSP blocks and their associated comparison circuit) is repeated four times vertically on a single column. The tripled logical sum of their four error-flags is sent to the master FPGA, which will issue an error if the three output-registers transition from ‘0’ to ‘1’ logic state. Fig. 12 shows the block diagram for the DSP-MULT test design. The same sub-design for the DSP-MULT test is replicated on the five DSP-columns to evaluate the SET mitigation for other filters with various delays.




Fig. 12: Block Diagram of the DSP-MULT Test Design

The same idea was applied to the DSP-ACC case, with a few additions. In this case, the counters’ outputs cycle through the multiplier’s inputs (A and B) as well as the other control signals (SUB, SLOAD_N and ARSHFT17). To initialize the accumulator, the SLOAD_N signal must be active so the SLOAD_DATA can be loaded into the DSP output-registers. This operation mode is called LOAD mode and SLOAD_DATA is a constant value that is alternating ‘0’ and ‘1’ logic states.

The activation of SLOAD_N allows the re-synchronization of the two DSP blocks at any time. When SLOAD_N is deactivated, each single ACC feeds back its resulting data (P[40:0]) and is then operated in an ACC mode. The SUB and ARSHFT17 signals are used to select between a subtraction, an addition or a shift of the loaded/accumulated data. If an SEE occurs on a DSP block operational at that time in ACC mode, this event might get locked and only a toggle of the SLOAD_N or ACLR signals can cancel this error. Therefore, during the beam test, two types of errors are expected for this DSP-ACC case: 1) a single error when operated in LOAD mode or 2) a multiple error when operated in ACC mode.

Type 2 errors are due to the locking of errors during the accumulator operation mode. The error-counter will stop incrementing as soon as the ACC is switched to operate in the LOAD mode. As the selection between these two operation-modes is automated by toggling the SLOAD_N through the counters, there is no need to manually reset or switch modes upon the detection of a multiple error. This speeds-up the test and minimizes the loss of beam time or underestimation of the cross-section of Type 2 errors as well as a clear differentiation between the two types of errors. The block diagram of this test design is depicted in Fig. 13.




Fig. 13: Block Diagram of the DSP-ACC Test Design

The same test design was also built for the DSP-ADD/SUB, except that the DSP blocks were feeding each other instead of being independent from each other as for the multipliers or the accumulators. As shown in Fig. 14, the eight DSP blocks are shared in this case in two sets. Each set is behaving as the DUT-set or the Golden-set and includes four DSP-ADD/SUB that are feeding each other, except for the first one. Each ADD/SUB is similar to an ACC in its functionality. In the case of an ADD/SUB DSP, SLOAD_N and SLOAD_DATA are named CDESEL and CIN respectively. The ADD/SUB DSP is initialized with CIN value when CDESEL is active. If CDESEL is deactivated, the DSP output-registers are fed with CDIN of the previous DSP block in the chain. CDIN of the first DSP block is at '0' logic state and is set automatically by the LIBERO software; this means that it is not accessible to the user.

The outputs of each two DSP blocks and their overflow signals are compared and the sum of their error flags is again transferred to the master FPGA, similarly to the case of the accumulator DSP tests. In this case, multiple errors are expected only if an SEE occurs on the first, second or third DSP in the chain. If it occurs on the first (second or third)

DSP, it will propagate to the three (two or one) consecutive DSPs and will cause the error-counter to jump by four (three or two). The probability of multiple errors is then equal to its probability of occurrence on the first three DSPs divided by the total number of DSPs in the chain which is four, also divided by two because of the toggling of the signal CDESEL. It should then be equal to 37.5% ($=100*3/(4*2)$).




Fig. 14: Block Diagram of the DSP-ADD/SUB Test Design

B. Heavy-Ions Beam Test Results

The test was done on three different DUTs in heavy-ion beams. The beam data showed SEE sensitivities in the three tested DSP configurations (MULT, ADD/SUB and ACC). The test results are presented per DSP-configuration and with the frequency. The data displayed in Fig. 15, is for the DSP-multiplier operated at 30, 60 and 120 MHz and shows a clear dependence on the frequency. Indeed, the SEE cross-sections of the DSP blocks increase almost linearly with the frequency. Moreover, no heavy-ion event caused a stuck at fault that caused the reset of the multipliers.

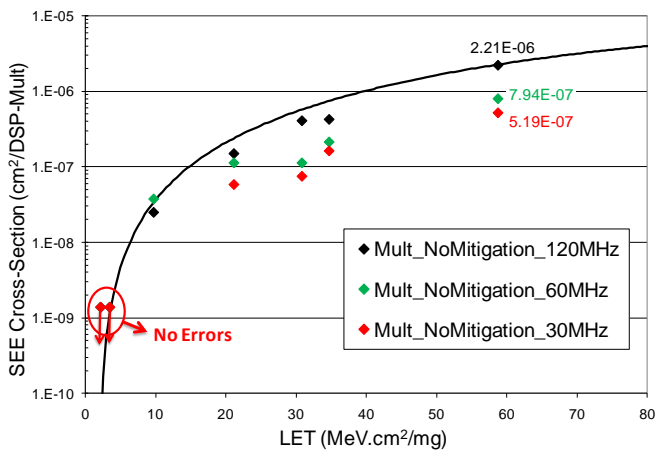


Fig. 15: SEE Cross-Sections for a DSP-Multiplier running at 30, 60 and 120 MHz

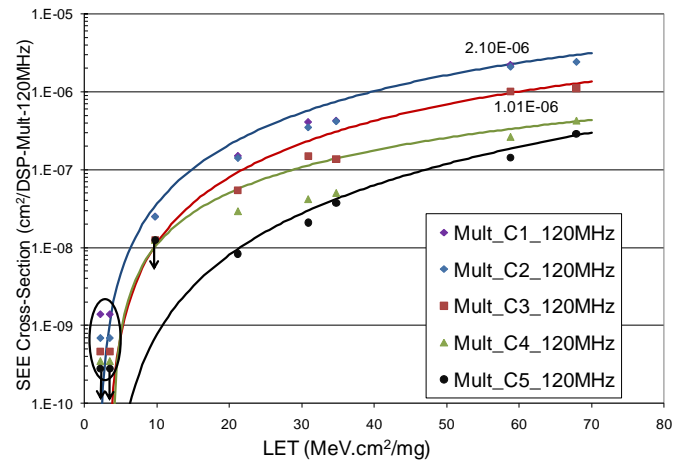


Fig. 16: SEE Cross-Sections for a DSP-Multiplier with different levels of SET Mitigation

Fig. 16 shows the SEE cross-sections of a DSP-MULT in the five columns, each with a different strength-level of SET Mitigation. Both of the 0 and 250 ps showed almost the same SEE cross-sections, while 750 ps and 1 ns showed also almost the same cross-sections. This means that most SETs are wider than 250 ns and shorter than 750 ps. Moreover, the 750ps delay appears to be efficient in reducing the DSP SET cross-sections. Note that because of the lack of statistics at the LET of 9.74 MeV.cm²/mg, the SEE cross-section with 750ps might have been overestimated, since all the other points are indicating lower cross-sections. More statistics are needed to confirm this statement.

Fig. 17 shows the SEE cross-sections of an ACC and an ADD/SUB DSP blocks both running at 80 MHz and a MULT running at 120 MHz. The beam data for the two other DSP-configurations (ACC and ADD/SUB) are similar but higher than in the case of the DSP-multipliers when compared at the same frequency. The difference in the cross-sections between an adder and a multiplier is most likely due to the added combinational logic for the feedback path. On the other side, the similarity in the measured cross-sections of the adder and the accumulator is most likely due to the same number of sensitive nodes in both of their circuits.

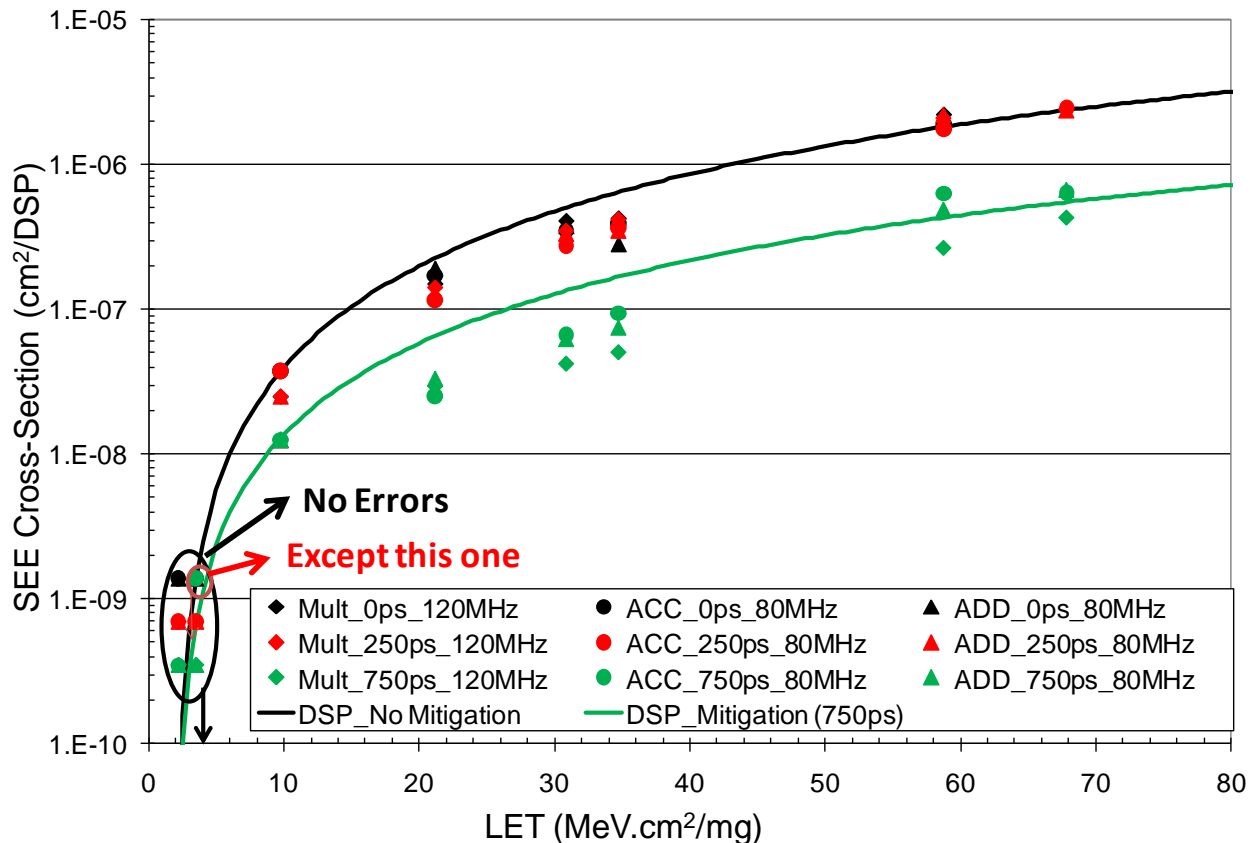


Fig. 17: SEE Cross-Sections for a DSP-MULT, a DSP-ACC and a DSP-ADD/SUB with different levels of SET Mitigation

