Characterisation of AMD Versal FPGA Transceivers Under Heavy-Ion Radiation

Alberto Gonzalez Villafranca, Albert Ferrer Florit, Marti Farras Casas, Steve Parkes, Pierre Maillard and Ken O'Neill

***Abstract*—The AMD Versal FPGA family includes the latest generation of radiation-tolerant FPGAs from AMD (XQRVC1902 and XQRVE2302), built on a 7-nm FinFET process. These provide unparalleled capabilities for space-qualified devices, featuring up to 44 integrated GTY high-speed transceivers that support lane speeds of 26.5 Gbit/s. These attributes make Versal ideal for implementing spacecraft communication protocols such as SpaceFibre, an advanced spacecraft on-board data-handling network technology that has been operating in space since 2021. This paper provides the results of a recent heavy-ion radiation campaign carried out in collaboration between STAR-Dundee and AMD. The campaign aimed to characterise the Versal transceivers under radiation and assessing the improvements in link reliability provided by the use of a robust protocol such as SpaceFibre. The results demonstrate that, as expected, different elements of the transceiver blocks are sensitive to radiation. Radiation-induced events disrupt the normal operation of communication protocols implemented using the high-speed transceivers. However, the results indicate that SpaceFibre automatically mitigates the majority of radiation-induced events affecting the transceivers, yielding a three order-of-magnitude reduction in the observed error rate, without requiring user intervention. Additionally, the campaign successfully demonstrated, for the first time, a 100 Gbit/s SpaceFibre link operating under radiation, with lanes running at 25 Gbit/s.**

***Index Terms*—FPGA, Heavy-Ion, Radiation Testing, Transceiver, Versal, SpaceFibre**

I. Introduction

The space-qualified AMD Versal XQRVC1902 and XQRVE2302 [1] are radiation-tolerant derivatives of the commercial SRAM-based Versal FPGA family. Manufactured using 7 nm FinFET technology, they provide a platform targeted at high-performance applications, offering 44 and 8 GTY transceivers respectively, each supporting data rates of up to 26.5 Gbit/s. An integrated configuration memory scrubbing mechanism (XilSEM), implemented with triplicated MicroBlaze processors in hardware, enables rapid repair of configuration memory (CRAM) upsets. Nevertheless, because the CRAM and programmable fabric are not inherently radiation-hardened, the use of Triple Modular Redundancy (TMR) may be necessary depending on the error-tolerance requirements of the application.

SpaceFibre (SpFi) [2] is a high-speed communication technology for use onboard spacecraft, supported by the European Space Agency (ESA) and standardised as ECSS-E-ST-50-11C in 2019. It enables both point-to-point and networked interconnections at multi-gigabit data rates while providing Quality of Service (QoS) and Fault Detection, Isolation, and Recovery (FDIR) capabilities. SpFi has been adopted by several spacecraft standards, including ADHA and SpaceVPX, and soon in SpaceVNX+ [3], and has been deployed in operational missions since at least 2021 (TRL-9). SpFi incorporates an error-recovery mechanism capable of automatically correcting both transient and persistent link errors, with recovery times for short transient faults of less than 3 μs. To further improve throughput and resilience, SpFi links support multi-lane operation, enabling the data stream of a single logical link to be distributed across multiple physical lanes. This approach provides inherent redundancy and, in the event of a lane failure, allows graceful degradation by redistributing traffic across the remaining lanes. Automatic reconfiguration of the link is typically completed within approximately 4 μs.

At the end of last year, STAR-Dundee and AMD completed a heavy-ion irradiation campaign on the Versal device. The primary objective of this campaign was to evaluate radiation effects on various FPGA elements, with particular emphasis on the high-speed transceivers, and to assess the mitigation provided by the SpFi protocol. The campaign characterised the radiation response of the transceiver blocks, including the determination of cross-sections for their internal components. The effectiveness of the XilSEM mechanism in correcting CRAM upsets was evaluated, together with its impact on link operation. The functionality of the Error Detection and Correction (EDAC) mechanism in the internal buffers was verified, along with the robustness of the clock and reset scheme and the reliability improvements achieved through distributed Triple Modular Redundancy (DTMR). Finally, the operation of the Fault Detection, Isolation, and Recovery (FDIR) mechanism was assessed to ensure rapid recovery from all Single Event Upsets (SEUs) causing transient errors, as well as to confirm that transceiver resets are automatically triggered

A. Gonzalez Villafranca is with STAR-Barcelona S.L., Spain (e-mail: alberto.gonzalez@star-barcelona.com).
A. Ferrer Florit is with STAR-Barcelona S.L., Spain (e-mail: albert.ferrer@star-barcelona.com).
M. Farras Casas is with STAR-Barcelona S.L., Spain (e-mail: marti.farras@star-barcelona.com).
S. Parkes is with STAR-Dundee Ltd., UK (e-mail: steve.parkes@star-dundee.com).
Pierre Maillard is with Adaptive Embedded Computing Group, AMD Inc. (e-mail: pierre.maillard@amd.com).
Ken O'Neill is with Adaptive Embedded Computing Group, AMD Inc. (e-mail: kenneth.oneill@amd.com).

and that the link reliably re-establishes following persistent failures.

II. Test Design

The VCK190 Evaluation Kit was selected as the test platform for the radiation campaign. This board incorporates a VC1902 Versal device, which is the commercial counterpart of the space-grade XQRVC1902. For the objectives of this campaign, both devices are regarded as functionally equivalent. Furthermore, the board offers a variety of high-speed interconnects (e.g., zSFP, zQSFP, FMC+) that allow the evaluation of the embedded GTY transceivers.

A. Test Architecture

Fig. 1 illustrates the architecture of the tested design. The block on the left (DUT) represents the portion of the design implemented within the Versal FPGA and exposed to radiation. The block on the right corresponds to an external STAR-Ultra PCIe unit, not subject to irradiation, which interfaced with the Versal through two SpFi links, I1 and I2. The STAR-Ultra PCIe provides two independent four-lane SpFi links via QSFP+ interfaces, with each lane supporting data rates of up to 7.8 Gbit/s. These interfaces can be used to transfer data at high speed to and from a host PC via PCIe Gen3. In this test configuration, I1 implemented a four-lane SpFi link operating at 25 Gbit/s and I2 a dual-lane link operating at 12.5 Gbit/s. Links I3 and I4 were looped back locally through an FMC+ card owing to external equipment limitations. Specifically, I3 implemented a four-lane SpFi link at 25 Gbit/s, whereas I4 implemented an experimental four-lane link operating at an aggregate data rate of 100 Gbit/s (25 Gbit/s per lane).

Internal data generators and checkers were implemented for each link, while an additional Test Monitor block was responsible for supervising and reporting the status of all links.

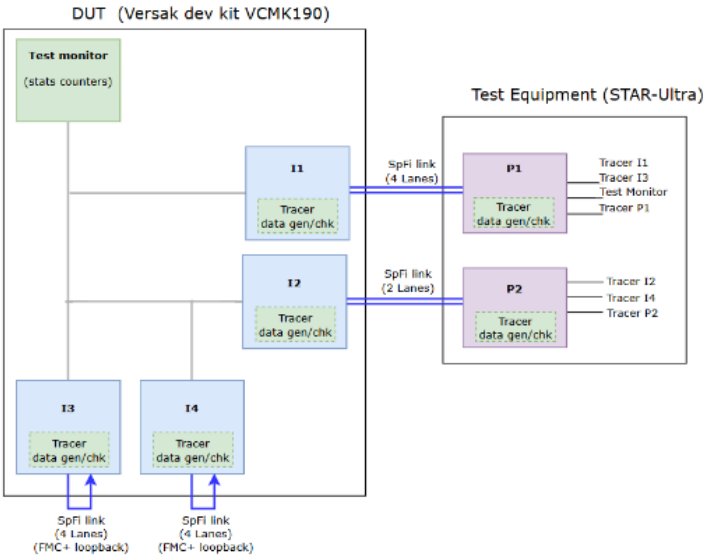

Fig. 1. Radiation test architecture.

B. Radiation Monitor

Two complementary STAR-Dundee tools were employed to maximise the value of the data collected, enabling classification of Single Event Effects (SEEs) by both source and effect. As these tools are not specific to SpFi or transceiver testing, they can be applied to the characterisation of other FPGA blocks as required.

One tool interfaced directly with the Test Monitor block shown in Fig. 1, providing event logging and recording of user interactions with sub-second resolution. System status was displayed to the operator in real time via a computer GUI, enabling continuous monitoring of the design and the detection of events requiring intervention, such as Single Event Functional Interrupts (SEFIs). In the event of a SEFI, the board was reprogrammed to restore normal functionality.

The second tool functioned as an embedded logic analyser, providing nanosecond-resolution measurements of relevant events within the design. The data processing chain incorporated an event classifier based on STAR-Dundee's internal verification tools, complemented by additional Python libraries specifically developed for the analysis of the radiation data collected.

III. Test Campaign

Testing was conducted in the autumn of 2024 at GANIL (France) and was funded by the RADNEXT European H2020 project [4]. The effective Linear Energy Transfer (LET) covered a range from 28 to 43 MeV·cm²/mg. Lower LET values could not be tested, as only Xenon ions (^{129}Xe with a 46^{+} charge state) were available during the two 8-hour irradiation windows. Consequently, the results reported here represent a worst-case heavy-ion radiation sensitivity measurement. An additional campaign is planned to investigate LETs below 25 MeV·cm²/mg to determine the onset threshold and intermediate cross-section values required for fitting Weibull curves, and estimating orbital error rates.

Fig. 2 shows the VCK190 board with an unlidded Versal device mounted on the support frame, ready for irradiation.

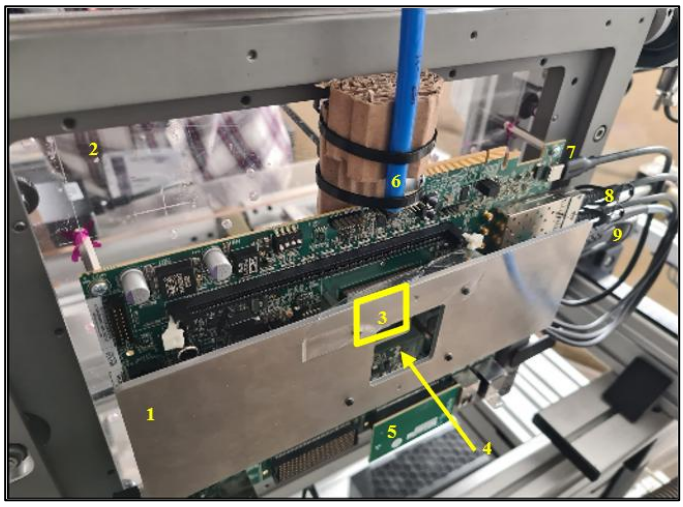

Fig. 2. Test Setup with the VCK190 board in place for radiation testing.

The key elements are described below:

1) Metal shield to protect other electronics in the board.
2) Lexan support board.
3) Versal VC1902 part.
4) Beam direction.
5) FMC+ loopback card.
6) Compressed air source used for cooling.
7) Programming cable.
8) 2x SFP+ connection to a two-lane SpFi link (I2).
9) QSFP+ connection to a four-lane SpFi link (I1).

IV. Results

Fig. 3 presents the aggregated cross-section for the four channels forming a transceiver Quad. This value accounts for multiple error cases affecting transceiver components, including the TX PLL and data path, and represents the cross-section experienced by a four-lane link employing such a Quad. The blue curve corresponds to measurements on Quads connected to external equipment, whereas the orange curve represents measurements on transceivers operating in physical loopback. The cross-section is approximately an order of magnitude higher when connected to external equipment than in loopback mode, highlighting the importance of testing with a representative setup. Green markers indicate SEFI events requiring board reprogramming for recovery. SpFi was able to successfully recover from all non-SEFI events affecting the transceiver without data loss. This represents approximately a three-orders-of-magnitude reduction in the Quad's observed error rate—difference between blue and green values.

Fig. 4 presents the cross-section for the individual elements comprising a transceiver Quad: the LCPLL (blue), the channel data path (orange), and the Quad shared logic (green). Events affecting the LCPLL or the channel data path trigger a SpFi retry, with recovery occurring in less than 5 µs, representing approximately 99 % of all events (see histogram of Fig. 5). In contrast, events affecting the Quad shared logic require a full transceiver reset lasting approximately 1.5 ms and account for about 1 % of events. Fig. 6 displays the histogram of transient error burst lengths induced by SEEs in the transceiver. The green bins correspond to error bursts that are correctable by the forward error correction (FEC) mechanism employed in high-speed commercial protocols such as Ethernet, Fibre Channel, and InfiniBand, but they represent only a small fraction (9%) of all recorded bursts. Consequently, although SpFi automatically recovers all recorded bursts, blue bins (91%) may still impact system performance when alternative protocols are used.

A protocol must be implemented in the FPGA fabric to enable data transmission through a transceiver. However, both the fabric and its configuration memory (CRAM) are susceptible to SEUs, introducing an additional challenge. Fig. 7 presents the cross-section of a four-lane SpFi link. Radiation-induced SEUs in the FPGA fabric degrade the operation of the SpFi IP, as indicated by the blue curve. The embedded XilSEM *scrubbing* mechanism recovers from the CRAM SEUs in about 15 ms, with the SpFi link self-recovering a few ms afterwards. However, in this case data loss may occur depending on the

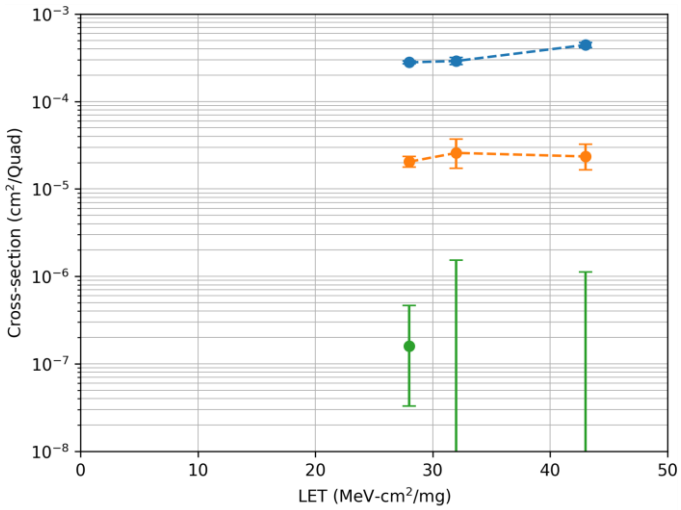

Fig. 3. Transceiver Quad cross-section. Blue: entire Quad. Orange: entire Quad operating in physical loopback. Green: Quad SEFI.

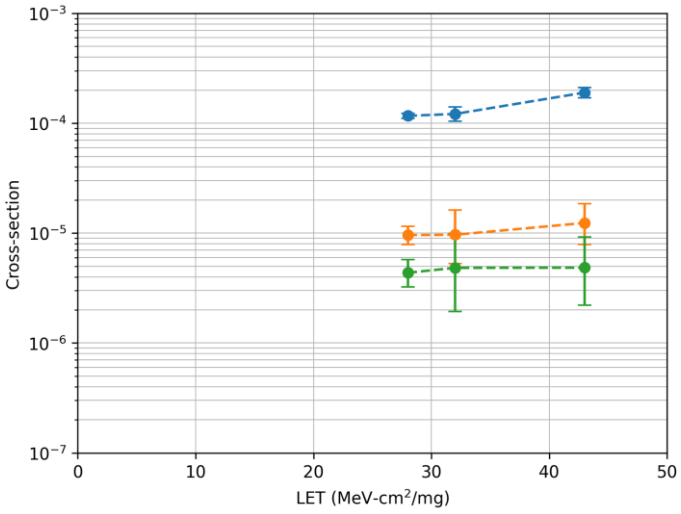

Fig. 4. Cross-section for the different elements of the transceiver Quad. Blue: Quad PLL. Orange: Channel data path. Green: Quad shared logic.

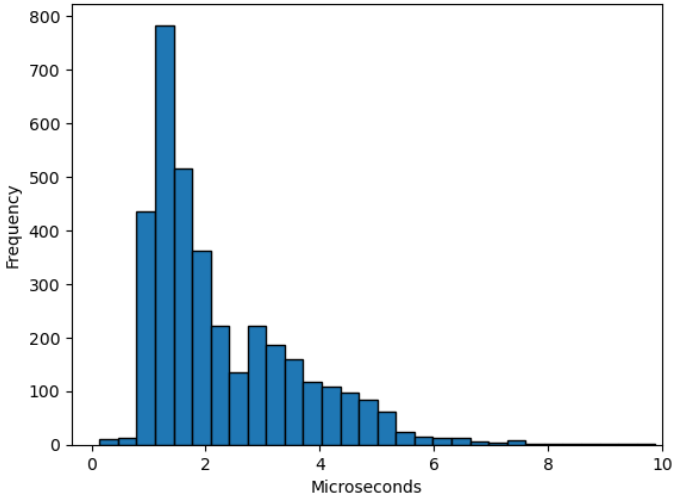

Fig. 5. SpaceFibre error recovery time histogram for transient SEEs.

specific logic part affected by the event. These cross-sections are approximately an order of magnitude lower than those of the transceiver Quad (blue curve in Fig. 3). This difference is partly attributable to the small footprint of the SpFi IP in the Versal device, which uses only about 0.5 % of the available fabric resources in an XQRVC1902, thereby reducing the susceptibility of the design to SEUs. The orange curve shows the cross-section for the same SpFi link with DTMR applied. This value remains to be confirmed, as only approximately 88 % of the link logic was covered by DTMR—closely matching the observed improvement factor. It is therefore reasonable to expect that a fully triplicated SpFi link would exhibit negligible susceptibility to SEUs in the FPGA fabric.

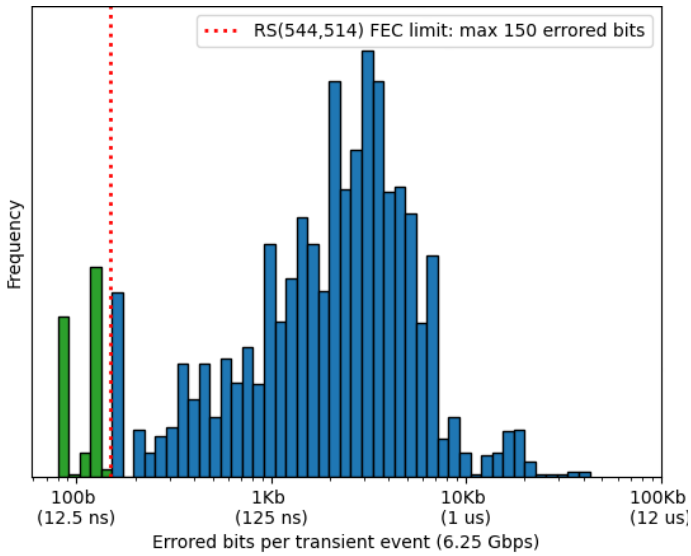

Fig. 6. Transient radiation error burst length distribution.

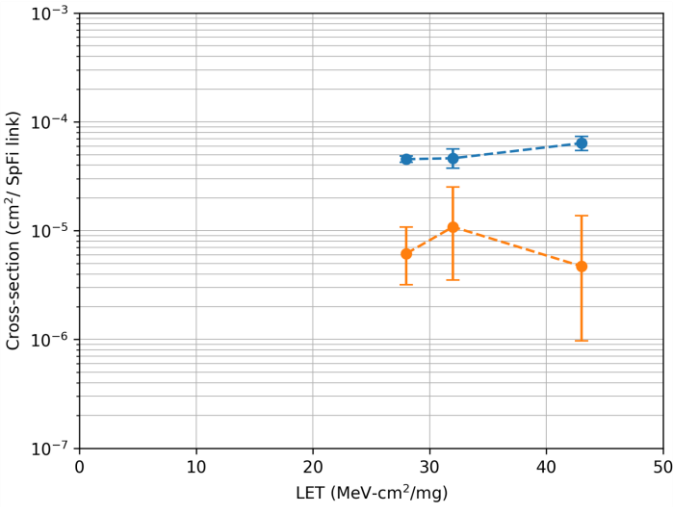

Fig. 7. Cross-section for SpaceFibre IP link fabric. Blue: nominal. Orange: DTMR applied to most (88%) of the link logic.

The cross-section of the SpFi link can be used to infer the configuration memory (CRAM) cross-section. Fig. 8 presents the inferred CRAM cross-section, superimposed with the values reported by AMD [5]. This value is inferred because direct measurement of the CRAM cross-section is extremely challenging. In this work, the value was obtained using the number of essential bits reported by Vivado for the design and estimating the number of essential bits associated with a single SpFi link based on the ratio of total design logic to the logic corresponding to one link. As reported in the literature [6][7], not all essential bits lead to errors when affected by an SEU. Only a fraction—between 9 % [6] and 15 % [7]—appear to be critical to the design. Using a central estimate of 12 % as reference, the CRAM SEU cross-section per bit is plotted in Fig. 8, showing good agreement with the values reported by AMD.

Additionally, the radiation campaign successfully demonstrated an experimental 100 Gbit/s SpFi link operating under radiation. This was achieved using a four-lane configuration, with each lane operating at 25 Gbit/s. The correct operation of the 100 Gbit/s link was validated under nominal conditions.

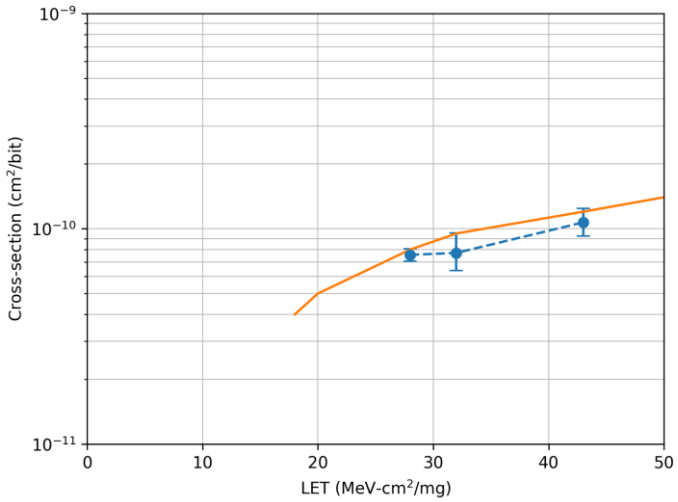

Fig. 8. Cross-section for the configuration RAM (CRAM) memory. Orange: Value reported by AMD [5]. Blue: Inferred value from the measured SpFi link cross-section.

V. Conclusion

A heavy-ion campaign with nanosecond-accuracy data collection was conducted to characterise the operation of Versal transceivers under radiation. This enabled precise measurement of error-burst lengths and identification of SEE sources within the transceiver. Results correspond to 28–43 MeV·cm²/mg. A follow-up campaign at lower LETs will refine onset and Weibull fits.

Cross-sections for the LCPLL, data path, and Quad shared logic were obtained. The analysis of the collected data confirmed that SpFi automatically recovers from all non-SEFI events affecting the transceivers in under 5 µs, without data loss. This represents approximately a three-orders-of-magnitude reduction in the Quad's observed error rate.

The analysis of fabric events showed that a four-lane SpFi link implemented in the FPGA fabric exhibits an order-of-magnitude lower SEU sensitivity than the Versal Quad transceiver used for the high-speed serial data transmission. These fabric SEUs can be mitigated using the FPGA's inbuilt scrubbing mechanism (XilSEM), which responds in

approximately 15 ms. Although data errors may occur, the SpFi link generally self-recovers, maintaining system functionality in most cases. Applying DTMR to the SpFi IP mitigates most fabric-induced upsets.

Additionally, the SpFi link cross-section was used to infer the CRAM SEU cross-section per bit. This estimation, based on the number of essential bits reported by Vivado and scaled to a single SpFi link, accounts for the fraction of essential bits considered critical (≈12 %). The resulting value aligns well with the CRAM cross-sections reported by AMD, supporting the validity of this indirect approach.

The campaign also successfully demonstrated a 100 Gbit/s SpFi link—implemented using four 25 Gbit/s lanes—operating under radiation, paving the way for future ≥200 Gbit/s implementations.

Overall, these results highlight the potential of the SpFi–Versal combination as a high-reliability solution for spacecraft data-handling systems, supporting the development of next-generation, high-throughput, radiation-tolerant communication architectures.

ACKNOWLEDGMENT

This activity received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 101008126, corresponding to the RADNEXT project.

The authors gratefully acknowledge the GANIL team for their support during the radiation test campaign, and Synopsys for providing a Synplify Elite license, which enabled the application of TMR to the SpaceFibre IP core.

REFERENCES

[1] AMD, "XQR Versal for Space 2.0 Applications Product Brief", AMD, 2023, available from https://www.xilinx.com/content/dam/xilinx/publications/product-briefs/xilinx-xqr-versal-product-brief.pdf (last visited 12th March 2025).

[2] ECSS Standard ECSS-E-ST-50-11C, "SpaceFibre – Very high-speed serial link", European Cooperation for Space Data Standardization, 15th May 2019, available from http://www.ecss.nl.

[3] S. Parkes et al., "SpaceFibre Onboard Interconnect: from Standard, through Demonstration, to Space Flight", IEEE Aerospace Conference, Big Sky, Montana, 2025.

[4] R. G. Alía et al., "Heavy Ion Energy Deposition and SEE Intercomparison Within the RADNEXT Irradiation Facility Network," in IEEE Transactions on Nuclear Science, vol. 70, no. 8, pp. 1596–1605, Aug. 2023, doi:10.1109/TNS.2023.3260309.

[5] Maillard, Pierre et al., "Protons and Heavy-ions Evaluation of 7nm Versal™ ACAP Programmable Logic (PL) and XilSEM for Space", Data Workshop Session, RADECS 2022, Venice.

[6] Luis Aranda et al., "Analysis of the Critical Bits of a RISC-V Processor Implemented in an SRAM-Based FPGA for Space Applications" in Electronics 9, no. 1: 175, 2020, https://doi.org/10.3390/electronics9010175.

[7] Ilya Tuzov et al., "Reversing FPGA Architectures for Speeding up Fault Injection: does it pay?", 2022, IEEE. 81-88. https://doi.org/10.1109/EDCC57035.2022.00023