

SpaceFibre Multi-Lane Routing Switch IP

Albert Ferrer Florit
STAR-Barcelona SL
Sant Cugat del Valles, Barcelona, Spain
albert.ferrer@star-dundee.com

Alberto Gonzalez Villafranca
STAR-Barcelona SL
Sant Cugat del Valles, Barcelona, Spain
alberto.gonzalez@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. SpaceFibre provides up to 6.25 Gbit/s per lane, with multi-lane allowing to reach up to 16 times the speed of a single lane.

In this work we present the SpaceFibre Multi-Lane Routing Switch IP Core developed by STAR-Dundee and its subsidiary STAR-Barcelona. This IP provides a highly flexible router comprising a number of ports and a fully configurable, non-blocking, high performance, routing switch matrix. The internal ports use AXI4-Stream protocol, and the external ports can implement SpaceFibre or SpaceWire interfaces. The SpaceWire ports include additional bridging logic for efficient interconnection between SpaceWire and SpaceFibre equipment. The core logic of the IP is technology independent but has been optimised to be easily implemented in radiation tolerant FPGAs.

The routing switch is fully compliant with all layers of the SpaceFibre standard, supporting up to 64 virtual networks and 256 broadcast channels. Among other features, it implements network time synchronisation, packet time-outs, and automatic translation between SpaceFibre broadcast messages and SpaceWire broadcast codes (SpaceWire Time-Codes or Interrupts). With up to 8 lanes per SpaceFibre interface, raw link rates of 50Gbps per port can be achieved.

The multi-lane routing Switch IP Core is implemented in the STAR-Tiger Routing Switch of the Hi-SIDE project.

Keywords—SpaceFibre, Routing Switch, IP Core, FPGA, Radiation Tolerant, RTG4.

I. INTRODUCTION

SpaceFibre (SpFi) [1] is the next generation of SpaceWire (SpW) [2] network technology for use onboard spacecraft. It supports high data-rate payloads, provides robust, long-distance communications for launcher applications, and supports avionics applications with deterministic delivery capability. SpaceFibre provides in-built Quality of Service (QoS) and Fault Detection, Isolation and Recovery (FDIR) capabilities and runs over electrical or fibre-optic cables.

A. SpaceFibre Lanes

New generation payloads, such as SAR and multi-spectral imaging instruments, require the use of multiple parallel high-speed links to fulfil the increasing bandwidth requirements. To accommodate these needs, SpFi supports multi-lane operation, thus allowing data to be sent over several individual physical lanes to enhance throughput and robustness. The multi-lane operation allows much 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 single link. Furthermore, when a lane fails the multi-lane mechanism supports hot redundancy and graceful degradation by automatically spreading traffic over the remaining working lanes.

B. SpaceFibre Link

A SpFi link is made up of one or more lanes. In a multi-lane link, some of the lanes can be unidirectional provided that at least one lane is bi-directional. The SpFi link provides QoS and error recovery. SpFi links carry traffic (application information) through one or more virtual channels (VCs). There is a maximum of 32 VCs on a link. Traffic entering VC N comes out of VC N at the other end of the link.

Each VC is provided with a QoS which has three components: bandwidth reservation, priority and scheduling. Bandwidth reservation, reserves a portion of the link bandwidth for the VC. Priority assigns a priority-level to the VC so that higher priority VCs are able to send before lower priority ones. Scheduling divides time into 64 sequential time-slots and specifies in which of those time-slots a VC is permitted to send information. These three different QoS components are not alternatives, they work together.

C. SpaceFibre Network

SpFi carries SpW packets over VCs and provides a broadcast feature similar to SpW time-codes but offering much more capability.

The information to be sent is packaged in the same packet format as SpW. SpFi also uses the same routing concepts as SpW including both path and logical addressing. SpFi broadcasts are short messages that are expected to be received by all nodes of the network with minimum latency and jitter. Each broadcast source must use a different broadcast channel.

Fig. 1 shows an example onboard network using SpFi. The control processor is used to configure and control all on-board data-handling equipment. It therefore needs a connection to

every instrument and to the mass memory. This could be provided using a separate command and control network, but this would result in additional mass and power consumption. The addition of a SpFi Routing Switch connecting to all of the on-board equipment allows the control processor to send commands and receive information from all of the on-board data handling units.

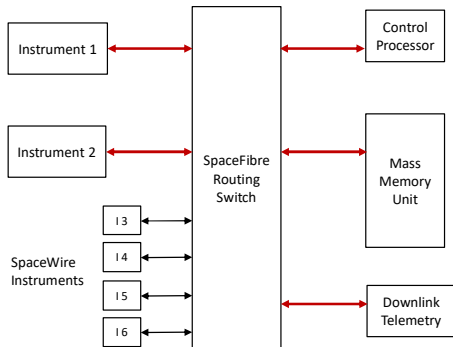


Fig. 1. Onboard network example using SpFi.

Each SpFi link implements multiple VCs, each one with specific QoS parameters. For example, the bandwidth allocation parameter of each VC can be set to the expected bandwidth of each instrument, so they do not interfere with one another or with the control network.

D. SpaceFibre Virtual Networks

The SpFi Network layer defines the concept of Virtual Network (VN). VNs are built from the interconnection between VCs of different ports. These VNs enable the creation of highly flexible SpFi routing switches comprising a number of SpFi interfaces and a fully configurable, non-blocking, high performance, routing switch. This routing switch can theoretically support an arbitrary number of VNs, each effectively behaving like independent SpW networks capable of working at multi-Gbps rates.

Fig. 2 shows a simple example of how the control network (blue path) and two instrument data flows (green and yellow paths) can be assigned to VNs.

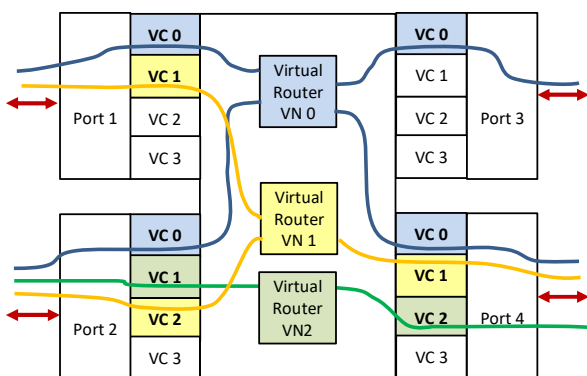


Fig. 2. Example of virtual networks within a SpFi Routing Switch

The traffic running over each VN is constrained by the SpFi QoS mechanism of each of its VCs. Traffic remains within its allocated bandwidth and follows the priority and schedule allocated to it. Data within a VN cannot flow to another VN.

A VN is also able to opportunistically use more bandwidth than it has been allocated, when no other VN has traffic to

send over the links of the SpFi network that the particular VN wants to use.

As far as the addressing of packets and their routing across the network is concerned, SpFi operates in the same way as SpW. This has the substantial advantage that existing application software or SpW equipment can be used with a SpFi network by simply tying a SpW link interface to a SpFi VC interface. The application does not need to know that it is running over SpFi, but gains all the QoS and FDIR advantages of SpFi. This makes the integration of existing SpW equipment both simple and advantageous [3].

E. SpaceFibre Multi-lane Routing Switch IP Core

The STAR-Dundee SpFi Multi-Lane Routing Switch IP is directly derived from the STAR-Dundee SpFi single-lane Routing Switch IP [4]. The multi-lane version provides the same functionality but with a configurable number of lanes on the SpFi ports. The main difference, apart from the multi-lane capable interfaces, is that the internal data path width is increased accordingly—including the AXI4-Stream internal ports—so that the internal clock frequency of the Routing Switch does not scale up with the number of lanes. Thus, the internal clock frequency of the multi-lane version is the same of the single-lane. This multiplies the bandwidth of the Routing Switch at the expense of more resources, but leaving timing largely unaffected.

This paper presents the STAR-Dundee SpFi Multi-Lane Routing Switch IP Core (the Routing Switch). Its architecture and features are described in Section II and Section III respectively. Section IV and V presents the synthesis and performance results. An example of a hardware implementation is described in section VI. Finally, conclusions are presented in section VII.

II. ROUTING SWITCH ARCHITECTURE

The Routing Switch architecture is built around a non-blocking routing switch matrix with a number of ports, which can implement a SpW, SpFi or AXI4-Stream interfaces. Each port has a number of VCs, each one comprising an input and an output VC buffer. Each VC has an associated VN number. The switch matrix interconnects one or more VCs with the same VN number, but each of these VCs must be located 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 allocated throughput and latency within the routing switch matrix. On the other hand, when multiple packets in the same VN need to be transferred from different ports to the same output port, packet-by-packet, round-robin arbitration is performed, similarly to a SpW router.

Fig. 3 shows a simplified Routing Switch architecture with a configurable number of ports. The SpW ports only have associated a single VC but have additional buffering resources. The configuration port implements the RMAP protocol [5] to configure the Routing Table, the VNs, and the SpFi and SpW interfaces.

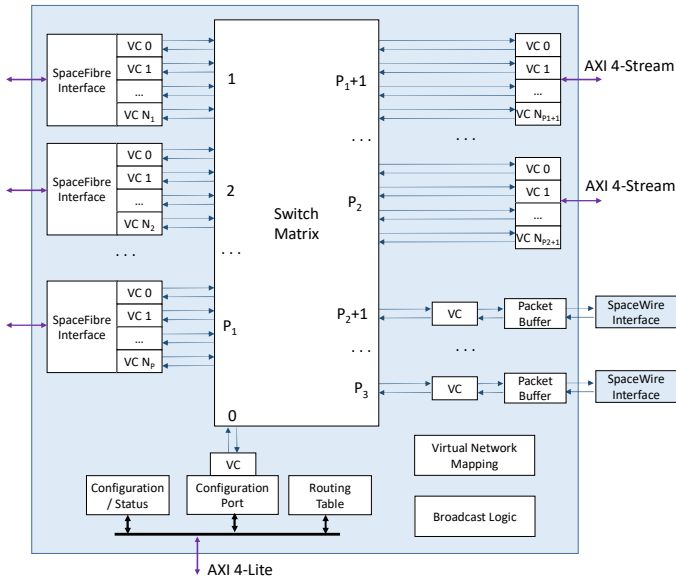


Fig. 3. SpaceFibre Routing Switch block diagram

The Switch Matrix allows to interconnect every VC from one port to any other VC from another port. The specific VN configuration will specify which connections are allowed in a particular application. The non-blocking switching logic allows to reach full bandwidth for each possible connection, independently of other simultaneous connections being used.

When a packet arrives at an input port from a particular VC, the output port is determined based on the packet address and the Routing Table, which translates logical address numbers to output port numbers. Then, the VC within this output port is determined by the VN Mapping block using the VN configuration information.

If an output VC has no space or it is busy transferring data from another packet coming from another input port within the same VN, the packet is stalled until the resource is freed. A round robin scheme ensures fair arbitration for packets from the same VN requesting the same output port at the same time.

Finally, the Broadcast Logic block handles how the broadcast messages received are sent across all allowed output ports, implementing the mechanism defined in the standard to avoid network level broadcast storms in the presence of switching loops.

III. ROUTING SWITCH FEATURES

The Routing Switch has the following main features:

- Technology independent (FPGA or ASIC) but optimised for radiation-hardened FPGAs.
- Configurable number of SpFi, SpW and internal AXI4-Stream ports.
- Configurable SpFi lane rate, number of lanes, and number of VCs per port.
- Configurable target technology (RTG4, PolarFire, Xilinx Kintex/Ultrascale/Versal, generic) for memory blocks and SerDes interface.
- Up to 64 VNs that can be statically or dynamically configured.

- Configuration registers can be accessed via a configuration port using RMAP or using a dedicated AXI4-Lite interface.
- High performance, full non-blocking switch matrix with deterministic switching latency. VNs do not share any switching resources.
- Round-robin arbitration with watchdog timeout for packets in the same VN requesting the same output port.
- SpW/SpFi network capabilities such as path and logical addressing with a routing table.
- Up to 256 broadcast channels with higher priority for time-critical broadcast messages.
- Simple and efficient integration with SpW networks using SpW packet buffers and automatic SpW to SpFi broadcast translation.
- Internal timer tracks time being distributed over the network.

A. Configuration

The main capabilities of the Routing Switch can be configured statically before the IP is synthesised. The most important are the target technology, number and type of ports, lane rate, lanes and VCs per port, the default value of the routing table, and the VN setup. This allows to use the Routing Switch with the default configuration, without using any software host, and to optimise router resources required for a specific application.

Each VN is configured by specifying the VC used for each port of the router in which the VN is used. Table I shows the VNs configuration table for the simple case shown in Fig. 2.

TABLE I. VIRTUAL NETWORK CONFIGURATION EXAMPLE

VN number	Virtual channel number			
	Port 1	Port 2	Port 3	Port 4
0	0	0	0	0
1	1	2	-	1
2	-	1	-	2

Once the Routing Switch has been implemented it can be configured after reset by accessing to the router registers using the AXI4-Lite interface or the configuration port zero and the RMAP protocol. The Routing Switch memory space uses 16-bits address with either bit 15 or bit 16 set, to support network discovery of legacy devices such as the SpW 10-X Router ASIC (AT7910E) [6]. Fig. 4 shows the memory map regions of the Routing Switch.

Region																	Start
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
Legacy support	0	0	Reserved													0x0	
Device Information	0	1	0	0	0	0	0	0	Register Select								0x4000
Routing Table	0	1	0	0	0	0	0	1	Logical Address				Sel				0x4200
Network Management	0	1	0	0	0	1	0	Register Select								0x4400	
Broadcast Notification	0	1	0	0	0	1	1	Register Select								0x4600	
Device Specific	0	1	1	Register Select													0x6000
VC Information	1	Port Number				0	VC Number			Register Select					0x8000		
Port Information	1	Port Number				1	0	0	Register Select					0x8200			
Link Information	1	Port Number				1	0	1	Register Select					0x8280			
Lane Information	1	Port Number				1	1	0	Lane Number		Regs Select			0x8300			
Reserved	1	Port Number				1	1	1	Reserved					0x8380			

Fig. 4. Memory map regions of the Routing Switch

B. Packet Addressing

The routing control logic interprets the first byte of each received packet as the packet address. If the leading byte of a packet has a value between 0 and 31, path addressing is used and the packet will be forwarded to the corresponding port number. Otherwise, logical addressing is used, and the output port is determined by the routing table. Therefore, this works the same as a SpW router.

However, SpFi links also support leading FILL characters in a SpFi packet. When leading FILLs are present, the first packet byte is then replaced by a FILL character in case it is a path address value. If leading FILLs are not present, the first packet byte is removed when it is a path address value. This behaviour ensures packet cargo is 32-bit aligned when it arrives at the destination when the packet is transferred across SpW and SpFi networks using path address bytes, as SpW does not support sending FILL characters.

Fig. 5. shows an example for these two scenarios that the router supports. Note that the number of path address bytes matches the number of network hops, the FILL character is represented by symbol \emptyset and that the SpFi standard specifies that the router should remove 4 consecutive leading FILL characters.

Source	1st Hop	Destination
01 02 FE 10 11 12 13 14	02 FE 10 11 12 13 14 15	FE 10 11 12 13 14 15 16
\emptyset \emptyset 01 02 FE 10 11 12	\emptyset \emptyset \emptyset 02 FE 10 11 12	FE 10 11 12 13 14 15 16

Fig. 5. Examples of path address processing.

C. SpaceWire Packet Buffers

Each SpW port has a packet buffer on its receive side and a FIFO buffer on its transmit side. The Packet Buffer buffers packets arriving over the SpW port. It only forwards full packets to the SpFi VN. The packets arrive at the buffer at SpW speeds and leave at SpFi speeds, so that the VC sending the SpW packet is not held up by the slower SpW interface. If the incoming packet is larger than the size of the Packet Buffer there is a configurable option to spill the remaining of the packet. The FIFOs on the transmit side of the SpW interface simply buffer the traffic arriving over SpFi. The size of the buffers and FIFOs is configurable.

D. Watchdog Timeout Mechanism

SpFi VNs decouple the traffic flowing in one VN from the traffic in another VN. Therefore, if a packet becomes blocked

in a VN it will not affect packets in another VN. Any congestion in a VN will not affect another VN.

However, within the same VN, package blocking can still occur, in the same way that there can be packet blocking and congestion in a SpW network. There are three main causes:

- Source stalls and stops transmitting bytes of a SpW packet while the packet is being routed.
- Destination stalls and stops receiving bytes of a SpW packet while the packet is being routed.
- A package is blocked due to another packet being blocked. This can only occur if both use the same VC at one of the links of the path to their destination.

The Routing Switch implements a watchdog timer to prevent indefinitely packet blocking. When the packets transfer stops the watchdog timer is started. When the maximum time elapses, the packet is spilled.

E. Broadcast messages

The Routing Switch forwards any broadcast message type but there are some broadcast types that are processed in a specific manner:

- Time:** Used to synchronise the local time with the network time. The CCSDS value of the time broadcast message is validated by comparing the value of two consecutive time broadcasts received. The time difference is compared with the difference in arrival time using the local clock. If the difference is very small the broadcast time value is accepted. The local time is then updated with this new validated value.
- Time-Slot:** Used to set the device time-slot and synchronise time-slots across the network for SpFi network scheduling.
- SpW Time-Code:** Contains a SpW broadcast code of type Time-Code. The Routing Switch generates this broadcast when a Time-Code is received in a SpW port. The broadcast generated is sent to all ports except the SpW port that received the Time-Code. Likewise, the Routing Switch distributes a SpW Time-Code to all the SpW ports when this broadcast is received from a SpFi port.
- SpW Interrupt:** Contains a SpW broadcast code of type Distributed Interrupt. The Routing Switch generates this broadcast when a Distributed Interrupt is received in a SpW port. The broadcast generated is sent to all ports except from the SpW port that received the Distributed Interrupt. Likewise, the Routing Switch sends a SpW Distributed Interrupt to all the SpW ports when this broadcast is received from a SpFi port.

The broadcast type value of each of these broadcast types can be configured, with lower broadcast type values being forwarded with higher priority. Fig. 6 shows the format of the last three broadcast message types. The second 32-bit data word is the bit-inverse value of the first 32-bit data word. The least significant bits hold the actual value. SpFi Broadcasts messages generated by the router from SpW broadcast codes use the Broadcast channel specified in register "Router BC Channel". The value in this register must be initialised to

enable the forwarding of SpW Time-Codes and Interrupts to SpFi broadcast messages.

Fig. 7 shows the Time broadcast message format which holds a CCSDS time information value.

0		7		8		15		16		23		24		31	
COMMA		SBF		Broadcast Channel		Broadcast Type									
Time-Slot or Time-Code		0x0		0x0		0x0									
Interrupt code		Bit-Inverse		Bit-Inverse		Bit-Inverse		Bit-Inverse							
EBF		RSVD/LATE		SEQ_NUM		CRC									

Fig. 6. Broadcast frame carrying a SpW broadcast code or SpFi time-slot.

0		7		8		15		16		23		24		31	
COMMA		SBF		Broadcast Channel		Broadcast Type = 0									
DATA 1 LS		DATA 1		DATA 1		DATA 1 MS		DATA 1 MS		DATA 1 MS		DATA 1 MS		DATA 1 MS	
Fractional Byte LS		Fractional Byte		Fractional Byte MS		Seconds Byte 0 LS									
DATA 2 LS		DATA 2		DATA 2		DATA 2 MS		DATA 2 MS		DATA 2 MS		DATA 2 MS		DATA 2 MS	
Seconds Byte 1		Seconds Byte 2		Seconds Byte 3		Seconds Byte 4 MS									
EBF		RSVD/LATE		SEQ_NUM		CRC									

Fig. 7. Broadcast frame carrying a CCSDS time information.

IV. SYNTHESIS RESULTS

The Routing Switch has been designed to achieve timing closure at the highest data rates supported by the transceivers available in existing radiation-tolerant technologies. The IP supports lane rates of 3.125 Gbps in RTG4 and 6.25 Gbps in PolarFire FPGAs. In UltraScale and Versal Xilinx devices, faster speeds are possible.

The Routing Switch IP has also been designed to scale well in both timing and area metrics when the number of lanes, ports and VCs are increased. This means that the same maximum lane rates can be achieved independently of these parameters. Regarding area, Fig. 8 and Fig. 9 show that area resources increase linearly with these parameters. Note that this same slope is common to the number of lanes and VCs per port. A different slope—slightly above 1—is associated with the number of ports.

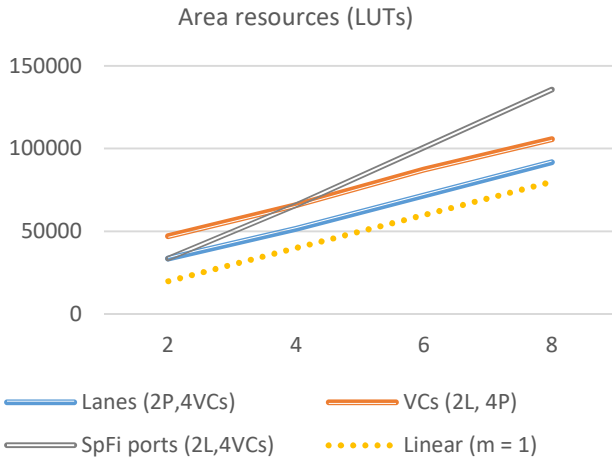


Fig. 8. Linear dependency of LUTs according to different parameters in a PolarFire implementation.




Fig. 9. Linear dependency of Flip-Flops according to different parameters in a PolarFire implementation.

Table II presents the resource usage after synthesis for different value of number of lanes, ports and VCs. The port count in the table includes two non-SpFi ports: one SpW port and one AXI4-Stream internal port with 2 VCs. The SpFi Multi-Lane Interface logic of the ports is included in the logic count, as well as the additional RMAP configuration port (see Fig. 3).

TABLE II. SpFi MULTI-LANE ROUTING SWITCH RESOURCE USAGE

	RTG4			XQRKU060 *		
	LUT	DFP	LSRAM	LUT	DFP	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	DFP	LSRAM	LUT	DFP	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	DFP	RAM	LUT	DFP	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.

V. PERFORMANCE RESULTS

The Routing Switch architecture supports the full bandwidth of the SpFi interfaces without any dependency on the number of simultaneous data flows within the router. There is only potential blocking when two input ports request the same output port within the same VN. Therefore, the two main performance figures, the packet latency and the data packet throughput, can be easily obtained.

Table III shows the packet latency measured in clock cycles since the start of a packet entering an input port until the start of the packet appearing at the output port. The switching latency only takes into account the time a packet stalls until the output port is determined. This time determines the packet data throughput. The use of logical addressing increases both latency values due to the need to access the routing table.

TABLE III. ROUTING SWITCH LATENCY

	Path addressing	Logical Addressing
Packet Latency	34 clock cycles	38 clock cycles
Switching Latency	22 clock cycles	25 clock cycles
Broadcast Latency	10 Clock cycles	

The Routing Switch has been optimised for timing and to scale well when the number of lanes, ports, and VCs is increased. The trade-off is more pipelining, which increases latency. For example, at 6.25 Gbit/s lane rate and with a core clock of 156.25 MHz, the packet latency is around 243 ns. However, for fast FPGAs or ASICs, the core clock of the Routing Switch can be much faster than the one used by the SpFi interfaces so the latency can be reduced in exchange for more power utilisation.

The broadcast latency of low priority broadcast types can be higher than the one shown if there are other higher priority broadcast pending to be sent. This increases the jitter of this lower priority broadcast but the broadcast message is then modified by the router with the delayed status flag set, so the destination knows that the broadcast has been delayed by the router.

Fig. 10 shows the data throughput depending on the packet size.




Fig. 10. Throughput depending on the size of the packets.

VI. HARDWARE IMPLEMENTATION

The SpFi Multi-Lane Routing Switch IP has been used to build the primary element of the Hi-SIDE project, the STAR-Tiger, a 10-port SpFi Multi-Lane Routing Switch with 4-lane and 2-lane ports [7].

The Hi-SIDE project is a European Union project carried out by several leading aerospace organisations from across Europe. It aims to develop satellite data-chain technologies for future Earth Observation and Telecommunication systems. The data chain elements are interconnected via a SpFi network.

Fig. 11 shows a photograph of the STAR-Tiger SpFi routing switch. STAR-Tiger is used in the Hi-SIDE project for transferring data at high data-rates between instruments, mass-memory, data compressor/processor and downlink transmitters. It is also used to provide the control network used by the control computer to control both the network and the equipment attached to the network. It has the following key features:

- 10 SpaceFibre ports
 - Two quad-lane ports
 - Eight dual-lane ports
 - Lane speed up to 6.25 Gbit/s
 - Port data rate 19.2 Gbit/s (quad-lane) and 9.6 Gbit/s (dual-lane port)
- 2 SpW ports
- 2 further SpW ports for programming STAR-Tiger
- Power consumption 13.5W typical at 20 °C
- 108 x 108 x 70 mm (excluding mounting brackets)
- Spaceflight TRL5/6 level design
 - Electronic components are EM flight parts or industrial/commercial equivalents of flight parts
 - Conduction cooled
 - Operating temperature range: -25 to +55 °C



Fig. 11. Photograph of a STAR-Tiger unit.

VII. CONCLUSIONS

The STAR-Dundee SpFi Multi-Lane Routing Switch IP has been designed for ease of use and to achieve the highest lane rates on space-grade FPGAs, independently on the number of ports, virtual channels and number of lanes. Synthesis results show that area resources also scale well when the values of these configuration parameters are increased.

The Routing Switch supports lane rates of 3.125 Gbps in RTG4 and 6.25 Gbps in PolarFire FPGAs. In UltraScale and Versal Xilinx devices, faster speeds are possible. The multi-lane capability multiplies the data throughput of SpFi ports and provides hot and warm redundancy, or graceful degradation of the link bandwidth when no redundant lanes are available. It allows to implement a SpFi routing switch with more than 250 Gbps of aggregated bandwidth.

The Routing Switch also provides additional functionality to easily integrate SpaceWire devices into SpaceFibre networks and to distribute accurate time information across the network using broadcast messages.

The Routing Switch has been tested and subsequently validated within a full satellite data-chain technology demonstrator. Both commercial and a radiation-tolerant FPGAs have been used for these validation activities, ensuring full compatibility and defining an easy adoption path for this technology.

STAR-Dundee has developed and demonstrated the critical SpaceFibre Router technology necessary for

cutting-edge on-board data-handling systems with very high data rate sensor and telecommunications systems.

ACKNOWLEDGMENT

The Hi-SIDE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 776151.

REFERENCES

- [1] 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>.
- [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] S. Parkes, A. F. Florit, A. G. Villafranca, C. McClements and D. McLaren, "SpaceFibre network and routing switch," 2017 IEEE Aerospace Conference, 2017, pp. 1-7, doi: 10.1109/AERO.2017.7943805.
- [4] A. Ferrer Florit, A. Gonzalez Villafranca, M. Farras Casas and S. Parkes, "SpaceFibre Routing Switch IP Implementation in Radiation-Tolerant FPGAs", DASIA 2019
- [5] ECSS Standard ECSS-E-ST-50-52C, "SpaceWire – Remote memory access protocol", European Cooperation for Space Data Standardization, February 2010, available from <http://www.ecss.nl>.
- [6] Atmel, "AT7910E SpW-10X SpaceWire Router Datasheet", Atmel, 2012, available from <http://www.microchip.com>.
- [7] S. Parkes et al., "SpaceFibre Payload Data-Handling Network", 2022 International SpaceWire Conference.