

SpaceFibre IP Cores for the Next Generation of Radiation-Tolerant FPGAs

Alberto Gonzalez Villafranca
STAR-Barcelona SL

Sant Cugat del Valles, Barcelona, Spain
alberto.gonzalez@star-dundee.com

Albert Ferrer Florit
STAR-Barcelona SL

Sant Cugat del Valles, Barcelona, Spain
albert.ferrer@star-dundee.com

Marti Farras Casas
STAR-Barcelona SL

Sant Cugat del Valles, Barcelona, Spain
marti.farras@star-dundee.com

Steve Parkes
STAR-Dundee

Dundee, Scotland, UK
steve.parkes@star-dundee.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 provides point-to-point and networked interconnections at Gigabit rates with Quality of Service (QoS) and Fault Detection, Isolation and Recovery (FDIR). SpaceFibre has been designed as a replacement of SpaceWire (ECSS-E-ST-50-12C)—it is backwards compatible with SpaceWire at the packet level—for next-generation space missions where very high throughput is required, providing up to 6.25 Gbit/s per lane, with multi-lane allowing to reach up to 16 times the speed of a single lane. NORBY and OPS-SAT technology demonstrators have already flown SpaceFibre, with more missions in both Europe and the USA currently designing or planning to use SpaceFibre.

STAR-Dundee has developed a complete family of SpaceFibre IP cores fully compliant with the SpaceFibre standard. This family is composed of four different IPs: Single-Lane Interface, Multi-Lane Interface, Single-Lane Routing Switch and Multi-Lane Routing Switch.

A new generation of radiation-tolerant FPGAs is emerging to cope with the ever-growing processing power required by newer missions. Microchip has released the PolarFire RTPF500, Xilinx the Versal XQRVC1902, and NanoXplore the BRAVE NG-Ultra. SpaceFibre operation requires serial transceivers, which are already inbuilt in modern FPGAs. The IPs have been adapted to take advantage of the specific transceivers and memory blocks offered by these new FPGAs.

In this work we analyse in detail the performance of STAR-Dundee SpaceFibre IP cores on this new generation of FPGAs considering several performance metrics, e.g. maximum lane speed, resource usage, etc. We also compare the performance of the IPs with current state-of-the-art space-grade FPGAs, i.e. Microchip RTG4 and Xilinx Kintex UltraScale XQRKU060. This analysis can also be used as a representative benchmark to compare the performances of the different FPGAs available for space.

Keywords—SpaceFibre, Interface, Routing Switch, IP Cores, PolarFire RTPF500T, Versal XQRVC1902, BRAVE, NG-Large, NG-Ultra, RTG4, XQRKU060

I. INTRODUCTION

SpaceFibre (SpFi) [1] is a communication technology for use onboard spacecraft which was released as an ECSS 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 interoperates seamlessly with a SpaceWire (SpW) [2] network over virtual

channels (VCs) as it uses the same data packet definition. Furthermore, SpFi provides broadcast capabilities and can operate over either copper or fibre optic cables. 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.

The Network layer in SpFi is responsible for transferring data packets over a link or network. The information to be sent uses the SpW format: <Destination Address> <Cargo> <End of Packet Marker>. The routing concepts are the same as in SpW including both path and logical addressing. The Network layer includes the definition of Virtual Networks (VN). These VNs are built from the interconnection between VCs of different ports. VNs enable the creation of flexible SpFi routing switches (also known as Routers) comprising SpFi interfaces and a fully configurable, non-blocking, high performance, routing switch. This routing switch typically supports up to 64 VNs, each VN effectively behaving like independent SpW networks capable of working at multi-Gbps rates.

STAR-Dundee has developed a range of SpFi IP cores compliant with the standard. The range is composed of four different IPs: Single-Lane Interface (SL Intf), Multi-Lane Interface (ML Intf), Single-Lane Router (SL Router) and Multi-Lane Router (ML Router). These IPs have been optimised for speed considering the timing constraints of the slower FPGAs for space. The family of SpFi IPs is also compatible with commercial FPGAs such as Microchip SmartFusion2 and PolarFire, or Xilinx 7-series, UltraScale, Versal, etc. The SL Intf IP has already been tested in orbit in two demonstrator missions: NORBY and OPS-SAT [3]. These collaborations have demonstrated operational SpFi links in space, thus providing fly heritage for this technology.

This paper is a follow-up from a previous paper presented in the 2018 International SpW Conference which analysed the existing SpFi IP cores performance in the state-of-the-art FPGAs at the time, the Microchip RTG4 and the Xilinx Virtex5-QV [4]. For this new work the ML Router IP has been added to the IP analysis. On the other hand, a new generation of radiation-tolerant FPGAs has emerged to cope with the growing processing power required by newer missions since

the paper was published in 2018. The RTG4 has been kept in the analysis as reference for non-volatile legacy FPGA. The volatile reference has been updated from Virtex5-QV to the newer and more powerful Xilinx XQRKU060.

This paper describes the new generation of space-qualified FPGAs in section II. Section III analyses the main features of the STAR-Dundee SpFi IP Cores. Sections IV, V, VI and VII focus on the specific features and performance analysis of the SL Intf, ML Int, SL Router and ML Router IPs respectively. Finally, conclusions are presented in section VIII.

II. THE NEW GENERATION OF FPGAS FOR SPACE

There are two main devices that belong to the previous generation of space-qualified FPGAs. Both devices offer inbuilt high-speed transceivers that allow for the different SpFi IPs to be implemented.

On the one hand, there is the non-volatile configuration memory Microchip RTG4 manufactured in a low power 65 nm process [5], which is radiation-hardened by design. Thus, the big advantage of using the RTG4 is that Triple Module redundancy (TMR) has been integrated in its fabric transparently to the user. TMR is a method consisting of using triple module redundancy or triple voting to implement registers. Each register is implemented by three flip-flops that “vote” to determine the final output signal of the register function. Using TMR increases the number of resources used by an IP, affecting area and potentially also timing because of the additional logic inserted.

On the other hand, there is the SRAM-based (volatile) Xilinx XQRKU060 [6] manufactured on a faster and more compact 20 nm process. This FPGA offers roughly four times the resources of the RTG4 but does not offer the same degree of hardening against radiation.

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 even higher performance applications, which makes them ideal targets for using very high-speed communications protocols such as SpFi.

A. Microchip PolarFire

The radiation-tolerant PolarFire (PF) RTPF500T FPGA is directly derived from its commercial counterpart, a non-volatile FPGA built on a 28 nm process [7]. 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. Like the RTG4, PF provides 24 transceivers but with a higher maximum speed, each capable of running up to 10 Gbit/s. Unlike radiation-hardened devices, depending on the application the Single Event Upset (SEU) rate of the PF registers may not be good enough. In this case, the use of some form of TMR is advised.

B. Xilinx Versal

The Xilinx Versal XQRVC1902 [8] is the radiation-tolerant version of the commercial SRAM-based XCVC1902 FPGA. It is manufactured in a 7nm FinFET technology and provides a platform aiming at high performance applications, offering 44 GTY transceivers, each capable of running at more than 25 Gbit/s. Using TMR in Versal is also advised depending on the application and its requirement for register SEU sensibility, as the FPGA is not radiation-hardened.

C. NanoXplore BRAVE

The NanoXplore BRAVE FPGA family is a the European addition to the available options of space-qualified devices. There are several members of the BRAVE family, although only the NG-Large [9] (65nm FD-SOI SRAM) and NG-Ultra [10] (28 nm FD-SOI SRAM) include the inbuilt SerDes blocks required by SpFi. Both devices are radiation hardened by design, which removes the need for TMR.

III. SPACEFIBRE IP CORES GENERAL FEATURES

The SpFi IP core family has been extensively tested in the different space FPGA families. IPs have been carefully designed to guarantee timing closure in all the temperature and voltage conditions required by the space devices, including EDAC in the memories and Single Event Transient (SET) filtering enabled—when available. These radiation mitigation techniques have an impact on the maximum speed of the designs and can potentially create problems to meet the targeted clock frequencies.

Effort has also been put to minimise the designer effort when adding the SpFi IP to a design. 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 are used. Specifically, they follow the AXI4-Stream (AXI4-S) protocol [11], which is a popular industry standard. This AXI4-S interface allows using independent user-defined read and write clocks, with clock synchronisation between user and SpFi IP clock domains managed by the IP. The AXI4-S width can be extended in 32-bit multiples in the ML Intf/Router IPs.

The IPs can be configured using generics. Different properties can be configured, e.g. transceiver interface, target technology for memory direct instantiation (for EDAC use), number and size of VCs, etc. Different high-speed transceiver interface options provide the set of signals to be directly connected to the selected transceiver. There are specific interfaces for the RTG4, PolarFire, 40-bit and 20-bit parallel, and Xilinx devices. Each of these takes into account whether 8B10B encoding/decoding, bit and symbol alignment, and clock correction can be done by the transceiver for better resource usage. Support for old FPGA technologies that require external transceivers is also provided through a dedicated TLK2711-SP (Wizardlink) [12] interface. A wrapper is supplied for each of the different transceiver interfaces for user convenience.

The QoS is independently and dynamically configurable for each VC, offering three mechanisms that work concurrently: scheduling, priority and bandwidth reservation. The FDIR mechanisms automatically recover from transient, persistent and permanent (when ML is used) errors on the SpFi link. A transient error takes less than 3 μ sec to recover. It 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.

A management interface allows real time configuration of the IP control and status parameters, also including optional statistics and debug signals. Two different types of management interfaces can be selected: AXI4-Lite and APB bus. A signal bus is also available in the interface IPs. The AXI4-Lite and APB bus have independent clock with clock

synchronisation managed inside the IP for convenience. Independent signals for each status and configuration field are useful when an FPGA design needs direct access to the IP. Accessing these fields over an AXI4-Lite/APB bus simplifies the interface for designs that use a CPU or want a centralised access point to several interfaces, for example. Power management options have been considered. For example, it is possible to start one end of the link in a low-power mode waiting for the other end to become active.

Two radiation testing campaigns have been carried out in collaboration with Microchip for the SpFi SL and ML Interfaces in the RTG4 [13, 14]. The information gathered during the test campaigns has allowed for assessing and further refining the robustness of the IPs under radiation and their associated RTG4 reference designs.

The IP Cores are also ready for ASIC implementation. For example, the SL Intf was used in the RC64 many-core DSP ASIC (12 SpFi interfaces) developed by Ramon.Space [15]. Other ASICs are currently under development implementing the ML Intf and the Router IPs.

Full TMR has been applied to the IPs in the PF to assess its impact in performance. As expected, the resulting synthesis takes ~2.8 times the number of initial registers (it is slightly below 3 because synchronisers are not TMR'ed). The number of LUTs is increased by a factor of ~2. One possibility to reduce the impact of TMR on the IP is to apply partial TMR. The idea is to only protect the most critical parts of the IP, in this case the control logic. This way, SEUs can induce sporadic data errors at the receiver, but the operation of the protocol itself is rugged against these events. This alternative is a compromise between full TMR and no TMR at all, and can be appropriate for certain applications which can tolerate a certain rate of data errors. This partial TMR is an ongoing development for the IPs here presented.

Finally, STAR-Dundee has adapted its SpFi IP Cores to be compatible with the BRAVE family. There is an ongoing activity to validate the operation of the IP inside the NG-Large FPGA. A successful SpFi link has been established, allowing the correct transmission of data between a STAR-Fire Mk3 unit and the NG-Large development board (Fig. 1). However, retry events were observed during the IP validation. The cause for these retries is probably related to the clock scheme adopted, which uses a fabric clock as the SerDes reference clock due to hardware limitations on the experiment set-up. A new set-up with an external SerDes reference clock is in the process of being tested. Nevertheless, it is worth highlighting that despite the retry events no data errors appeared. This is an example of the resilience of SpFi against errors on the link.

A. Timing Performance

Timing provided by the synthesis tools is not accurate, as the final timing depends on the routing and placement of the IP inside the FPGA fabric. Testing these IPs in different configurations on different development boards have confirmed that the maximum lane speeds can be achieved with the RTG4 (3.125 Gbit/s) even for congested designs. The rest of FPGAs achieve lane rates beyond 6.25 Gbit/s with plenty of margin, thus allowing to operate at these high speeds even when using TMR on the design. The figures shown in the tables of next sections all include transmit and receive FIFOs.

IV. SPACEFIBRE SINGLE-LANE INTERFACE IP CORE

The resources required by the SL Intf IP are detailed in Table I for a different number of VCs. As the table shows, the IP offers a compact design only requiring a small percentage of area for implementing a SpFi interface, even when multiple VCs are used. Even with full-TMR, the impact in area usage of a SpFi link will be limited. Note that adding an additional VC to the design has also a limited impact on the overall resource usage.

The IP resource usage for both NG-Large and NG-Ultra has been obtained with the latest tool release—NXMap v22.1.0.1. Usage is the same for both devices because the fabric of the FPGAs is essentially the same, so no differences are expected between them. Due to the continuous evolution of NanoXplore tools, timing results continue to improve although they are still trailing those of Microchip and Xilinx devices. Final results will be presented once the IPs have been fully validated in BRAVE.

TABLE I. SPFI SINGLE-LANE INTERFACE RESOURCE USAGE

	RTG4			XQRKU060*		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1 VC	3316 2.2%	2365 1.6%	4 1.9%	1823 0.5%	2346 0.4%	4 0.4%
2 VCs	3960 2.6%	2946 1.9%	6 2.9%	2162 0.7%	2969 0.4%	6 0.6%
4 VCs	5389 3.5%	4114 2.7%	10 4.8%	2960 0.9%	4214 0.6%	10 0.9%

	RTPF500T*			XQRVC1902*		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
1 VC	2796 0.6%	2226 0.5%	8 0.5%	1687 0.2%	2272 0.1%	2 0.2%
2 VCs	3400 0.7%	2801 0.6%	12 0.8%	1985 0.2%	2824 0.2%	3 0.3%
4 VCs	4653 1.0%	3972 0.8%	20 1.3%	2796 0.3%	3923 0.2%	5 0.5%

	NG-Large			NG-Ultra		
	LUT	DFP	RAM	LUT	DFP	RAM
1 VC	2703 2.0%	2496 1.9%	8 4.2%	2703 0.5%	2496 0.5%	8 1.2%
2 VCs	3275 2.4%	3068 2.4%	12 6.3%	3275 0.6%	3068 0.6%	12 1.8%
4 VCs	4350 3.2%	4220 3.3%	20 10.4%	4350 0.8%	4220 0.8%	20 3.0%

* TMR not included.



Fig. 1. BRAVE SpFi Interoperability Test with a STAR-Fire Mk3 unit.

V. SPACEFIBRE MULTI-LANE INTERFACE IP CORE

A. Specific Features

Multi-lane is an optional capability of the SpFi link. The Multi-Lane layer coordinates the operation of multiple lanes acting as a single SpFi link, providing higher data throughput and redundancy. Each lane can be initialized and operated independently from each other.

The number of lanes 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. SL implementations must be bidirectional even if the end-user data flow is unidirectional, because of the feedback required by the protocol. However, in a Multi-Lane implementation only one bidirectional lane is enough for the interchange of protocol related information. Therefore, other lanes can be unidirectional to save power and mass in asymmetric data flows. The width of the AXI4-S interface of the VCs is configurable in multiples of the SpFi word size (32-bits). This allows supporting slower user clocks and still being able to send or receive data at the maximum speed over a single VC.

Hot redundant lanes allow the link to fully recover not only from transient errors (like the SL Intf), but also from persistent or permanent lane failures in less than 3 μ s without user intervention and without any data loss. This 3 μ s time is close to the round-trip delay of the lane. 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. Warm redundant lanes save power with respect the always-on hot-redundant alternative, but they take around 20-40 μ s to reach a working state. Bandwidth overprovision and dynamic power management are also possible. These capabilities are very useful for space applications where strict power constrains and a high level of reliability is required on the harsh space environment.

Fig. 2 shows the ML Interface IP being tested in a PF with 2 lanes coming out of a STAR-Dundee SpW/SpFi FMC daughterboard. Similarly, Fig. 3 also shows a ML testing on Versal with a set-up allowing for up to 8 lanes out of the FPGA by using 2 QSFP+ connectors on an FMC daughterboard.

B. Area Resources

Table II provides the FPGA resource usage for a combination of different number of lanes (2, 4 and 8) and VCs (1, 2 and 4). Individual lanes can operate up to 3.125 Gbps in the RTG4 and in excess of 6.25 Gbit/s in the rest of devices. This means aggregate rates with 8 lanes of up to 25 Gbit/s in the RTG4 and 50+ Gbit/s in the other devices. The user data rate (removing 8B10B and protocol overheads) that can be achieved in a full-duplex 8-lane scenario in each direction is 18.5 Gbit/s for the RTG4 and 37 Gbit/s for the rest of FPGAs. Multi-lane is a convenient way of multiplying the link bandwidth. It provides additional advantages, e.g. graceful degradation, unidirectional operation, redundancy, that are automatically managed by the link without a big increase in resources.

TABLE II. SpFi MULTI-LANE INTERFACE RESOURCE USAGE

	RTG4			XQRKU060 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
2 Ln	6870	5166	8	3390	4771	8
1 VC	4.5%	3.4%	3.8%	1.0%	0.7%	0.7%
2 Ln	7792	6166	12	3928	5962	12
2 VCs	5.1%	4.1%	5.7%	1.2%	0.9%	1.1%
2 Ln	9492	8087	20	4948	7441	20
4 VCs	6.3%	5.3%	9.6%	1.5%	1.1%	1.9%
4 Ln	12776	9007	16	6020	7942	12
1 VC	8.4%	5.9%	7.7%	1.8%	1.2%	1.1%
4 Ln	13908	10497	24	6744	9259	18
2 VCs	9.2%	6.9%	11.5%	2.0%	1.4%	1.7%
4 Ln	15969	13390	40	8100	11792	30
4 VCs	10.5%	8.8%	19.1%	2.4%	1.8%	2.8%
8 Ln	27203	16739	32	12957	14305	20
1 VC	17.9%	11.0%	15.3%	3.9%	2.2%	1.9%
8 Ln	28565	19197	48	13995	16409	30
2 VCs	18.8%	12.6%	23.0%	4.2%	2.5%	2.8%
8 Ln	31076	24014	80	15976	20494	50
4 VCs	20.5%	15.8%	38.3%	4.8%	3.1%	4.6%

	RTPF500T *			XQRVC1902 *		
	LUT	DFP	LSRAM	LUT	DFP	RAMB36
2 Ln	4778	4332	12	3195	4611	8
1 VC	1.0%	0.9%	0.8%	0.4%	0.3%	0.8%
2 Ln	5681	5242	18	3699	5444	12
2 VCs	1.2%	1.1%	1.2%	0.4%	0.3%	1.2%
2 Ln	7357	6980	30	4629	6995	20
4 VCs	1.5%	1.5%	2.0%	0.5%	0.4%	2.1%
4 Ln	8619	7381	20	5828	7640	12
1 VC	1.8%	1.5%	1.3%	0.6%	0.4%	1.2%
4 Ln	9724	8681	30	6539	8786	18
2 VCs	2.0%	1.8%	2.0%	0.7%	0.5%	1.9%
4 Ln	11733	11189	50	7597	10954	30
4 VCs	2.4%	2.3%	3.3%	0.8%	0.6%	3.1%
8 Ln	19287	13498	36	12703	13712	20
1 VC	4.0%	2.8%	2.4%	1.4%	0.8%	2.1%
8 Ln	20386	15572	54	13410	15676	30
2 VCs	4.2%	3.2%	3.6%	1.5%	0.9%	3.1%
8 Ln	23073	19604	90	14876	18865	50
4 VCs	4.8%	4.1%	5.9%	1.7%	1.0%	5.2%

	NG-Large **			NG-Ultra **		
	LUT	DFP	RAM	LUT	DFP	RAM
2 Ln	5702	5373	16	5702	5373	16
1 VC	4.2%	4.2%	8.3%	1.1%	1.1%	2.4%
2 Ln	7878	8410	40	7878	8410	40
4 VC	5.7%	6.5%	20.8%	1.5%	1.7%	6.0%
4 Ln	10604	9367	32	10604	9367	32
1 VC	7.7%	7.3%	16.7%	2.0%	1.9%	4.8%
4 Ln	13254	13926	80	13254	13926	80
4 VC	9.7%	10.8%	41.7%	2.5%	2.8%	11.9%
8 Ln	22578	17409	64	22578	17409	64
1 VC	16.5%	13.5%	33.3%	4.2%	3.4%	9.5%
8 Ln	25793	24975	160	25793	24975	160
4 VC	17.3%	19.4%	83.3%	4.8%	4.9%	23.8%

* TMR not included.

** Inferred values.

Regarding the BRAVE values indicated in Table II, for this IP and the Router IPs the resource usage has been inferred from the RTG4 values. Both BRAVE and RTG4 use LUT4 elements. It has been verified with the SL Intf IP that the resource increase ratio from 1 VC to 2 or 4 VCs is almost identical for BRAVE and RTG4. Hence, the usage ratio between RTG4 and BRAVE SL Intf IP has been used to infer the resources for the NG-Large/Ultra. Only results for 1 and 4 VCs have been included in the table for simplicity.




Fig. 2. SpFi ML on a PolarFire connected to a STAR-Ultra PCIe unit.




Fig. 3. SpFi Multi-Lane interface running on a Versal.

VI. SPACEFIBRE SINGLE-LANE ROUTER IP CORE

The Router 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. 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, round-robin arbitration is performed packet by packet, like a SpW Router.

Fig. 4 shows a simplified Router. Note that SpW ports only have associated a single VC. The configuration port uses the RMAP protocol [16] to configure the routing table, VNs,

Router ports, etc. Nevertheless, a dedicated AXI4-Lite interface can also be used to access the same configuration registers.

A. Specific Features

The STAR-Dundee SpFi SL Router IP is a scalable, fully configurable non-blocking router. 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. This Router implements path and logical addressing, VNs, time distribution and message broadcasting. In addition, it also fully supports the QoS and FDIR capabilities native to SpFi. The maximum number of VNs is 64, but each of these VNs is completely flexible: any VC of any port can be configured to any VN. VNs can be statically or dynamically configured. The VNs can be configured statically during FPGA programming using VHDL constants—allows using the Router IP without using any software host —, or they can be dynamically modified by the user using logic connected to the configuration port or the RMAP protocol, which can be accessed over one of the ports of the Router. The high flexibility of the SpFi Router IP Core ensures that different user needs can be accommodated with ease.

There are up to 256 broadcast channels with higher priority for time-critical broadcast messages. The Router 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 Router IP presents a deterministic low latency switching. Round-robin packet arbitration can only occur within each VN. When arbitration is required, it only takes place when two or more VCs request to access to the same output port within the same VN. A timeout controls if the source or the sink stall in the middle of a packet, or when there is a babbling node. Upon timing-out, the router performs automatic packet spilling of the blocking packet.

B. Area Resources

Table III presents the resource usage for two different Router configurations which have been adopted as reference. The table shows that even a “large” Router of 8 SpFi ports with 4 VC each, plus SpW and AXI4-S ports, would fit inside an RTG4. The port count in the table includes two non-SpFi ports: one SpW port and one AXI4-S port with 2 VCs, that is 3 more VCs. So, for example, the 6 Port Router has 4 SpFi ports plus the SpW and AXI4-S ports. The SpFi SL Interface logic of the ports is included in the table figures, as well as the additional RMAP configuration port—one extra VC—and all the configuration logic (see Fig. 4). Hence, the total number of VCs of the Router is the total number of SpFi VCs plus 4.

Dividing the total number of VCs (12 and 36) of the two scenarios by the resource usage of the different FPGAs produce an interesting result. For all scenarios and devices, the average number of registers per VC is 2200 ± 100 . Regarding LUTs, their number per VC is 2600 ± 100 for LUT4 (RTG4 and PF) and 1500 ± 50 for LUT6 (KUS and Versal). This provides an easy method to calculate a rough estimation for a Router with a different number of ports and VCs.




Fig. 4. SpaceFibre Router Block Diagram.

There are options to reduce the resource usage by tailoring the behaviour of the Router VNs. For some applications, provided that the network requirements are fixed, it is possible to limit the number of ports that can be accessed by a certain VN. This way VNs are pre-defined during the design phase and hardcoded into the Router design. This knowledge can be ported to the IP as a constraint, thus helping to reduce the area of the design.

VII. SPACEFIBRE MULTI-LANE ROUTER IP CORE

The STAR-Dundee SpFi ML Router IP is directly derived from the SL Router. It provides the same functionality but with a configurable number of lanes on the SpFi ML ports. The main difference, apart from the ML interfaces, is that the internal data path width is increased accordingly—including the AXI4-S internal ports—so that the internal clock frequency of the Router does not scale up with the number of lanes. Thus, the internal clock frequency of the ML Router is the same of the SL Router. This multiplies the bandwidth of the Router at the expense of more resources but leaving timing largely unaffected. Note that the increase in congestion can have an impact on the final timing.

This IP has been used to build the primary element of the Hi-SIDE project, the STAR-Tiger, a 10 SpFi port ML Router with 4- and 2-lane ports [17]. In [18] there is an in-depth technical analysis of the ML Router architecture, operation and performance measurements, including latency (packet, switching, broadcast) or throughput depending on the packet size.

The scenarios analysed for the ML Router (Table IV) are identical to the SL Router, with the difference that all SpFi ports have either 2 or 4 lanes. The port count in the table also includes the SpW and AXI4-S ports. The values reported already include the SpFi ML Interface IP Cores used by each port and the additional configuration port (see Fig. 4).

TABLE III. SPFI SINGLE-LANE ROUTER RESOURCE USAGE

	RTG4			XQRKU060*		
	LUT	DFF	LSRAM	LUT	DFF	RAMB36
6 Port	31782	27393	47	17984	28090	48
2 VCs	20.9%	18.0%	22.5%	5.4%	4.2%	4.4%
10 Port	98540	76035	127	55917	78051	128
4 VCs	64.9%	50.1%	60.8%	16.9%	11.8%	11.9%

	RTPF500T*			XQRVC1902*		
	LUT	DFF	LSRAM	LUT	DFF	RAMB36
6 Port	29938	26943	93	17098	27652	48
2 VCs	6.2%	5.6%	6.1%	1.9%	1.5%	5.0%
10 Port	93526	75905	253	53800	75867	128
4 VCs	19.4%	15.8%	16.6%	6.0%	4.2%	13.2%

	NG-Large**			NG-Ultra**		
	LUT	DFF	RAM	LUT	DFF	RAM
6 Port	26379	28489	94	26379	28489	94
2 VCs	19.2%	22.1%	49.0%	4.9%	5.6%	14.0%
10 Port	81788	79076	254	81788	79076	254
4 VCs	59.7%	61.3%	132.3%	15.2%	15.6%	37.8%

SpFi Interface IP resources are included.

* TMR not included.

** Inferred values.

TABLE IV. SPFI MULTI-LANE ROUTER RESOURCE USAGE

	RTG4			XQRKU060*		
	LUT	DFF	LSRAM	LUT	DFF	RAMB36
2L6P	48043	44434	59	28579	42829	33.5
2 VCs	31.6%	29.3%	28.2%	8.6%	6.5%	3.1%
2L10P	139644	116463	171	82625	109168	101.5
4 VCs	92.0%	76.7%	81.2%	24.9%	16.5%	9.4%
4L6P	77279	69216	117	46808	65607	61.5
2 VCs	50.9%	45.6%	56.0%	14.1%	9.9%	5.7%
4L10P	-	-	-	128600	158420	185.5
4 VCs	-	-	-	38.8%	23.9%	17.2%

	RTPF500T*			XQRVC1902*		
	LUT	DFF	LSRAM	LUT	DFF	RAMB36
2L6P	47042	44809	67	26492	42844	33.5
2 VCs	9.8%	9.3%	4.4%	2.9%	2.4%	3.5%
2L10P	135690	117563	203	77366	109295	101.5
4 VCs	28.2%	24.4%	13.4%	8.6%	6.1%	10.5%
4L6P	76001	69750	123	43106	65608	61.5
2 VCs	15.8%	14.5%	8.1%	4.8%	3.6%	6.4%
4L10P	212500	174216	371	120865	158510	185.5
4 VCs	44.2%	36.2%	24.4%	13.4%	8.8%	19.2%

	NG-Large**			NG-Ultra**		
	LUT	DFF	RAM	LUT	DFF	RAM
2L6P	46278	49729	118	46278	49729	118
2 VCs	33.8%	38.5%	61.5%	8.6%	9.8%	17.6%
2L10P	130332	128230	342	130332	128230	342
4 VCs	95.1%	99.4%	178.1%	24.3%	25.4%	50.9%

SpFi Multi-Lane Interface IP resources are included.

* TMR not included.

** Inferred values.

The RTG4 does not allow for many SpFi ML ports as it has not enough resources for the design to fit. However, the rest of the FPGAs can implement large Routers using 4-lane SpFi ports with a total of 36 VCs (no TMR) while still having a considerable amount of free space for other applications if required. Comparing the resources for 2-lane and 4-lane ML Routers against the SL Router implementation shows that switching from single to 2-lane SpFi ports requires roughly ~50% more resources. Moving to a 4-lane version instead

increases resources demand by ~150%. Note that in a 2-lane version the bandwidth of the Router is doubled and that in the 4-lane case the bandwidth is multiplied by 4.

Regarding the ratio of resources/VC, for the 2-lane a rough order of magnitude is ~3500 registers/VC and ~3900 LUT4/VC or ~2200 LUT6/VC. For the 4-lane case, the ratios are ~4500-5500 registers/VC (the more VCs the lower the ratio) and ~6000 LUT4/VC or ~3500 LUT6/VC.

Finally, note that the last Router configuration scenario (4L10P) requires 32 lanes. The PF FPGA, despite having ample margin to implement such configuration only offers 24 SerDes lanes, so it is not possible to get all these 4-lane SpFi ports out of the FPGA. The values for this configuration have not been added to the RTG4 table because it exceeds available resources.

VIII. CONCLUSIONS

The STAR-Dundee SpaceFibre Single-Lane Interface, Multi-Lane Interface, Single-Lane Router and Multi-Lane Router IP Cores have been designed to be easy to implement in radiation-tolerant FPGAs. In this article 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.

A simple SpFi Single-Lane Interface (1 VC) can be integrated in radiation-hardened FPGAs by using only a 2% of an RTG4 and less than 0.5% in the other FPGAs analysed. This offers a simple way of having a high-speed and resilient communication channel with an FPGA.

The multi-lane capability increases the data throughput of SpFi and the addition of multiple lanes provide hot and warm redundancy, or graceful degradation of the link bandwidth when no redundant lanes are available. Therefore, if more bandwidth or additional robustness is required out of a SpFi link, the Multi-Lane Interface IP is a convenient choice, keeping resource usage at a mere 4% for the RTG4 and 1% or less for the other devices. Finally, the SpFi Routers (SL and ML) can also be integrated in space-qualified FPGAs even when many ports are required.

All the STAR-Dundee IP Cores have been verified in simulation and subsequently validated in hardware prototypes. Both commercial and the main radiation-hardened FPGAs have been used for these validation activities, ensuring full compatibility, and defining an easy adoption path for this technology. IPs come with specific reference designs for each FPGA, and these can directly be implemented in the FPGA to assist the end-user and allow an easy adoption. A comprehensive end user test bench for ModelSim/Quarta simulators is also provided, which can be used as a reference for test integration.

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, a European Union project involving several European aerospace organisations that have developed satellite data-chain technologies for future Earth observation and telecommunication systems [17]. The different elements of the data chain are interconnected via a SpFi network. SpFi is currently being implemented in FPGA and ASIC designs by different missions and products all over the world.

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] 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>.
- [3] M. Farras Casas, S. Parkes, A. Ferrer Florit and A. Gonzalez Villafranca, "Testing SpaceFibre in Orbit: the OPS-SAT and NORBY Technology Demonstrators", 2022 International SpaceWire Conference.
- [4] A. Ferrer Florit, A. Gonzalez Villafranca, S. Parkes and C. McClements, "SpaceFibre Interface and Routing Switch IP Cores", 2018 International SpaceWire Conference, available from https://www.star-dundee.com/wp-content/star_uploads/conference_papers/spacefibre/2018_SPW_Space_Fibre_Interface_Routing_Switch_IP_Cores.pdf (last visited 11th September 2022).
- [5] Microsemi, "RTG4 FPGAs Technical Brief", rev B, Microsemi, October 2021, available from https://www.microsemi.com/document-portal/doc_download/134430-pb0051-rtg4-fpgas-product-brief (last visited 11th September 2022).
- [6] Xilinx, "Radiation Tolerant Kintex UltraScale XQRKU060 FPGA Data Sheet DS882", v1.3, Xilinx, 11th April 2022, available from <https://docs.xilinx.com/v/u/en-US/ds882-xqr-kintex-ultrascale> (last visited 11th September 2022).
- [7] Microchip, "P00145 Product Overview Radiation-Tolerant PolarFire FPGA", Microchip, 2019, available from https://www.microsemi.com/document-portal/doc_download/1244478-rt-polarfire-fpga-product-overview (last visited 11th September 2022).
- [8] Xilinx, "XQR Versal ACAP Product Table", Xilinx, 2021-2022, available from <https://china.xilinx.com/content/dam/xilinx/support/documents/selecti-on-guides/xqr-versal-product-selection-guide.pdf> (last visited 11th September 2022).
- [9] NanoXplore, "NG-Large Space NX1H140ATSP Datasheet", ver 1.0.3, NanoXplore, May 2022, available from <https://files.nanoxplore.com/f/db44862a31ff4521bb4a/?dl=1> (last visited 11th September 2022).
- [10] NanoXplore, "NG-Ultra NX2H540TSC Datasheet", ver 2.1, NanoXplore, June 2022, available from <https://files.nanoxplore.com/f/5785df61b4f545a08ae8/?dl=1> (last visited 11th September 2022).
- [11] ARM, "AMBA AXI Protocol Version 2.0 Specification", ARM, Ed., 2010.
- [12] Texas Instruments Inc., "1.6-Gpbs to 2.5-Gbps Class V Transceiver", Texas Instruments Inc., 2018, available from <http://www.ti.com/product/tlk2711-sp> (last visited 11th September 2022).
- [13] 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.
- [14] 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.
- [15] R. Ginosar, P. Aviely, T. Israeli and H. Meirov, "RC64: High Performance Rad-Hard Manycore", IEEE Aerospace Conference, Big Sky, Montana, 2016.
- [16] 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>.
- [17] S. Parkes et al., "SpaceFibre Payload Data-Handling Network", 2022 International SpaceWire Conference.
- [18] A. Ferrer Florit, A. Gonzalez Villafranca, M. Farras Casas and S. Parkes, "SpaceFibre Multi-Lane Routing Switch IP", 2022 International SpaceWire Conference.