

Mitigation and Recovery of Single Event Effects in RTG4™ Transceiver Links Using SpaceFibre

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

Chris McClements
STAR-Dundee Ltd.
Dundee, UK
chris.mcclements@star-dundee.com

Nadia Rezzak
Microchip
San Jose, USA
nadia.rezzak@microchip.com

Abstract— SpaceFibre (ECSS-E-ST-50-11C) is a very high-performance, high-reliability and high-availability network technology specifically designed to meet the needs of space applications. It requires serial transceivers for implementation. The Microchip RTG4™ FPGA is equipped with 24 serial transceivers, each capable of supporting rates up to 3.125 Gbit/s. STAR-Dundee provides SpaceFibre communication IPs optimized for deployment on RTG4 FPGAs. Multiple lanes can be grouped to achieve aggregate bandwidths up to 25 Gbit/s. In this paper we present the results of radiation tests examining heavy-ion single event effect behavior of the serial transceivers used on the RTG4 FPGA while implementing SpaceFibre data links. We study both the native upset rates of the RTG4 transceivers and the resulting SpaceFibre link error rates. Results show that the use of SpaceFibre mitigates the effects of radiation allowing to create reliable links in harsh conditions.

Keywords—SpaceFibre, SpaceWire, Radiation Testing, RTG4, SerDes, Transceiver, FPGA

I. INTRODUCTION

Radiation tolerant devices for space applications are essential. These devices must operate in the toughest environmental conditions without compromising a mission, including extreme temperature ranges and ionising radiation. Therefore, for its reliable use in space applications the Microchip RTG4 FPGA [1] shall present as much immunity to radiation-induced effects as possible [2].

The RTG4 features 6 transceiver (SerDes) blocks, each with 4 separate lanes (24 high-speed communication interfaces), each running up to 3.125 Gbit/s. Effects of the radiation in the transceivers include bit-flips and error bursts in data reception, loss of lock of the clock data recovery mechanism (CDR) or the PLL, and bit-flips on some of the transceiver configuration registers. The configuration registers of the transceivers are not radiation hardened [3] and a mechanism to minimise the effect of the radiation in these registers is required. The registers can be accessed over a APB configuration interface in each transceiver, through the FPGA fabric.

SpaceFibre (ECSS-E-ST-50-11C) is a very high-speed serial link developed by the University of Dundee for the European Space Agency which is intended for use in data-handling networks with high data-rate payloads [4][5][6][7]. The aim of SpaceFibre (SpFi) is to provide point-to-point and networked interconnections for Gigabit rate instruments, mass-memory units, processors and other equipment, on

board a spacecraft. SpaceFibre can operate over fibre-optic and copper cable and currently supports lane rates of 3.125 Gbit/s in the RTG4 FPGA (with EDAC and SET filter enabled and worst-case conditions). Its high data rate per lane coupled with novel multi-lane technology enables SpFi to achieve very high performance: an 8-lane link provides up to 25 Gbit/s in the RTG4. Its in-built error detection, isolation and recovery mechanisms enable rapid recovery from transient errors, without loss of data, providing high availability. Its multi-lane hot and cold redundancy features support high reliability. These capabilities are built into the hardware of each SpFi interface

The main goal of the test campaign was to best characterise the transceiver performance under radiation while minimising the test complexity. Furthermore, improvements to the transceiver operation reliability provided by the SpFi protocol were also to be assessed. Finally, an additional goal of this campaign was to validate the correct performance of the SpFi and SpaceWire [8] (SpW) interfaces under radiation using the IPs available from STAR-Dundee.

II. TEST SET-UP

The test set-up was originally designed to use all 6 transceiver blocks available in the RTG4. The test board selected was the RTG4 Development Kit, which offers two FMC-HPC connectors for interfacing some of the RTG4 transceivers and pins. All lanes ran at 2.5 Gbit/s, and the links were saturated with data to maximise the probability of detection of single event effects (SEE) affecting the data integrity. Due to physical limitations, not all the transceivers were interfaced externally, with some using either physical loopback cables or internal loopback functionality (PMA near-end loopback):

- 2x Transceivers (i.e., 8 lanes) operated in PRBS mode to observe the effects of radiation natively on the transceivers. All were configured in PMA near-end loopback.
- 3x Transceivers operated using SpFi Multi-Lane links. One Transceiver used PMA near-end loopback. The other two were connected via SpW-SpFi FMC Boards [9] to an external STAR Fire [10] unit and external loopback cables respectively.
- 1x Transceiver operated using a SpFi Single-Lane link. It was connected in loopback via SMA interfaces.

Apart from that, SpW links were also tested:

- 1x SpW was used for the Status & Control monitor of the test.
- 6x SpW interfaces used the dedicated Clock Recovery Circuits in the RTG4.
- 1x SpW Interface used a fabric clock recovery circuit.

Data generators and checkers controlled by the monitor were connected to all the SpFi and SpW interfaces to verify their correct operation.

Fig. 1 shows a block diagram of the set-up. HPC1 and HPC2 are used by the SpW/SpFi FMC boards to feed reference clocks in the transceivers and to provide SpW and SpFi physical interfaces.

Fig 2 shows a block diagram of the SpW test set-up. A total of 6 test Interfaces are connected on FMC HPC1 and HPC2 via SpW/SpFi FMC boards. A loopback setup is used to pair test interfaces.

Six SpW interfaces are instantiated in the RTG4 fabric with built-in self-test functions including packet checking and status monitoring. The SpW interfaces are configured to always initiate a link connection and to send data continuously. The SpW bit rate is 100 Mbit/s resulting in a bi-directional user data transfer rate of approximately 76 Mbit/s. A SpW interface is reserved as a redundant software access interface to the monitor via an RMAP target. One of the SpW interfaces is instantiated without dedicated data and clock recovery CCC due to pin connection constraints when an FMC board is connected on HPC2.

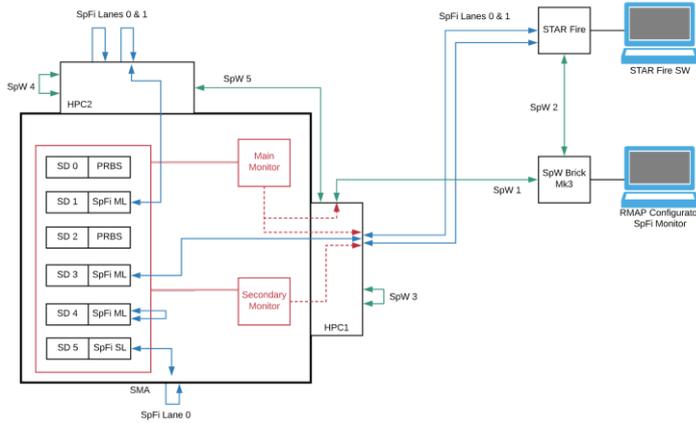


Fig. 1. SpFi radiation test set-up block diagram.

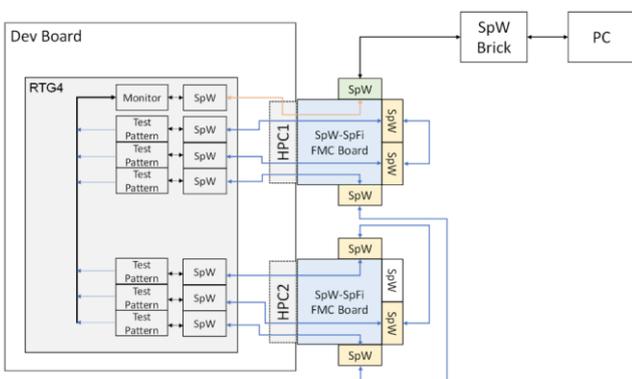


Fig. 2. SpW radiation test set-up block diagram.

Each of the SpW test interfaces has monitoring logic to capture the number of events from test start including link events, SRAM events and packet checker events. A link event is recorded on detection of a disconnect, a decoder error, or a credit error. Link events cause the link to exit from the Run state and reinitialise exit from the Run state. SRAM events including single bit error correction and double-bit error detection flags generated by the embedded RTG4 μ SRAM EDAC are recorded. A self-test packet transfer monitor is used to check received data and can record packet anomalies including out of sequence, data error, error end of packet received, and packet length errors.

III. TEST CAMPAIGN

The test campaign was carried out at the TAMU Cyclotron in-air testing facility located in Texas (USA). Three different ions were used: Nitrogen, Argon and Krypton with four different effective linear energy transfer (LET) tested: 0.91, 2, 6.29 and 24.5 MeV*cm²/mg. Higher LETs were initially planned but not tested due to the lack of beam time. The different runs have been summarised in Table I. The final 3 runs were affected by oscillations in the beam flux and the final fluence of 10⁷ ions/cm² was not achieved in each of these runs.

Fig. 3 shows an image of the test board with an unlidged RTG4 FPGA in position for one of the runs. On the top of the RTG4 test board, the SpFi and SpW external connections through an FMC daughterboard can be seen.

The RTG4 is an especially large device. This caused a problem with the device irradiation, as it was discovered after the tests that the beam aperture was not large enough for radiation to affect the entire RTG4 die. Consequently, not all the transceiver blocks were fully irradiated, causing differences in the number of events observed. The transceivers blocks are all placed along the north side of the RTG4. The corners of the device did not receive the total fluence of the test. Therefore, the external transceivers have been removed from the data analysis. The resulting set-up used in the analysis of the data consists of two RTG4 transceivers (placed close to the centre of the north side), both used with SpFi Multi-Lane interfaces. Of those, one was configured in PMA near-end loopback (a 3-lane link), and the other one was connected to an external 2-lane SpFi interface (STAR Fire). No significant statistical differences were measured in the events of these two SpFi links.

TABLE I. CAMPAIGN RUN SUMMARY

Run	Ion	Energy (MeV)	Effective LET (MeV*cm ² /mg)	Flux (ions/cm ² /s)	Fluence (ions/cm ²)
1	N	25	0.91	1.30E+04	1.16E+07
2	N	25	0.91	1.38E+04	1.00E+07
3*	N	25	2	1.38E+04	1.00E+07
4*	N	25	2	1.33E+04	1.00E+07
5	Ar	25	6.29	1.02E+04	1.00E+07
6	Ar	25	6.29	1.37E+04	1.00E+07
7	Kr	25	24.5	1.24E+04	9.99E+06
8	Kr	25	24.5	1.27E+04	1.44E+06
9	Kr	25	24.5	9.57E+03	7.76E+05
10	Kr	25	24.5	4.68E+03	3.98E+06

* Degraded used (24 mil)

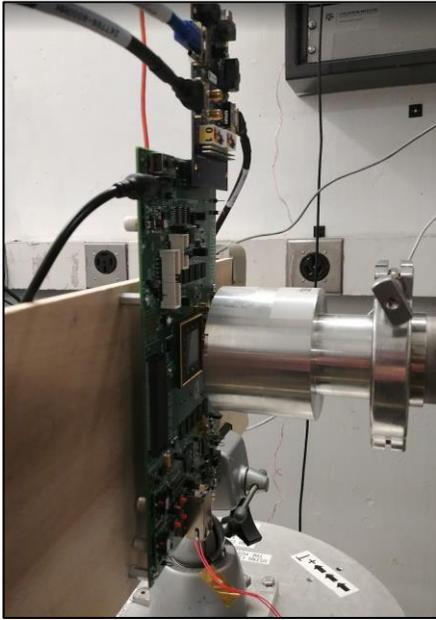


Fig. 3. Unshielded RTG4 under the beam.

IV. RESULTS

The following subsections detail the radiation effects observed in the subset of transceivers deemed relevant for analysis as explained in Section III, and the SpW links. Note that for events reported per SpFi lane, the probability will increase directly with the number of lanes composing a link.

A. SEFI Lane Events

Single-event functional interrupts (SEFI) are defined as events leading to a potentially persistent failure of the transceiver lane, meaning these cannot be recovered by resetting the lane. In the tests, lanes were recovered by a dedicated SerDes Recovery Block (SRB). This block is a bespoke development by STAR-Dundee. The SRB rewrites the important transceiver lane block registers (those are the ones not hardened in the RTG4) with their correct values. To assess its effectiveness, one of the SpFi links had its SRB enabled, while the other link had it disabled. The transceiver using the recovery block never experienced a SEFI. The transceiver not using the recovery block, on the other hand, experienced several SEFIs. However, the transceiver was able to immediately recover after enabling the SRB.

It is worth mentioning that there was one SEFI in the whole campaign in which the SRB was not able to immediately recover the transceiver lane. In this case, the lane spontaneously recovered after 200 seconds. One hypothesis for its cause is that an important register may be missing from the SRB register list. A SEE affected this register which in turn caused the SEFI. The resulting cross-section for this event was 3.1×10^{-8} for an LET of $24.5 \text{ MeV} \cdot \text{cm}^2/\text{mg}$.

Fig. 4 shows the cross-section for lane SEFI events.

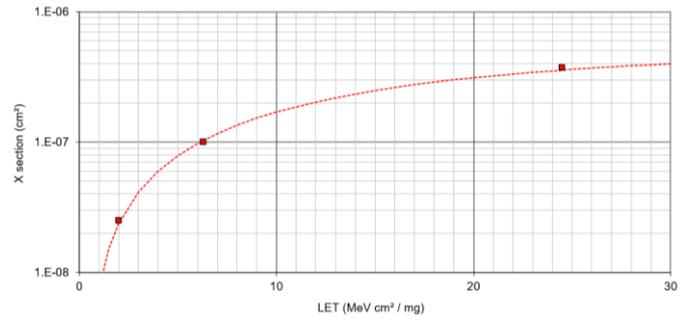


Fig. 4. Transceiver lane SEFI events cross-section.

B. Lane Disconnections

A lane disconnection is a transient event that can be triggered by:

- The loss-of-signal/electrical idle circuitry activated in the lane receiver.
- A burst of errors (> 255) received in a short period. In this case, SpFi reinitialises the lane to recover from the burst and ensure that the connection is reliable.

Fig. 5 shows the cross-section for the lane disconnection events.

When a lane reconnects, a recovery procedure is automatically initiated by SpFi: the data affected by the disconnection is automatically resent. Recovery from a disconnection is typically very fast. Fig. 6 shows the recovery time histogram for the lane disconnections. Three regions have been identified:

- A) Typically (62% of the cases), it takes less than 110 μsec to recover from a lane disconnection. These cases (e.g., CDR fail) are recovered by SpFi resetting the lane. Note that this recovery time is dependent on the lane speed, so faster recovery times will be obtained for faster lanes.
- B) Other cases (33%) show a longer recovery of $\sim 350\text{--}500 \mu\text{sec}$. These cases correspond to the SEFI Lane Events described in subsection A. As previously discussed, Lane SEFIs are recovered by the SRB. In this case, most of the delay difference with respect to case (A) corresponds to the data gathering procedure for test analysis. Thus, similar recovery times to (A) are expected for the final version of the SRB.
- C) 4 cases (5%) were measured at $\sim 2\text{--}2.5 \text{ msec}$. These need further investigation.

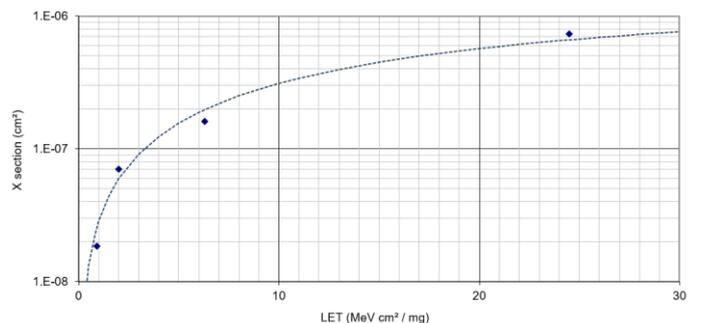


Fig. 5. Transceiver lane disconnections cross-section.

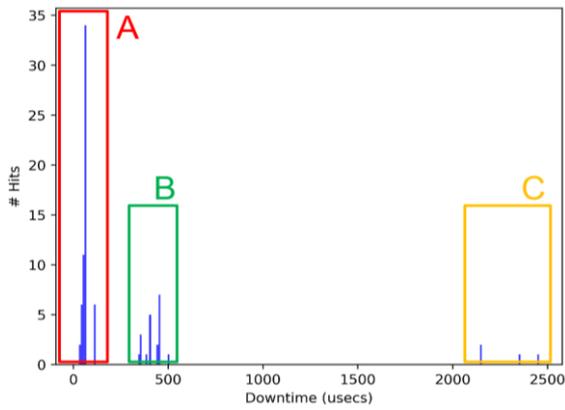


Fig. 6. Lane disconnection recovery time histogram.

In summary, the recovery from a lane disconnection typically requires less than 110 μsec , with a worst-case of 2.5 msec.

C. Retry Events

A Retry event happens when the link detects an error in a data frame (e.g., a bit-flip affecting a data character). In this case, the corrupted data frame is automatically resent (retry). This procedure is very fast, less than 3 μsec , which makes it transparent to the application.

A Retry event provides a good estimation of the number of SEE affecting the data and lanes/link, as they are by far the most common SEE affecting the Transceiver. Therefore, if no protocol were to be used, the number of retry events would represent the number of error bursts received. SpFi provides a means to filter these out so that their impact on the application is greatly reduced.

Fig. 7 shows the cross-section for the retry events.

D. Data Errors

A few data errors were detected by the checker application. These corresponded to errors in the data that were not detected by the SpFi link.

There were no errors detected for LETs equal to or less than 2 $\text{MeV}\cdot\text{cm}^2/\text{mg}$; they were first observed at 6.29 $\text{MeV}\cdot\text{cm}^2/\text{mg}$. Further investigation is needed, but one possibility is that they are caused by single-event transients (SET) affecting the data path of the link (Fig. 8).

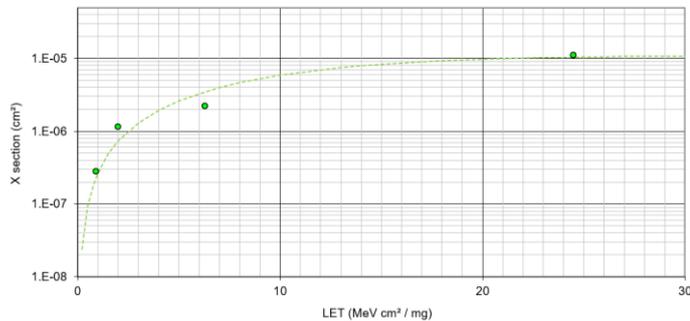


Fig. 7. Retry events cross-section.

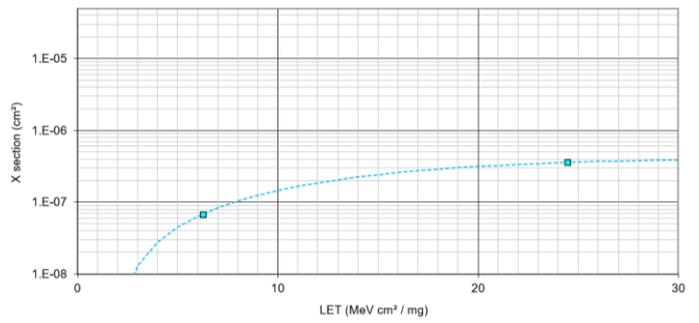


Fig. 8. Data error events cross-section.

E. Combined Errors

All previous cross-sections have been combined in a single figure (Fig. 9) for convenience.

The most likely event is the Retry, which also provides an estimation of the number of SEEs affecting the link. This value represents the number of burst of errors received by the Transceiver, although these errors are automatically corrected by SpFi. As these errors are contained at the link level — managed in hardware by SpFi— recovery is so fast that the application does not notice.

More than an order of magnitude less likely are the Disconnections events, which are also automatically managed by SpFi. In this case, a lane reinitialises typically in less than 100 μsec (lane downtime). However, for Multi-Lane links, the link automatically reconfigures to operate with any remaining lanes in ~ 2 μsec through a mechanism known as graceful degradation. This provides a working link with a reduced throughput which depends on the number of remaining lanes, but the highest priority data can still be transferred thanks to the embedded quality of service capabilities of SpFi. As soon as the lane is reinitialised, the Multi-Lane link reconfigures again to continue operation with the recovered lane. Any number of lanes between 1 and 16 is supported by SpFi, providing maximum flexibility and robustness in terms of lane usage.

SEFI rates are even lower, but the effect of these events in the SpFi link is similar to lane disconnections when the SRB is used: SEFIs will not occur if properly addressed. In fact, the Disconnection cross-section also includes SEFI-caused disconnections. Note that there was one SEFI case observed which needs further investigation as the SRB did not help to recover.

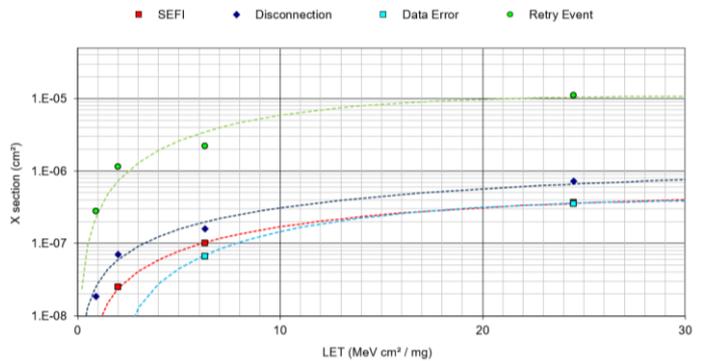


Fig. 9. Transceiver lane retry events cross-section.

Finally, we have the data errors, which are critical, as they compromise the data received by the application. No data error was observed at low LETs, with the first occurrences observed at 6.29 MeV*cm²/mg. Additional investigation is required to see whether they can be prevented.

F. SpaceWire

The SpW interface test results are shown in Table II below. During the test, it was observed that no link errors occurred for LETs less than 24.5 MeV*cm²/mg. The SpW interface core has a small footprint which may explain why errors were not observed at lower LET.

TABLE II. SPACEWIRE TEST INTERFACES

Effective LET (MeV*cm ² /mg)	Fluence (ions/cm ²)	Bit Rate (Mbit/s)	Link Error Count	Packet Error Count	Cross Section (Link Errors/Fluence)
24.5	1.62E+07	100	12	22	7.41E-07

The link error count indicates the number of times a link exited the Run state due to the detection of a parity or decoder error at the receiver. On detection of an error, the link will reinitialise and interrupt any packets in progress using the error recovery mechanism defined in the SpW protocol [8]. Packet transfer is bi-directional therefore the number of observed packet errors is expected to be larger than the number of recorded link errors as packet reception at each end of the link is interrupted.

V. CONCLUSION

In this paper we have shown how heavy ion induced SEEs on the RTG4 transceivers can produce temporary and permanent errors on communication links when simple protocols are used. The protection added when using SpaceFibre has been demonstrated. Most SEE effects on the data can be transparently fixed by SpaceFibre, many without any noticeable effect on the application. However, a few data errors induced on the IP Core logic have been observed. Further tests should allow elucidating their causes.

REFERENCES

- [1] <https://www.microsemi.com/product-directory/rad-tolerant-fpgas/3576-rtg4> (last accessed 07/04/2022).
- [2] N. Rezzak, J.J. Wang, D. Dsilva and N. Jat, "TID and SEE Characterization of Microsemi's 4th Generation Radiation Tolerant RTG4 Flash-Based FPGA," 2015 IEEE Radiation Effects Data Workshop (REDW), 2015, pp. 1-6, doi: 10.1109/REDW.2015.7336739.
- [3] RTG4 FPGA High-Speed Serial Interfaces, User Guide (UG0567).
- [4] ECSS Standard ECSS-E-ST-50-11C, "SpaceFibre – Very high-speed serial link", European Cooperation for Space Data Standardization, 15th May 2020. Available from <http://www.ecss.nl>.
- [5] S. Parkes, C. McClements and M. Suess, "SpaceFibre", International SpaceWire Conference, St Petersburg, Russia, 2010, ISBN 978-0-9557196-2-2, pp 41-45.
- [6] S. Parkes et al, "SpaceFibre: Multi-Gigabps Interconnect for Spacecraft On-board Data Handling", IEEE Aerospace Conference, Big Sky, Montana, 2015.
- [7] A. Ferrer Florit, A. Gonzalez Villafranca and S. Parkes, "SpaceFibre Multi-Lane", International SpaceWire Conference, Yokohama, Japan, 2016, ISBN 978-0-9954530-0-5.
- [8] SpaceWire – Links, Nodes, Routers and Networks (ECSS-E-ST-50-12C Rev.1).
- [9] <https://www.star-dundee.com/products/fmc-spacewire-spacefibre-board/> (last accessed 07/04/2022).
- [10] <https://www.star-dundee.com/products/star-fire-mk3/> (last accessed 07/04/2022).