

SpaceFibre Radiation Testing on Versal

Alberto Gonzalez Villafranca
STAR-Barcelona S.L.
Barcelona, Spain
alberto.gonzalez@star-dundee.com

Albert Ferrer Florit
STAR-Barcelona S.L.
Barcelona, Spain
albert.ferrer@star-dundee.com

Marti Farras Casas
STAR-Barcelona S.L.
Barcelona, Spain
marti.farras@star-dundee.com

Steve Parkes
STAR-Dundee Ltd.
Dundee, United Kingdom
steve.parkes@star-dundee.com

Pierre Maillard
Adaptive Embedded Computing Group
AMD Inc.
pierre.maillard@amd.com

Ken O'Neill
Adaptive Embedded Computing Group
AMD Inc.
kenneth.oneill@amd.com

Abstract—SpaceFibre (ECSS-E-ST-50-11C) is an advanced spacecraft on-board data-handling network technology. It builds upon its predecessor, SpaceWire (ECSS-E-ST-50-12C), to meet the increasing demands for higher data transfer rates and improved reliability in space applications. SpaceFibre allows links with different numbers of lanes to seamlessly interoperate and provides aggregate rates of 100 Gbit/s in existing space-qualified technology, targeting 200 Gbit/s in the short term. Consequently, this international open protocol has been integrated into numerous spacecraft standards such as ADHA, SpaceVPX and soon in SpaceVNX+.

STAR-Dundee has developed a comprehensive suite of SpaceFibre IP cores with optimized footprint and speed, specifically targeting space applications. These IP cores have achieved TRL-9, having been deployed in at least six operational missions since 2021 and currently being designed into more than 60 spacecraft. Support has been added for new radiation-tolerant FPGAs, including AMD Versal. Versal is the latest generation of radiation-tolerant FPGAs from AMD, built on a 7-nm FinFET process. It provides unparalleled capabilities on space-qualified devices, featuring up to 44 integrated GTY high-speed transceivers that support lane speeds of up to 25 Gbit/s. These attributes make Versal ideal for implementing advanced spacecraft communication protocols such as SpaceFibre.

This paper provides the results of a recent high-LET heavy-ion radiation campaign carried out in collaboration between STAR-Dundee and AMD. The campaign aimed to test the Versal transceivers and evaluate the improvements in link reliability achieved by the SpaceFibre protocol. Results demonstrate that SpaceFibre automatically mitigates most transceiver events, effectively reducing the error rates experienced by the user at these LETs by three orders of magnitude, without requiring user intervention. Such performance cannot be achieved with standard forward error correction techniques alone, such as Reed-Solomon codes. Furthermore, the results also demonstrate that applying distributed Triple Modular Redundancy to the SpaceFibre IP removes most single event effects affecting the FPGA fabric.

Additionally, the campaign successfully demonstrated for the first time a 100 Gbit/s SpaceFibre link operating under radiation. This was achieved using a quad-lane configuration, with each lane operating at 25 Gbit/s.

Keywords—SpaceFibre, SpFi, Versal, Transceiver, Radiation Testing, Heavy-Ion, FPGA

I. INTRODUCTION

SpaceFibre (SpFi) [1] is a communication technology for use onboard spacecraft which was instigated by the European Space Agency (ESA) and released as an ECSS open standard in 2019 (ECSS-E-ST-50-11C). It provides point-to-point and

networked interconnections at multi-Gigabit rates while offering Quality-of-Service (QoS) and Fault, Detection, Isolation and Recovery (FDIR) capabilities. SpFi has been integrated into numerous spacecraft standards including ADHA, SpaceVPX and soon in SpaceVNX+ [2].

SpFi implements an error recovery mechanism that automatically recovers from transient and persistent errors on the SpFi link, with typical recovery from transient errors in under 3 μ s. To enhance throughput and robustness, SpFi links can also operate as multi-lane, thus allowing data of a single logical link to be spread over several independent physical lanes. Multi-lane operation provides higher data rates through lane aggregation—supporting any number of lanes, up to 16—unidirectional operation, and hot and warm redundancy. Furthermore, when a lane fails, the multi-lane mechanism supports graceful degradation by automatically spreading traffic over the remaining working lanes, with automatic reconfiguration of the link requiring about 4 μ s.

The space-qualified AMD Versal XQRVC1902 and XQRVE2302 [3] are radiation-tolerant versions of the commercial SRAM-based Versal FPGA family. They are manufactured using a 7nm FinFET technology and provide a platform aiming at high performance applications, offering 44 or 8 GTY transceivers respectively (each capable of 25 Gbit/s). An internal scrubbing mechanism (XilSEM) has been implemented to quickly repair any configuration memory (CRAM) upset. Due to the CRAM and fabric not being specifically radiation-hardened, using a Triple Modular Redundancy (TMR) technique may be required depending on the application error tolerance requirement.

STAR-Dundee and AMD completed an initial heavy-ion test campaign with the Versal at the end of last year. The goal of the campaign was to evaluate the radiation effects on different elements of the FPGA, and to assess the mitigation measures provided by SpFi:

- Transceiver: Examine the effects of radiation on the transceiver blocks and characterise the cross-sections of their internal components.
- FPGA fabric: Assess how efficiently XilSEM corrects CRAM upsets, and how this impacts the operation of the link.
- SpFi IP: Verify the operation of the Error Detection and Correction (EDAC) mechanism in the internal buffers, and evaluate the robustness of the clock and reset scheme in the implementation. Also, assess the reliability improvements provided by distributed TMR (DTMR) and ensure that the transceiver reset is

SpFi protocol: Check how the FDIR mechanism operates in the Versal. Ensure quick recovery from all single event upsets (SEUs) causing transient and persistent errors.

II. TEST DESIGN

The VCK190 Evaluation Kit was selected as a test vehicle for the radiation campaign. This board features a VC1902 Versal part, which is the commercial equivalent of the space-grade XQRVC1902. For the purposes of this radiation campaign, both parts are considered equivalent. Additionally, this board provides a series of high-speed interconnects (e.g. zSFP, zQSFP, FMC+) that can be used for testing the embedded high-speed GTY transceivers.

A. Test Architecture

Fig. 1 outlines the architecture of the design tested. The left square (DUT) represents the part of the design implemented inside the Versal FPGA and subject to radiation. The right square corresponds to an external—not irradiated—STAR-Ultra PCIe unit that interfaced with two SpFi links in the Versal (I1 and I2). The STAR-Ultra PCIe provides two independent 4-lane SpFi links via QSFP+ interfaces, with each lane currently supporting rates up to 7.8 Gbit/s. These interfaces can be used to transfer data at high speed to/from a host PC (over PCIe Gen3). In this case, I1 implemented a quad-lane SpFi link (25 Gbit/s) and I2 a dual-lane SpFi link (12.5 Gbit/s). I3 and I4 were links connected to themselves via an FMC+ loopback card, due to limitations on the external equipment available. Specifically, I3 consisted of a quad-lane SpFi link (running at 25 Gbit/s) whereas I4 implemented an *experimental* quad-lane SpFi link running at 100 Gbit/s (25 Gbit/s per lane).




Fig. 1. Radiation test architecture.

Internal data checkers and generators, plus a *Tracer* monitor providing high resolution measurements, were implemented for each link. An additional *Test Monitor* block oversaw the status of each of the links.

B. Radiation Monitors

Two complementary STAR-Dundee monitor tools were used to maximise the value of the data collected. These allowed single event effect (SEE) classification by source and effect. These tools are not specific to SpFi or transceiver testing and can therefore be used for testing other FPGA blocks if needed.

One tool interfaced directly with the *Test Monitor* block in Fig. 1, enabling the logging of all events and user interactions with the system at sub-second resolution. The system status was presented to the operator in real time through a computer GUI (Fig. 2), allowing to monitor the design and detect events requiring user intervention, such as single event functional interrupts (SEFIs). Upon encountering a SEFI, the board was reprogrammed to reestablish functionality.




Fig. 2. Test Monitor information displayed in the GUI.

The second tool, the *Tracer* monitor, is an embedded logic analyser-like tool, providing nanosecond-accuracy measurements for relevant events in the design. It uses timestamping to record signal transitions. This tool enables the analysis of SEEs with maximum detail, thereby allowing a more accurate determination of the event source. Fig. 3 shows an actual event captured by the Tracer. The data was converted into a VCD file and loaded into ModelSim for visualisation.




Fig. 3. Tracer displaying an event using the Modelsim GUI.

The data processing chain included an event classifier based on internal STAR-Dundee verification tools, along with additional Python libraries tailored to the nature of the radiation data collected.

III. TEST CAMPAIGN

The tests were carried out in the autumn of 2024 at GANIL (France) and were funded by the RADNEXT European H2020 project [4]. The effective LET used covered a range from 28 to 43 MeV·cm²/mg. Lower LETs were unfortunately not tested as only Xenon ions were available during the two 8-hour test windows. Therefore, the results presented can be considered a worst-case scenario of heavy-ion radiation sensitivity measurement. An additional campaign is planned to test LETs < 25 MeV·cm²/mg in late 2025, which will provide the onset threshold and intermediate cross-section

values necessary for estimating the Weibull curves required to calculate the error rates for different orbit profiles.

Fig. 4 shows the VCK190 board with an unlidded Versal fixed to the supporting frame ready for irradiation. The important 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+ cables connecting to a 2-lane SpFi link (I2).
- 9) QSFP+ cable connected to a 4-lane SpFi link (I1).




Fig. 4. Test Setup with the VCK190 board in place for radiation testing.




Fig. 5. Equipment setup in the radiation chamber.

Fig. 5 shows the radiation chamber with the board prepared for irradiation—hidden from view, approximate location defined by the red parallelogram—and the rest of the equipment located in the platform below the board.

IV. RESULTS

A. Transceiver Effects

Fig. 6 shows the aggregated cross-section for the four channels composing a transceiver Quad. This value accounts for various error cases affecting transceiver components such as the TXPLL and the data path, and indicates the resulting cross-section that would be experienced by a quad-lane link using such Quad. The blue line corresponds to the cross-section measured on Quads connected to external equipment, whereas the orange line corresponds to the same measurement on transceivers connected via a physical loopback. The cross-section is an order of magnitude greater when connected to external equipment than when using a loopback, demonstrating the importance of testing designs using a representative setup. Green point and error bars correspond to SEFI events that required board reprogramming.




Fig. 6. Aggregated transceiver Quad cross-section. Blue: entire Quad. Orange: entire Quad connected in physical loopback. Green: Quad SEFI.

Fig. 7 provides the cross-section for different elements composing a transceiver Quad. In blue, the Quad PLLs; in orange, the channel data path; in green, the Quad shared logic. Each of the Quad PLLs is shared among two Quad channels. These Quad PLLs are the most common source of radiation errors, producing small bursts of errors or skew changes in the channels. Either a Quad PLL or a channel data path event triggers a SpFi retry event, with recovery occurring in less than 4 μ s. This represents about 99% of the total number of events. An event affecting the Quad shared logic, on the other hand, requires a full transceiver reset, lasting ~1.5 ms. This represents ~1% of the total number of events. Note that SpFi automatically recovers all these events, but they may impact operation when alternative protocols are used.




Fig. 7. Cross-section for the transceiver Quad elements. Blue: Quad PLL. Orange: Channel data path. Green: Quad shared logic.

B. Fabric Effects & DTMR

A protocol is typically implemented in the FPGA fabric to enable data transmission through a transceiver. However, both the fabric and its associated CRAM are vulnerable to SEUs, adding another source of errors. Fig. 8 shows the cross-section of a quad-lane SpFi link. As expected, SEUs in the FPGA CRAM and fabric affect the operation of the SpFi IP, as shown by the blue line. 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 specific logic part affected by the event. Nevertheless, the SpFi IP fabric events are still an order of magnitude less frequent than those of the transceiver Quad—blue line of Fig. 6. This difference is partly explained by the small footprint of the SpFi IP in the Versal, with a quad-lane link using only about 0.5% of the fabric resources available in an XQRVC1902, thus minimising the number of SEUs on the FPGA fabric.




Fig. 8. Cross-section for the fabric of a SpaceFibre IP link. Blue: nominal. Orange: DTMR applied to most of the link.

The orange line in Fig. 8 represents the cross-section for the same SpFi link when distributed TMR (DTMR) is applied. A design with DTMR was generated using Synplify Elite and tested in some of the runs. Due to a design flaw, DTMR was not applied to the data generators and checkers, which

accounted for 12% of the link logic. Consequently, only about 12% of the original events would be expected if DTMR provided full protection against fabric SEUs. This aligns with the improvement factor observed in the cross-section between blue and orange curves of Fig. 8. Therefore, it is reasonable to assume that applying DTMR to the SpFi link would eliminate the impact of SEUs on the FPGA fabric and CRAM.

Table I reports the resource usage for different configurations of the SpFi IP Core, with and without DTMR. As shown, there is a significant penalty in complexity (i.e. resource usage) when DTMR is applied. This highlights the importance of using IP cores with a small footprint in critical applications where DTMR is needed.

TABLE I. SpFi IPs Usage BEFORE AND AFTER TMR IS APPLIED

	XQRVC1902 (Nominal)			XQRVC1902 (DTMR)		
	LUT	DFF	RAMB36	LUT	DFF	RAMB36
1 Lane	1797	2151	4	11555	6459	12
1 VC	0.2%	0.1%	0.4%	1.3%	0.4%	1.2%
1 Lane	2164	2604	6	14487	7713	18
2 VCs	0.2%	0.1%	0.6%	1.6%	0.4%	1.9%
4 Lanes	5456	6726	12	42130	23949	36
1 VC	0.6%	0.4%	1.2%	4.7%	1.3%	3.7%
4 Lanes	6980	9375	30	61492	31557	90
4 VCs	0.8%	0.5%	3.1%	6.8%	1.8%	9.3%

C. Link Radiation Effects

Very high-speed commercial communication protocols (e.g. InfiniBand, Ethernet or Fibre Channel) use forward error correction (FEC) to reduce channel errors. When operating at speeds of about 25 Gbit/s or higher per lane, these protocols have adopted Reed-Solomon (RS) coding as FEC. Specifically, RS(544, 514)—capable of correcting a maximum error burst length of 150 bits—is used by the fastest versions running up to 100 Gbit/s per lane [5]. FEC is essential at such high speeds, where techniques like PAM-4 modulation are adopted but result in a degraded channel bit-error rate. In terrestrial applications, RS(544, 514) has proved to be a good compromise between implementation complexity and error-correction performance, hence its widespread adoption in these commercial protocols.

In space, however, radiation introduces an additional source of errors. For Versal transceivers, the dominant radiation effect has been shown to be bursts of errors. The length of these bursts is characterised in the histogram of Fig. 9. Bursts shorter than 150 bits (green bars) can be corrected by RS(544, 514), while longer bursts (blue bars) cannot. This histogram shows that only about 9% of transceiver transient events are correctable by the RS(544, 514) used in Ethernet or Fibre Channel. For the remaining events, hot-link redundancy or a reliable protocol such SpFi is required to ensure fast recovery.

In all these events SpFi achieves a rapid recovery by avoiding timeouts, using only the round-trip delay to signal errors. Fig. 10 presents the SpFi error recovery time for the transceiver transient SEEs. This histogram shows that most transient SEEs—about 99 % of all transceiver events observed—are recovered automatically within a few microseconds, with the vast majority completing in under 5 μ s.




Fig. 9. Transceiver transient error burst length distribution.

Considering all error types measured in a SpFi quad-lane link, the statistics are as follows:

- Transient transceiver errors account for 85% of events. SpFi always recovers them in under 10 μ sec, with most of the events recovered in less than 5 μ sec (Fig. 10).
- Fabric errors represent 14% of events. These are automatically recovered by the embedded XilSEM scrubber in about 15 msec, although data loss may occur. Results suggest that DTMR in the SpFi IP can protect against these errors, but further testing is needed to confirm full mitigation.
- Persistent transceiver errors represent 1% of events. SpFi recovers all of them automatically in under 3 ms, as they involve a full transceiver reset.
- Transceiver SEFI events are very rare (less than 0.1% of events) and require FPGA reprogramming.

Overall, SpFi provides a robust solution for reliable data transfer inside a spacecraft. Combined with DTMR, it automatically mitigates all observed errors except SEFIs, which are extremely rare.




Fig. 10. SpaceFibre error recovery time for transceiver transient SEEs.

D. 100G SpFi Link

Additionally, the radiation campaign demonstrated an experimental 100 Gbit/s SpFi link operating under radiation. This was implemented using a quad-lane configuration, with each lane operating at 25 Gbit/s. While its operation was validated under nominal conditions, further testing is needed as its error recovery under radiation was less robust than the rest of the SpFi links. Further details of this upgraded SpFi link will be presented at this same Conference [6].

V. CONCLUSION

A heavy-ion campaign was conducted to validate SpFi running on Versal. The analysis of the collected data confirms that SpFi automatically recovers from all non-SEFI events affecting the transceiver without data loss, typically in less than 5 μ sec. The transceiver cross-section remained consistent for both 100 Gbit/s and 25 Gbit/s links, despite the use of different lane rates. Due to limitations on the facility ion cocktail available, only high LET testing was performed (28–43 MeV·cm²/mg). The results presented correspond to the saturation part of the Weibull curves but allow nonetheless to have a clear idea of the different effects induced by radiation in both the transceiver and fabric of the Versal, and an initial estimation of their likelihood. Another test campaign is planned for late 2025 to complete the results at lower LETs.

The fabric results demonstrate that a quad-lane SpFi link presents an order of magnitude lower SEU sensitivity than the Versal Quad transceiver. These SEUs can still be mitigated using the FPGA inbuilt scrubbing mechanism (XilSEM), which responds within tens of milliseconds. While data loss may occur during this process, the SpFi link self-recovers, ensuring that system functionality remains intact. Moreover, applying DTMR to the SpFi IP appears to eliminate its radiation sensitivity (pending confirmation), as SEUs are automatically corrected by XilSEM without affecting protocol operation, thanks to the triplication mechanism.

Additionally, the campaign successfully demonstrated a 100 Gbit/s SpFi link—implemented using 25 Gbit/s lanes—operating under radiation. This paves the way for 200 Gbit/s SpFi links in the near future.

These results underscore the potential of the SpFi-Versal combination as a high-reliability solution for spacecraft data-handling systems, enabling the development of next-generation, high-throughput, radiation-tolerant communication architectures.

ACKNOWLEDGMENT

This activity has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101008126, corresponding to the RADNEXT project.

We would like to 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] 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>.

- [2] S. Parkes et al., “SpaceFibre Onboard Interconnect: from Standard, through Demonstration, to Space Flight”, IEEE Aerospace Conference, Big Sky, Montana, 2025.
- [3] 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 15th September 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] P. Anslow, RS(544,514) FEC performance including precoding, IEEE P802.3cd Task Force, July 2016. [Online]. Available: https://www.ieee802.org/3/cd/public/July16/anslow_3cd_01_0716.pdf (last visited 15th September 2025).
- [6] A. Ferrer et al., “100 Gbit/s-plus SpaceFibre on Space-Qualified FPGAs”, European Data Handling & Processing Conference (EDHPC), Elche, Spain, 2025.