

SpaceFibre IP Cores for Fast Adoption of Next-Gen FPGA Communication Architectures

Alberto Gonzalez Villafranca
STAR-Barcelona S.L.,
Av. Cerdanyola, 79-81, p.1 of.11
08172 Sant Cugat del Vallès
Barcelona, Spain
alberto.gonzalez@star-dundee.com

Albert Ferrer Florit
STAR-Barcelona S.L.,
Av. Cerdanyola, 79-81, p.1 of.11
08172 Sant Cugat del Vallès
Barcelona, Spain
albert.ferrer@star-dundee.com

Marti Farras Casas
STAR-Barcelona S.L.,
Av. Cerdanyola, 79-81, p.1 of.11
08172 Sant Cugat del Vallès
Barcelona, Spain
marti.farras@star-dundee.com

Steve Parkes
STAR-Dundee Ltd.,
STAR House, 166 Nethergate
Dundee, DD1 4EE, United Kingdom
steve.parkes@star-dundee.com

Abstract—SpaceFibre is an advanced spacecraft on-board data-handling network technology, building upon its predecessor, SpaceWire, to meet the increasing demands for higher data transfer rates and improved reliability in space applications. This open standard (ECSS-E-ST-50-11C), promoted by the European Space Agency (ESA), has been integrated into numerous spacecraft standards including SpaceVPX, SpaceVNX+, and ADHA. SpaceFibre’s built-in quality of service (QoS) and fault detection, isolation, and recovery (FDIR) capabilities ensure high reliability and availability – critical for spacecraft operations. It operates over both copper and fibre-optic cables and supports any number of lanes (up to 16) in multi-lane mode, allowing links with different numbers of lanes to seamlessly interoperate and providing rates exceeding 50 Gbit/s in existing space-qualified technology, and soon 200 Gbit/s. Its network architecture supports virtual networks and arbitrary packet sizes, enhancing the scalability and flexibility of spacecraft data systems. Simple nodes can achieve maximum throughput since SpaceFibre is designed to operate without the need for CPU intervention. Additionally, it maintains backwards compatibility with SpaceWire at the network level, facilitating integration with existing systems. Capable of running on up to 100 meters over fibre optics and inherently designed to handle high data volumes from sophisticated payloads, SpaceFibre is a critical technology for existing and future spacecraft communication architectures.

STAR-Dundee has developed a comprehensive suite of SpaceFibre IP cores, consisting of the Single-Lane Interface, Multi-Lane Interface, and Routing Switch. Optimized for space applications, these cores target existing space-qualified FPGAs for improved footprint and speed. This paper presents recent updates to these IP cores. It analyzes their performance using metrics such as maximum lane rates and resource usage, and provides the latest results on heavy-ion radiation testing. Support has been added for new radiation-tolerant FPGAs, including Frontgrade CertusPro-NX-RT, Microchip RT-PolarFire SoC, AMD Versal, and NanoXplore NG-Ultra. A novel and significant architectural improvement is the encapsulation of all transceiver logic within a single block. This simplifies user interaction by requiring only a minimal interface between the transceivers of supported FPGAs and the SpaceFibre IP. This design enables easy setup for various architectures via a configuration file, streamlining integration and reducing complexity.

The SpaceFibre IP cores presented have achieved TRL-9, having been deployed in at least six operational missions since 2021 and currently being designed into more than 60 spacecraft. The paper will demonstrate that the IP cores enable efficient and reliable data handling and processing, which is essential for mission success.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. SPACEFIBRE IP CORES COMMON INTERFACES ..	3
3. SPACEFIBRE INTERFACE IP CORE.....	4
4. SPACEFIBRE ROUTING SWITCH IP CORE.....	5
5. OPERATION UNDER RADIATION	8
6. CONCLUSIONS	10
ACKNOWLEDGEMENTS	10
REFERENCES.....	11
BIOGRAPHY	11

1. INTRODUCTION

SpaceFibre (SpFi) [1] is a communication technology for use onboard spacecraft which was promoted 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 Gigabit rates while offering QoS and FDIR capabilities. SpFi has been integrated into numerous spacecraft standards including SpaceVPX, SpaceVNX+, and ADHA [2].

SpaceFibre has an error recovery mechanism that automatically recovers from transient and persistent errors on the SpaceFibre link, with recovery from transient errors in typically less than 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 individual physical lanes. This multi-lane operation provides higher data rates through lane aggregation, supporting any number of lanes (up to 16) and unidirectional operation. This effectively multiplies the throughput of the interface by

combining several lanes into a link. Furthermore, when a lane fails, the multi-lane mechanism supports hot and warm redundancy and graceful degradation by automatically spreading traffic over the remaining working lanes.

SpFi interoperates seamlessly with a SpaceWire (SpW) [3] network over virtual channels (VCs) as it uses the same data packet definition. Furthermore, SpFi provides broadcast capabilities and can operate over either electrical or fibre optic cables. The Network layer in SpFi is responsible for transferring data packets over a link or network. The routing concepts are the same as in SpW including both path and logical addressing. The Network layer includes the definition of Virtual Networks (VNs). These VNs are built from the interconnection between VCs of different ports within a SpFi routing switch. Traffic within a VN does not create network congestion to another VN.

Table 1 shows how the main SpFi capabilities are compared against other OSI Layer 2 protocols that can be implemented in FPGAs. Three main advantages of SpaceFibre are:

- 1) The arbitrary packet size simplifies the host interface and avoids the need to implement packet segmentation by software or hardware means.
- 2) A reliable link handles nominal bit error rate (BER) and most single-event upsets (SEUs) produced by the radiation environment [4, 5].
- 3) A CPU is not required at each endpoint or switch.

In other networks like Ethernet or TSN, packet segmentation and guaranteed delivery are managed by layer 4 transport protocols like TCP and UDP, which also requires the implementation of the layer 3 network IP protocol. Due to their complexity, these protocols are usually implemented in software. When working at high data rates, this leads to high CPU utilisation, higher latency and increased power consumption.

In contrast, SpaceFibre is a simpler layer 2 protocol that handles segmentation, link-layer error correction, and flow

control in hardware. This hardware-based approach significantly reduces CPU load at the SoC based endpoints, resulting in lower latency and power usage even under demanding conditions. Upper layer protocols can still be implemented in software for specific use cases.

STAR-Dundee has developed a range of SpFi IP cores compliant with the standard. These IPs have been optimised for speed considering the timing constraints of the slower FPGAs for space and have undergone several radiation test campaigns in different FPGA families. The range is composed of two main IPs:

- SpFi Interface: Single-Lane (SL) and Multi-Lane (ML) are supported.
- SpFi Routing Switch: supports SpFi SL and ML ports, together with SpW and internal AXI4-Stream ports.

STAR-Dundee SpFi IPs are TRL-9, currently flying in at least six operational missions, including the Airbus Pleiades Neo constellation [2] and being implemented into more than 60.

New Space-Qualified FPGAs

A new generation of FPGAs specifically targeting space applications has recently been or is in the process of being released. This new generation aims at higher performance applications, which makes them ideal targets for using very high-speed communications protocols such as SpFi.

Microchip PolarFire SoC—The radiation-tolerant PolarFire (PF) SoC RTPFS460ZT / RTPFS160ZT and its close relative RTPF500ZT are directly derived from their commercial counterparts, non-volatile FPGAs built on a 28 nm process [6]. PF uses low-power SONOS configuration switches that have been demonstrated to be robust at 100 krad of total dose and having an absence of configuration upsets under heavy-ion single event tests. PF provides 20 / 8 / 24 transceivers with a maximum speed of 10 Gbit/s.

Table 1. Capabilities of Different Data Link Protocols

	Packet Size (bytes)	Flow control	Reliable	QoS	Redundancy	Switching	Broadcast	CPU required
SpaceWire	1 to ∞	Yes	No	No	No	Yes	No	No
SpaceFibre	1 to ∞	Yes	Yes	Yes	Redundant lanes	Yes	Yes	No
Ethernet	64 to 9038	Optional	No	No	No	Yes	Yes	Yes
TSN	64 to 9038	Optional	No	Yes	Redundant paths	Yes	Yes	Yes
PCI Express	128 to 4096	Yes	Yes	Yes	No	Limited	No	Yes
Aurora	1 to ∞	Optional	No	No	No	No	N/A	No
Interlaken	1 to ∞	Yes	No	No	No	No	N/A	No
JESD204	1 to 256	No	No	No	No	No	N/A	No

Unlike radiation-hardened devices, depending on the application the SEU rate of the PF registers may not be good enough. In this case, the use of some form of Triple Modular Redundancy (TMR) is advised. STAR-Dundee, in collaboration with other partners, has successfully completed a heavy-ion test campaign to validate the performance of the SpFi IPs under radiation.

AMD Versal—The AMD Versal XQRVC1902 / XQRVE2302 [7] 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 / 8 GTY transceivers (each capable of 25 Gbit/s). An internal scrubbing mechanism (XilSEM) has been implemented to quickly repair any configuration memory (CRAM) SEU. Using TMR in Versal is also advised depending on the application and its requirements regarding register SEU sensibility, as the FPGA is not radiation-hardened. STAR-Dundee recently completed a heavy-ion test campaign to validate the performance of the SpFi IPs under radiation in Versal. The initial analysis of the data gathered shows promising results (Section 5), and the full analysis will be published when completed.

FrontGrade CertusPro-NX-RT—The radiation-tolerant version of the CertusPro-NX is built on a 28 nm FD-SOI process [8]. It uses SRAM-based configuration memory and includes active reliability features (e.g. configuration memory frame-based soft error correction). This FPGA offers up to 8 transceivers, each capable of 6.25 Gbit/s and a compact footprint format. Again, using TMR is recommended, depending on the required SEU tolerance of the design.

NanoXplore BRAVE—The NanoXplore BRAVE FPGA family represents a European alternative among space-qualified devices. The family includes several members, with the NG-Ultra [9] (28 nm FD-SOI SRAM) as its flagship. It features 32 integrated SerDes blocks, supporting data rates of up to 12.5 Gbit/s. Its configuration memory is inherently radiation-hardened, eliminating the need for TMR. The SpFi Interface IP was validated in the previous generation, NG-Large, and efforts are ongoing to complete the validation of the SpFi IPs for the newer NG-Ultra.

Previous Generation FPGAs—In the resource usage tables, the two main FPGAs of the previous generation have been included for reference: the Microchip RTG4 (65 nm flash-based) [10] and AMD Kintex UltraScale XQRKU060 (20 nm SRAM-based) [11]. The RTG4 configuration memory is immune to radiation, and all its registers are internally triplicated, meaning that it is radiation-hardened by design, unlike the XQRKU060. Two radiation test campaigns have been carried out in collaboration with Microchip involving the SpFi Interfaces and the RTG4, which have demonstrated how SpFi adds robustness to operation of the transceivers under radiation [4, 5].

2. SPACEFIBRE IP CORES COMMON INTERFACES

This section describes the two main interfaces used by the STAR-Dundee SpFi IP Cores, the host interface for VC data transfer and the transceiver interface.

Host Interface (AXI)

The host interface for data transfer is based on the AXI4-Stream (AXI4-S) specification, with a dedicated AXI4-S interface assigned to each virtual channel [12]. This interface is well suited for the SpaceFibre protocol as AXI4-S supports the transfer of packets with an arbitrary number of bytes. Packets of any size, transmitted as a byte stream, are accepted. However, like most AXI4-S implementations, null bytes in the middle of the packet are not supported. Table 2 outlines the functionality of the interface signals. Backpressure or flow control is managed using TVALID and TREADY signals.

Table 2. Functionality of AXI4-Stream Interface Signals

Function	Signal
Data content	TDATA, TKEEP
Handshaking and backpressure	TVALID, TREADY
End of packet	TLAST
Packet ends with error	TUSER

A SpaceFibre VC acts as a FIFO across both ends of the link. The same data sent using a slave AXI4-S interface from one side of the link will appear at the other side of the link in the master AXI4-S corresponding to the same virtual channel. For a point-to-point application, this means that the user can just push the data to transfer to this interface without any other modification. In case there is a network, a header must be added containing the path or logical address of the destination. There is no need to provide any other information in the header related to the operation of the link. The QoS of the virtual channel is configured using the management interface. The QoS is related to the VC number, not a specific packet. The content of the packet is to be consumed by the end-user or by upper layer protocols, if any.

Transceiver Interface

A recent addition to simplify the use of the SpFi IPs consists of a new block (named XCVRS) that encapsulates the technology-related logic associated with the transceivers of the target FPGA. This block integrates all the transceiver functionality required by each technology, including the transceiver channels and its associated control and monitoring logic. It simplifies the process of interfacing with different FPGA transceivers by encapsulating the complexity within this block, effectively acting as a black box. The XCVRS block groups in two record types all the transceiver interface signals required for the operation of the SpFi Interface. One record is used as input port and the other one

as output port. Therefore, the user simply connects the two ports between the SpFi and XCVRs blocks, without having to worry about the specifics of each transceiver technology. This architecture provides three main advantages that reduce the time required for integrating SpFi in a design:

- 1) Decouples the complexity of the IP interface with each of the supported technologies. Different technologies have different interfaces, for example, the RTG4 offers a basic 20-bit parallel interface, while newer FPGAs typically offer line encoding/decoding and clock compensation functionality embedded inside the transceiver block. Using the encapsulation XCVRs block simplifies the integration of the IP as it simply requires connecting two ports, input and output (Figure 1).
- 2) This novel architecture allows to transparently replace the transceiver technology model by a simplified version in simulation. The transceiver encapsulation block can be configured to implement the transceiver channels using a generic simulation model. This allows simulating a complete design without using the detailed simulation model of the transceiver provided by the FPGA vendor. The benefit is faster simulation times when using the specific technology transceiver model is not a requirement, for example for system-level simulations of the design.
- 3) Allows directly porting a design to any of the supported FPGAs. The only thing to do when switching technologies is to update a configuration package file with the corresponding technology values. The same RTL instantiation can be reused without any further change.




Figure 1. SpaceFibre IP Transceiver Interface Interconnection Diagram

3. SPACEFIBRE INTERFACE IP CORE

The STAR-Dundee SpFi IP core family has been extensively tested in different space-grade FPGA families. These IPs are provided with a protocol agnostic data interface, so that no prior knowledge of the SpFi standard is required. Simple data interfaces based on standard 32-bit input and output FIFO interfaces following the AXI4-S protocol are used [12], with independent user-defined read and write clocks. The IPs are configured using generics, e.g. transceiver interface, target technology for direct memory instantiation (ECC support in internal block RAMs), number and size of VCs, etc. A management interface allows real-time configuration of the IP control and status parameters, also including optional statistics and debug signals. Power management options have

been considered. For example, it is possible of start one end of the link in a low-power mode, waiting for the other end to become active.

Multi-lane (ML) is an optional capability of a SpFi link. The ML coordinates the operation of multiple lanes acting as a single SpFi link, providing higher data throughput and redundancy. SpFi Single-Lane (SL) and ML Interface IPs are fully interoperable. Whenever a SL link is connected to a ML link, the ML detects this condition during initialisation and adapts seamlessly. Moreover, ML links implementing different number of lanes are also fully compatible, with the link with the greater number of lanes adapting to the link with fewer lanes. Links presenting different internal data path widths at each end are also compatible.

The number of lanes in a link is fully configurable, with any number of lanes supported (up to 16). Each lane can independently be selected as uni/bidirectional and hot/warm redundant. In case of lane failure in a link without redundant lanes, the link is automatically reconfigured to continue with the remaining working lanes, hence producing an automatic graceful degradation of the link bandwidth. The QoS mechanism ensures that the most important data is sent first, i.e. higher priority VCs or scheduled traffic are less affected. These capabilities are very useful for space applications where strict power constrains and a high level of reliability are required.

The QoS is independently and dynamically configurable for each VC. The link FDIR mechanisms automatically recover from transient, persistent and permanent (when ML is used) errors on the SpFi lanes. A transient error takes less than 3 usec to recover and does not affect the user data rate thanks to the embedded buffering inside the IPs. Other protections against errors include data and broadcast babbling node protection. A lane is automatically disconnected when the BER is worse than 10^{-5} to prevent a potential protocol breakdown.

Figure 2 shows the types of host and transceiver interfaces available in the SpFi IP.

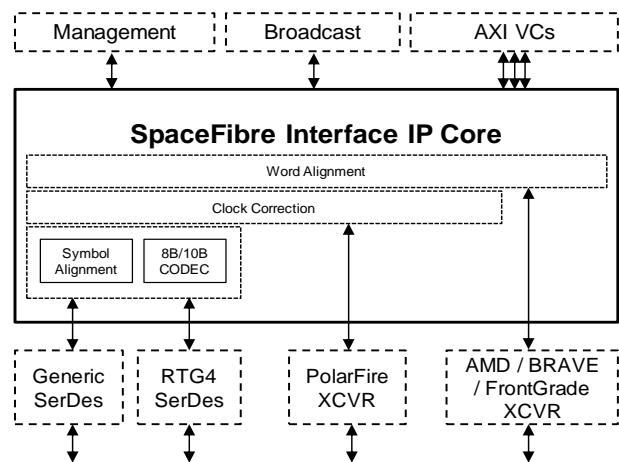


Figure 2. SpaceFibre IP Interfaces

Implementation

The resources required by the SpFi SL Interface IP for different number of VCs are detailed in Table 3. Table 4 provides resource values for SpFi ML Interfaces using a combination of different number of lanes and VCs. Synplify Elite v2024.09 has been used to obtain the results provided in Tables 3, 4 and 5. Note that the usage values in all the tables in this paper include the transmit and receive FIFOs. These tables show as the IP produces a compact block, only requiring a small percentage of fabric resources even when multiple lanes and VCs are used.

Table 3. SpFi Single-Lane Interface Resource Usage

	RTG4			XQRKU060 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1 VC	3765 2.5%	2683 1.8%	4 1.9%	1954 0.6%	2151 0.3%	4 0.4%
2 VCs	4403 2.9%	3139 2.1%	6 2.9%	2452 0.7%	2611 0.4%	6 0.6%
4 VCs	5622 3.7%	4039 2.7%	10 4.8%	3289 1.0%	3491 0.5%	10 0.9%

	RTPFS460ZT *			XQRVC1902 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1 VC	3178 0.7%	2226 0.5%	8 0.5%	1797 0.2%	2151 0.1%	4 0.4%
2 VCs	3846 0.8%	2758 0.6%	12 0.8%	2164 0.2%	2604 0.1%	6 0.6%
4 VCs	5145 1.1%	3797 0.8%	20 1.4%	2874 0.3%	3491 0.2%	10 1.0%

	NG-Ultra			LFCPNX-100 *		
	LUT	DFP	BRAM	LUT	DFP	EBR
1 VC	2523 0.5%	2308 0.5%	8 1.2%	2074 2.6%	2405 3.0%	8 3.8%
2 VCs	3035 0.6%	2913 0.6%	12 1.8%	2463 3.1%	3002 3.8%	12 5.8%
4 VCs	4049 0.8%	4098 0.8%	20 3.0%	3266 4.1%	4179 5.2%	20 9.6%

* TMR not included



Figure 3. SpaceFibre Multi-Lane Interface on a PolarFire

The timing estimations provided by the synthesis tools are not accurate, as the final timing depends on the routing and placement of the IP inside the FPGA fabric. However, testing

these IPs in different configurations on different boards have confirmed that maximum lane speeds can be achieved with RTG4 (3.125 Gbit/s) even in congested designs. Other supported FPGAs achieve lane rates exceeding 6.25 Gbit/s, with additional margin, enabling operating at these high speeds even with TMR applied. Note that the NG-Ultra maximum speed is an ongoing work, as the tools are still improving.

Figure 3, Figure 4 and Figure 5 show the SpFi Interface IP implemented in the development boards of different FPGAs.

Table 4. SpFi Multi-Lane Interface Resource Usage

	RTG4			XQRKU060 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
2 Ln 1 VC	6711 4.4%	5325 3.5%	8 3.8%	3400 1.0%	4002 0.6%	8 0.7%
2 Ln 4 VCs	8760 5.8%	7331 4.7%	20 9.6%	4898 1.5%	5778 0.9%	20 1.9%
4 Ln 1 VC	12670 8.3%	9531 6.3%	16 7.7%	5750 1.7%	6726 1.0%	12 1.1%
4 Ln 4 VCs	15168 10.0%	12195 8.0%	40 19.1%	7585 2.3%	9374 1.4%	30 2.8%

	RTPFS460ZT *			XQRVC1902 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
2 Ln 1 VC	5337 1.2%	4368 0.9%	12 0.8%	3194 0.4%	4008 0.2%	8 0.8%
2 Ln 4 VCs	7584 1.6%	6590 1.4%	30 2.1%	4386 0.5%	5778 0.3%	20 2.1%
4 Ln 1 VC	9909 2.1%	7606 1.6%	20 1.4%	5456 0.6%	6726 0.4%	12 1.2%
4 Ln 4 VCs	12529 2.7%	11135 2.4%	50 3.4%	6980 0.8%	9375 0.5%	30 3.1%

	NG-Ultra			LFCPNX-100 *		
	LUT	DFP	BRAM	LUT	DFP	EBR
2 Ln 1 VC	5670 1.1%	5719 1.1%	16 2.4%	4084 5.1%	4812 6.0%	12 5.8%
2 Ln 4 VCs	7563 1.4%	8069 1.6%	40 6.0%	5303 6.6%	7453 9.3%	30 14.4%
4 Ln 1 VC	10811 2.0%	10086 2.0%	32 4.8%	7085 8.9%	8342 10.4%	20 9.6%
4 Ln 4 VCs	13032 2.4%	13343 2.6%	80 11.9%	8713 10.9%	12719 15.9%	50 24.0%

* TMR not included

4. SPACEFIBRE ROUTING SWITCH IP CORE

The Routing Switch architecture is built around a non-blocking routing switch matrix with a configurable number of ports. Ports can be either SpFi, SpW or AXI4-S interfaces. Each port implements a configurable number of VCs, and each VC has an associated VN number. The switch matrix interconnects one or more VCs with the same VN number, with the limitation that each of these VCs must be in a different port. The output port is selected using path or logical addressing, indicated by the leading byte of each packet and the configuration of the internal routing table. Packets belonging to different VNs never interfere with one another and do not impact the throughput and latency within the routing switch matrix. On the other hand, when multiple

packets in the same VN need to be transferred to the same output port, packet-by-packet round-robin arbitration is performed, exactly like in a SpW Routing Switch.



Figure 4. SpaceFibre Single-Lane Interface on a BRAVE FPGA




Figure 5. SpaceFibre Single-Lane Interface on a CertusPro-NX

Implementation

The STAR-Dundee SpFi Routing Switch IP is a scalable, fully configurable non-blocking switch. The IP is very flexible, allowing to select the number of VCs, ports and target technology, among other options, using generics. The SpFi lane rates are also configurable. The Routing Switch IP presents a deterministic low latency switching, implementing path and logical addressing, VNs, time distribution and message broadcasting. VNs can be statically (fixed using VHDL constants) or dynamically configured, with up to 64 different VNs supported. There are 256 broadcast channels

with higher priority for time-critical broadcast messages. The Routing Switch offers a simple and efficient integration with SpW networks using SpW packet buffers and automatic SpW to SpFi broadcast translation. An internal timer tracks time being distributed over the network.

The configuration port uses the RMAP protocol [13] to configure the routing table, VNs, ports, etc. An internal AXI4-Lite master [10] allows accessing an external memory space using the extended address (0x1) capability of RMAP. Whenever the configuration port receives an RMAP command with extended address 0x1, the register access is carried out via this external AXI4-Lite interface. This can be useful when integrating the SpFi Routing Switch IP in a design with a companion instrument, for example. Another AXI4-Lite interface (slave) is available and allows doing the opposite access: accessing the internal memory space of the routing switch from an external device. This can be useful to configure the routing switch from a companion CPU. Figure 6 shows a simplified Routing Switch block diagram. Note that SpW ports only have associated a single VC.




Figure 6. SpaceFibre Routing Switch Block Diagram

The configuration of the Routing Switch is done via a configuration package. This package allows defining the default VN mapping (loaded immediately after reset) and the number and types of ports. For example, the following configuration would produce the routing switch depicted in Figure 7.

```

constant SPFI_PORT_LANES : integer := 1;
constant SPFI_PORTS      : integer := 4;
constant SPW_PORTS       : integer := 0;
constant AXI_PORTS       : integer := 0;

constant INIT_VN : [...]
-- For each port, the VC number used by each VN
--
-- VN      P1,P2,P3,P4,
0  => (0, 0, 0, 0),
1  => (1, 2, N, 1),
2  => (N, 1, N, 2),
others => (others => N));

```

This simple example illustrates how the control network (blue path) and two instrument data flows (green and yellow paths) can be assigned to VNs. Note that the VN configuration can be modified after reset using the configuration port or the management interface.

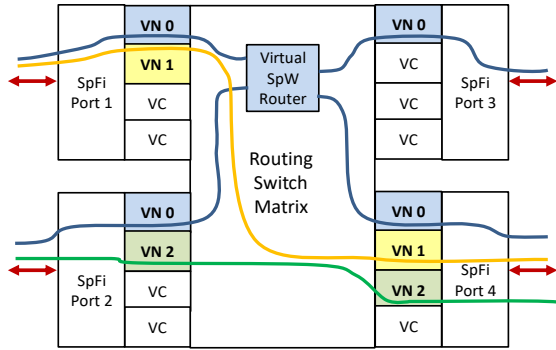


Figure 7. SpFi Routing Switch VN Configuration Example

Figure 8 shows the SpFi SL Routing Switch IP running on the Versal VCK190 board using four SpFi ports connected to a STAR-Ultra PCIe Interface unit.



Figure 8. SpaceFibre Single-Lane Routing Switch on a Versal

Table 5 presents the resource usage for two reference Routing Switch configurations implementing SpFi ports only: a 4-port Routing Switch, each port supporting 2 VCs (4P 2 VC), and an 8-port Routing Switch, each port with 4 VCs (8P 4 VC). Examples with SpFi ports with 1, 2 and 4 lanes are provided. The table shows that even a “large” Switch of 8 SpFi single-lane ports with 4 VCs (1L8P 4 VCs) fits inside an RTG4. Note that the SpFi Interface logic of the ports is included in the table figures, as well as the additional RMAP configuration port and all the configuration logic.

There are options to reduce the resource usage by tailoring the behavior of the VNs, e.g. setting static VNs. Additionally,

for some applications, provided that the network requirements are fixed, it is possible to limit the number of ports that can be accessed from a certain VN. Finally, additional resource savings will be provided with a new revision of the IP currently under development.

Table 5. SpFi Routing Switch Resource Usage

	RTG4			XQRKU060 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1L4P	31405	27740	36	19920	25516	38
2 VCs	20.7%	18.3%	17.2%	6.0%	3.8%	3.5%
1L8P	98236	78521	120	65181	74059	122
4 VCs	64.7%	51.7%	57.4%	19.7%	11.2%	11.3%
2L4P	48448	43402	63	29629	37966	63
2 VCs	31.9%	28.6%	30.1%	8.9%	5.7%	5.8%
2L8P	141388	113221	203	90655	102290	203
4 VCs	93.1%	74.6%	97.1%	27.3%	15.4%	18.8%
4L4P	80624	69090	113	46172	57667	88
2 VCs	53.1%	45.5%	54.1%	13.9%	8.7%	8.1%
4L8P	N/A	N/A	N/A	134571	148359	284
4 VCs				40.6%	22.4%	26.3%

	RTPFS460ZT *			XQRCV1902 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1L4P	28908	26229	68	18055	25572	38
2 VCs	6.3%	5.7%	4.7%	2.0%	1.4%	3.9%
1L8P	94061	76684	232	60430	74112	122
4 VCs	20.4%	16.6%	15.9%	6.7%	4.1%	12.6%
2L4P	43253	40119	95	27690	38022	63
2 VCs	9.4%	8.7%	6.5%	3.1%	2.1%	6.5%
2L8P	132823	109042	315	85138	102339	203
4 VCs	28.8%	23.7%	21.6%	9.5%	5.7%	21.0%
4L4P	70110	62728	145	43820	57722	88
2 VCs	15.2%	13.6%	9.9%	4.9%	3.2%	9.1%
4L8P	203106	163251	477	127717	148419	284
4 VCs	44.1%	35.4%	32.7%	14.2%	8.2%	29.4%

	NG-Ultra			LFCPNX-100 *		
	LUT	DFP	BRAM	LUT	DFP	EBR
1L4P	22087	26615	72	19072	27519	76
2 VCs	4.1%	5.3%	10.7%	23.9%	34.5%	36.5%
1L8P	70543	77339	240	N/A	N/A	N/A
4 VCs	13.1%	15.3%	35.7%			
2L4P	38002	44625	126	30796	42975	95
2 VCs	7.1%	8.8%	18.8%	38.6%	53.8%	45.7%
2L8P	110357	117070	406	N/A	N/A	N/A
4 VCs	20.6%	23.2%	60.4%			
4L4P	64396	70405	226	48046	67069	142
2 VCs	12.0%	13.9%	33.6%	60.2%	84.0%	68.2%

* TMR not included
SpFi Interface IP resources are included

Next-Gen Communication Architecture Use Case

The Hi-SIDE project (Innovation in Space Award 2022 at the European Space Forum) was a European Union project carried out by several leading aerospace organisations from across Europe. It aimed to develop satellite data-chain technologies for future Earth Observation and Telecommunication systems. The data chain elements were interconnected via a SpFi network.

The SpFi Multi-Lane Routing Switch IP was the main component in the design of the primary element of the Hi-

SIDE project: the STAR Tiger, a 10-port Routing Switch with 4-lane (25 Gbit/s) and 2-lane (12.5 Gbit/s) SpFi ports [14]. Figure 9 shows a photograph of the STAR-Tiger SpFi Routing Switch, with a size of 108 x 108 x 70 mm and a power consumption of 13.5 W (20 °C).



Figure 9. STAR-Tiger SpaceFibre Multi-Lane Routing Switch

In Hi-SIDE, the different data chain elements were interconnected using a SpFi ML Routing Switch (Figure 10). This switch, implemented in a Kintex UltraScale FPGA, was used for transferring data at very high data-rates between instruments, mass-memory, data compressor, data processor and downlink transmitters. It was also used to provide the control network, which the control computer employed to manage both the network, and the equipment attached to it.

The various elements of the on-board data chain are shown on the left and right sides of Figure 10, interconnected via SpFi links (red lines) to the routing switch in the middle. Each element contains a SpFi interface which connects it to the SpFi network. The interfaces are either quad-lane or dual-lane interfaces and have two or more VCs which are mapped to VNs by the switch. The VNs are colour coded and a key to the type of traffic they handle is shown. The configured VNs are also illustrated inside the SpaceFibre routing switch (X equals to VN switch).

5. OPERATION UNDER RADIATION

The SpaceFibre Interface IP Cores have been specifically designed for space applications, including the fact that they shall operate under radiation. Not all radiation-tolerant FPGA technologies present the same robustness against radiation and, as a consequence, additional mitigation measures may be required to meet the reliability goal for a design on a particular FPGA and operational conditions.

Triple Modular Redundancy (TMR) Insertion

The use of TMR is advised on most of the technologies discussed in Section 1, with the exception of RTG4 and NG-Ultra. Depending on the technology, two types of TMR are recommended: Local TMR (LTMR) when the configuration memory is not sensitive to radiation (e.g. PolarFire) and Distributed TMR (DTMR) when the configuration memory is susceptible to SEUs (e.g. Versal or CertusPro-NX-RT).

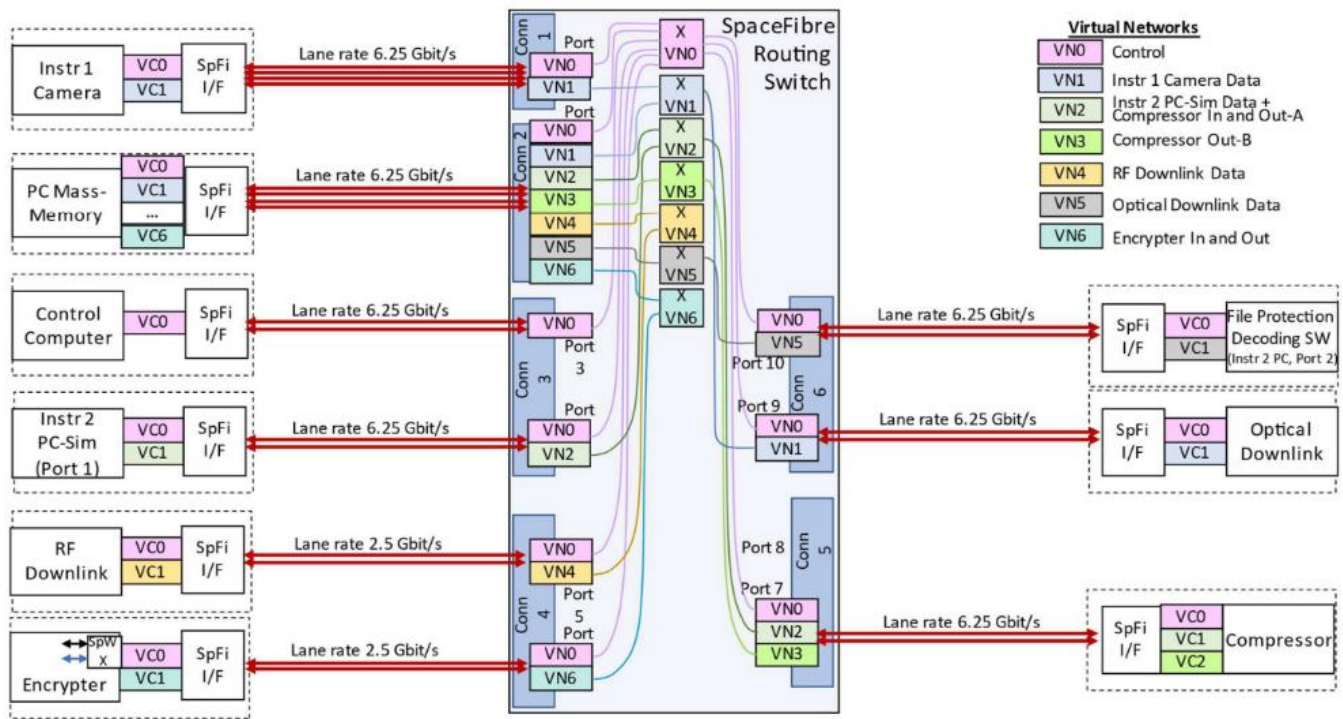


Figure 10. STAR-Tiger Virtual Networks Used by Hi-SIDE

Table 6. SpFi IPs Usage After TMR is Applied

	RTPFS460ZT (LTMR)			XQRVC1902 (DTMR)		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1 Ln	9868	10694	20	20006	10143	30
4 VCs	2.1%	2.3%	1.4%	2.2%	0.6%	3.1%
2 Ln	16830	19016	30	36437	18894	60
4 VCs	3.7%	4.1%	2.1%	4.0%	1.0%	6.2%
4 Ln	28326	32546	50	61492	31557	90
4 VCs	6.1%	7.1%	3.4%	6.8%	1.8%	9.3%
1L4P	61680	73738	68	128218	74428	114
2 VCs	13.4%	16.0%	4.7%	14.2%	4.1%	11.8%
2L4P	95461	114698	95	206255	118890	189
2 VCs	20.7%	24.9%	6.5%	22.9%	6.6%	19.5%
4L4P	152909	181156	145	323283	184600	264
2 VCs	33.2%	39.3%	9.9%	35.9%	10.3%	27.3%

Table 6 shows some IP resource usage examples (SpFi Interface and Routing Switch IPs) for two representative FPGA technologies, PolarFire using LTMR and Versal using DTMR. These results have been obtained with Synplify Pro v2023.09 (PolarFire) and Synplify Elite v2024.09 (Versal) and can be compared against the nominal values presented in Tables 3, 4 and 5. LTMR produces an increase of the resource usage by a factor of ~2x for LUTs and ~3x for registers. Block BRAM (BRAM) use is not affected. On the other hand, for DTMR the increase factor of the resource usage is ~7-8x for LUTs, and ~3x for registers and BRAMs. Note that these factors are mainly dependent on the nature of the TMR process and not so much to the IP architecture. Therefore, especially when DTMR is needed, it is crucial to keep complexity down. Otherwise, even for large devices such as the Versal XQRVC1902, a design can quickly fill the target device when DTMR is applied. This illustrates why the small footprint of the SpFi IPs is a key advantage in space applications.

It is worth noting that timing closure was successfully achieved after applying LTMR/DTMR in both PolarFire and Versal. The designs tested featured 25 Gbit/s SpFi 4-lane links.

Heavy-Ion Radiation Testing Results

Several heavy-ion radiation campaigns have been carried out in collaboration with Microchip and AMD to validate the IPs operation when implemented in different FPGAs: RTG4, PolarFire and, very recently, Versal devices have been tested. The main goal was to characterise the effects of Single Event Effects (SEEs) on the FPGA transceivers, and to assess the resilience of SpaceFibre under these conditions.

For RTG4, the observed saturation cross-section value for each transceiver lane was approximately 10^{-5} cm² (LET > 20 MeV*cm²/mg). The protection provided by the SpFi protocol was demonstrated, effectively mitigating the majority of SEEs on the data. In most cases, these effects were resolved transparently and without significant impact

on the application, owing to the rapid recovery mechanism of the protocol [5].

In the PolarFire case, the testing was conducted in collaboration with Microchip and an industrial partner. While the specific results remain confidential, the outcome was encouraging.

For Versal, a recent radiation campaign has been recently carried out at GANIL (France) and its results are still under analysis, with some preliminary results presented below. The effective LET used was high, reaching beyond 40 MeV*cm²/mg. Figure 11 shows the test board with an unlidged Versal in position for one of the runs.




Figure 11. Versal Board Placed Inside the Radiation Chamber at GANIL

The preliminary saturation cross-section for a quad transceiver (XCVR) is approximately 10^{-4} cm². This value accounts for various error cases affecting components such as the TXPLL and the data path. To enable data transmission using this XCVR, a protocol must be implemented in the FPGA fabric. However, both the fabric and its configuration memory (CRAM) are vulnerable to SEUs, adding an additional challenge.

In the tests, SpFi consistently recovered the XCVR automatically in all but two observed outliers, which are still under investigation. When DTMR is applied, SpFi link recovery occurs automatically without user data errors, as DTMR mitigates the risk of data corruption. In such cases, the only impact of a SEE on the link is a retry event, typically completed in less than 4 μs.

Without DTMR, the small implementation footprint of SpFi reduces its susceptibility to SEEs when compared to other protocols. SEEs can still be corrected using the FPGA inbuilt scrubbing mechanism (XiISEM), which typically reacts within tens of milliseconds. While data loss might occur in such cases, the SpFi link generally self-recovers, ensuring system functionality remains intact.

The time required for data recovery depends on the specific XCVR component affected. The XCVR recovers within 160 ns to 3 μ s. Approximately 99% of error events observed were transient, with recovery times under 4 μ s (SpFi retry event). Persistent errors, which occurred in about 1% of the events, required a reset of the XCVR receive logic. This process is automatically managed by the SpFi protocol and takes approximately 2 ms to complete. The two outlier events, representing approximately 0.3% of measured events, required FPGA reprogramming for recovery.

Additionally, the campaign successfully demonstrated 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.

6. CONCLUSIONS

The STAR-Dundee SpaceFibre Interface and Routing Switch IP Cores have been designed to be easy to implement in radiation-tolerant FPGAs. The host interface to the user is a simple AXI4-Stream for each virtual channel where the user can push arbitrary packet content of any size, appearing at the AXI-Stream far-end interface. The connection of the SpFi IP to the transceiver has been simplified with the encapsulation of all transceiver technology dependent logic into a generic transceiver wrapper. This allows switching among different target FPGAs or the simulation model with ease.

We have detailed the performance and capabilities of the different IP Cores, and discussed the resources required depending on several parameters, namely the number of VCs, lanes and ports. The compactness and simplicity of the IPs allows for integration into existing designs with minimal area impact, enabling quick and efficient addition of high-speed links to a design. All these IP cores have undergone extensive verification through simulation and have been successfully validated across multiple hardware prototypes.

Two separate heavy-ion campaigns have been successfully carried out to assess the operation of the SpFi Interface IPs in the RTG4, plus an additional one for PolarFire. In Q4 2024 another campaign to validate SpFi running on Versal was completed, with very encouraging preliminary results. The combination of SpFi and DTMR provides a high-performance, highly reliable solution for mitigating radiation-induced errors.

These IPs provide the all the necessary building blocks for creating next generation of onboard networks. This has been demonstrated in the Hi-SIDE project, an award-winning European Union project involving several European

aerospace organisations that have developed satellite data-chain technologies for future Earth observation and telecommunication systems. The different elements of the data chain were interconnected via a SpFi network.

SpFi is at TRL 9 – already in operation in at least 6 spacecrafts, and it is also currently being implemented in FPGA and ASIC designs by different missions and products all over the world. Support for higher lane rates will be provided soon, with SpFi running at 100 Gbit/s (25 Gbit/s per lane) already successfully tested under radiation, thus enabling multi-lane links running at 200 Gbit/s in currently available technology in the near future.

ACKNOWLEDGEMENTS

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

We acknowledge the GANIL team for their support during the radiation test campaign.

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] ECSS Standard ECSS-E-ST-50-12C, “SpaceWire, Links, Nodes, Routers and Networks”, Issue 1, European Cooperation for Space Data Standardization, July 2008, available from <http://www.ecss.nl>.
- [4] A. Gonzalez Villafranca, A. Ferrer Florit, M. Farras Casas, S. Parkes and C. McClements, “SpaceFibre for FPGA: IPs and Radiation Test Results”, SEE-MAPLD 2020
- [5] A. Gonzalez Villafranca, A. Ferrer Florit, M. Farras Casas and N. Rezzak, “Mitigation and Recovery of Single Event Effects in RTG4 FPGA Transceiver Links Using SpaceFibre”, RADECS 2022.
- [6] Microchip, “DS5395 Radiation-Tolerant PolarFire® SoC FPGA Product Overview”, Microchip, 2024, available from <https://ww1.microchip.com/downloads/aemDocuments/documents/FPGA/ProductDocuments/ProductBrief/Radiation-Tolerant-PolarFire-SoC-FPGA-Product-Overview-DS00005395.pdf> (last visited 4th October 2024).
- [7] 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 4th October 2024).
- [8] FrontGrade, “UT24CP1008 CertusPro-NX-RT Family Datasheet”, ver. 2.0.2, 12th May 2023, available from https://www.frontgrade.com/sites/default/files/documents/datasheet-ut24cp1008_0.pdf (last visited 4th October 2024).
- [9] NanoXplore, “NG-Ultra NX2H540TSC Datasheet”, ver. 2.1, NanoXplore, June 2022, available from <https://files.nanoxplore.com/f/5785df61b4f545a08ae8/?dl=1> (last visited 4th October 2024).
- [10] Microsemi, “RTG4 FPGAs Technical Brief”, rev D, Microsemi, January 2024, available from <https://ww1.microchip.com/downloads/aemDocuments/documents/FPGA/ProductDocuments/UserGuides/RTG4-FPGA-Technical-Brief.pdf> (last visited 4th October 2024).
- [11] AMD, “Radiation Tolerant Kintex UltraScale XQRKU060 FPGA Data Sheet DS882”, v1.4, AMD, 10th September 2024, available from <https://docs.xilinx.com/v/u/en-US/ds882-xqr-kintex-ultrascale> (last visited 4th October 2024).
- [12] ARM, “AMBA AXI Protocol Version 2.0 Specification”, ARM, Ed., 2010.
- [13] ECSS Standard ECSS-E-ST-50-52C, “SpaceWire – Remote Memory Access Protocol”, European Cooperation for Space Data Standardization, 5th February 2010, available from <http://www.ecss.nl>.
- [14] S. Parkes et al., “SpaceFibre Payload Data-Handling Network”, 2022 International SpaceWire Conference.

BIOGRAPHY



Alberto Gonzalez Villafranca holds a PhD in data compression for space applications and has been connected to the space field throughout his entire professional career. He began by collaborating with the Gaia mission and subsequently worked on the hardware implementation of a deterministic variant of the SpaceWire protocol at the European Space Agency. Since then, Alberto has been deeply involved in the definition and implementation of SpaceFibre for more than 10 years, first with STAR-Dundee Ltd and later with its subsidiary, STAR-Barcelona, as a Chip Designer.



Albert Ferrer Florit has a PhD in high-speed interconnection networks for space applications awarded by the University of Dundee. His PhD research was funded by ESA's Networking/Partnering Initiative after he worked in the on-board data processing group (TEC-EDP) in ESTEC. He is specialised in SpaceWire and SpaceFibre networks, being one of the key developers of the SpaceFibre standard. He started his career at CERN in the Summer Student Programme, worked for STAR-Dundee and is currently working for STAR-Barcelona as a Network and FPGA Engineer.



Marti Farras Casas holds a BSc. in Electronics and already began working in the space industry while at university, first with an internship on ESA's Gaia mission, followed by his degree thesis developed at STAR-Dundee Ltd. For over eight years, Marti has been deeply involved with the development of SpaceFibre, first with STAR-Dundee Ltd and later with its subsidiary, STAR-Barcelona, as a Hardware Design Engineer.



***Steve Parkes** received a BSc. and MSc. in Applied Physics and Electronics from Lancaster University, and a PhD from University College Cardiff. He is a Fellow of the Royal Academy of Engineering. He worked on Underwater Acoustic Systems for six years and at BAE Space*

Systems for seven years before moving to Academia. He worked at Dundee University for 24 years where he became Professor of Spacecraft Electronic Systems carrying out research on spacecraft onboard networks, vision-based navigation and related planet surface simulation, and signal processing. He is the author of the widely used ECSS SpaceWire and SpaceFibre standards, with inputs from international engineers. In 2002 he founded STAR-Dundee, which supports industry and space agencies using SpaceWire and SpaceFibre technology, where he is now Chief Technology Officer.