SpaceFibre Multi-lane

SpaceFibre, Long Paper

Albert Ferrer Florit, Alberto Gonzalez Villafranca
STAR-Dundee Ltd
STAR House, 166 Nethergate
Dundee, DD1 4EE, UK
albert.ferrer@star-dundee.com

Steve Parkes
Space Technology Centre, University of Dundee,
Dundee, DD1 4EE, UK
sparkes@computing.dundee.ac.uk

***Abstract*— SpaceFibre is a multi-Gbits/s, on-board network technology for spaceflight applications, which runs over electrical or fiber-optic cables. SpaceFibre supports multi-lane, thus allowing data to be sent over several individual physical lanes to enhance throughput and robustness. This is required by new generation payloads, such as SAR and multi-spectral imaging instruments. This paper describes the development of the multi-lane capabilities of SpaceFibre and its successful hardware implementation on space-qualified devices. The protocol has been designed to work with an arbitrary number of bidirectional or unidirectional lanes. In the event of a lane failing, SpaceFibre multi-lane mechanism supports hot redundancy and graceful degradation by automatically spreading traffic over the remaining working lanes. User data transfer is resumed in just a few microseconds without any data loss. These advanced capabilities are not provided in other high-speed link protocols available for space applications.**

***Index Terms* — SpaceFibre, multi-lane, hot redundancy, SpaceWire, Networking, Spacecraft Electronics.**

I. Introduction

SpaceFibre is a new technology for use onboard spacecraft that provides point-to-point and networked interconnections at Gigabit rates with Quality of Service. SpaceFibre interoperates seamlessly with a SpaceWire network over virtual channels, as it uses the same packet definition. It provides broadcast capabilities and it is able to operate over a copper or fiber-optic communication medium.

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 [1]. To accommodate these needs, SpaceFibre supports multi-lane operation, thus allowing data to be sent over several individual physical lanes to enhance throughput and robustness.

This paper describes the development of the multi-lane capabilities specified in the SpaceFibre Standard [2] and its hardware implementation on radiation hardened space-qualified FPGAs.

Multi-lane is an optional capability of a SpaceFibre link defined in the Multi-Lane layer of the SpaceFibre protocol stack. As shown in Fig. 1, the Multi-Lane layer is defined between the Data Link layer and the Lane layer implemented for each available lane.

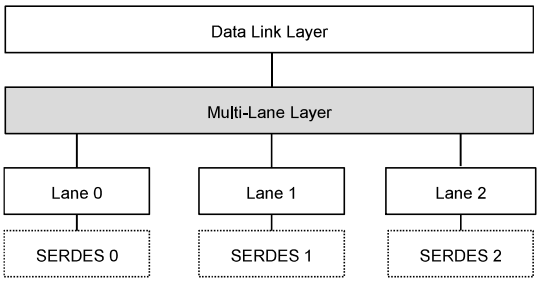

Figure 1. SpaceFibre Multi-Lane layer

The Data Link layer provides quality of service and flow control for a SpaceFibre link. It frames the information to be sent over the link to support QoS and multiple virtual channels. It also provides error recovery capabilities, detecting any frames or control words that go missing or arrive containing errors and resending them.

The Lane Layer establishes a connection across a SpaceFibre lane using a lane initialization state machine. This ensures that bit, symbol and word synchronizations are achieved and that the two ends of the lane are both ready to send and receive data with a nominal Bit Error Rate. The Lane Layer also encodes data and control words into 8B/10B symbols, sends and receives symbols over the lane and decodes the received symbols into data and control words.

The Multi-lane layer coordinates the operation of multiple lanes as a single SpaceFibre link, providing higher data throughput and redundancy. Because the logic that initialize a lane and monitor its status is located below the multilane layer, each lane can be initialized and operated independently of each other.

This architecture also supports the implementation of graceful degradation, which means that in the event of one or more lanes failing, traffic is spread over the remaining working lanes automatically. When combined with Data Link layer QoS, the bandwidth allocated to lower priority virtual channels is reduced when required to ensure that most important information gets through and deterministic traffic is maintained. Bandwidth overprovision and dynamic power management is 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.

The multi-lane requirements are expanded and consolidated in section 2. Section 3 describes the protocol analysis and the design of the SpaceFibre multi-lane capabilities. Section 4 shows the hardware implementation. Finally, conclusions are made in section 5.

II. Multi-lane requirements

The multi-lane capabilities of SpaceFibre have been designed to meet the following end-user requirements:

- Support an arbitrary number of lanes. This allows redundancy and graceful degradation without any restriction on the number of lanes.
- Re-synchronize both ends when the number of lanes changes, without resetting any lane, and as fast as possible. This way user data can be easily buffered when a lane is added or removed, until the link is again ready.
- Support hot redundancy.
- Support dynamic unidirectional lanes to save power and mass for asymmetric user data flows.
- Robust against lane errors and misconfiguration.
- Keep the same protocol overhead than single lane configuration.
- Number of lanes must be independent on the port width of the end-user interface.

These requirements enable unique capabilities for SpaceFibre. Other high speed protocols have limitations in the redundancy mechanisms. For example, in RapidIO [3] when one lane fails, the link falls back to a single lane. PCI Express [4] allows the link to continue using more than one lane, but it takes time as the link needs first to be reset. Interlaken [5] allows an arbitrary number of lanes to operate but does not define a mechanism for link reconfiguration when a lane fails, as this is expected to be done by software.

Fig. 2 shows a use case enabled by the above requirements. In this setup, unidirectional lane 6 can be enabled when one lane fails or higher data rate is required and bidirectional lane 2 can be set as a unidirectional lane for power saving reasons. Note that at least one bidirectional lane must be working for the link to operate.

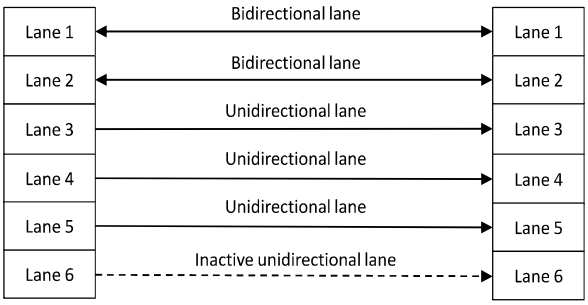

Figure 2. SpaceFibre Multi-Lane layer use case.

There are additional requirements related with SpaceFibre standardization efforts:

- Single-lane SpaceFibre implementation must not be affected by new rules added by multi-lane capabilities. This ensures that legacy single-lane implementations are still compatible with future single-lane implementations.
- Minimize modifications to the definition of other layers in the SpaceFibre standard.

Finally, there are the requirements regarding SpaceFibre implementation on space qualified devices:

- Must be feasible to implement in radiation hardened FPGAs, which are slower than state of the art COTS components.
- Resource usage in radiation hardened FPGAs must be minimized.

III. Multi-lane Design

The multi-lane specifications of SpaceFibre were designed with the following methodology:

1. Identify and evaluate the key concepts that could allow the requirements to be met (e.g. define generic protocol sequence diagrams).
2. Constrain these concepts to work with the more specific set of rules already specified in the SpaceFibre standard (e.g. adapt to control word definitions)
3. Validate the concepts and derived new set of rules in a prototype using a software simulator that can easily be modified. If an issue is found, rework the concept and/or associated set of rules.
4. Refine the simulation engine until it validates with high accuracy the proposed multilane specifications.

The resulting main concepts and specifications are explained in the following subsections.

A. Distribution of control and data words over sending lanes

In this specification, a row is defined as the set of words sent over all sending lanes simultaneously. These words can be data or control words. Data words contain data from the user interface and control words are generated by the Data Link layer to support the protocol operation.

In a single-lane SpaceFibre implementation, control words are processed at a rate of 62.5Mhz for a 2.5Gbps link rate, as there can only be one single control word for every 40 bits. If a row can contain multiple control words, then the processing rate required would increase with the number of lanes [6]. It would then not be possible to implement multi-lane in some radiation hardened devices. Therefore we must enforce that a row can only contain one control word.

The simplest solution is to replicate each control word to be sent across all lanes in a row. This avoids to have complex rules that deal with mixing data and control words in the same row. Fig. 3 shows this solution and the use of the PAD control word when the size of a data frame is not a multiple of the number of lanes available.

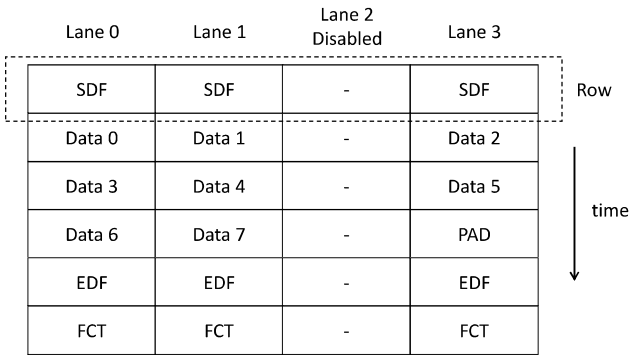

Figure 3. Words forming a row across a multi-lane link

In order to keep the same protocol overhead than single lane implementations, the maximum data frame size needs to be increased. The maximum data frame size was defined as a trade off between latency and protocol overhead. In a multi-lane solution, the maximum data frame size can be increased without modifying these metrics.

Another advantage of this solution is that the data frame CRC can be computed lane by lane as the CRC is provided in the EDF control word. A different CRC value can be included in the EDF of each lane. This keeps intact the error detection capabilities of the CRC for burst errors and for the amount of data covered by the CRC. More important, this solution simplifies the computation of the CRC in slow FPGAs as the incoming rate of the data covered by the CRC is kept the same than in single-lane implementations.

B. Lane alignment at reception

At the receiver side, the set of words received from each lane need to be aligned to compensate for small differences in lane delays. This delays are due to different cable lengths or line driver delays. Therefore, data and control words can not be passed to the Data Link layer until the multi-lane layer has compensated this skew and it is processing the same original rows sent by the sender side.

The lane alignment is usually done with a set of FIFOs that compensate the delays of each lane, using a specific control word, called ALIGN, that is known to be sent over all lanes simultaneously. To cover the case of lane errors, the ALIGN word needs to be sent periodically and with a minimum separation in between. When the alignment is no longer needed, the sending of ALIGNs can be disabled to avoid increasing the protocol overhead.

Fig. 4 shows the set of words received with a skew between lanes, which is compensated in Fig. 5 using ALIGN control words.

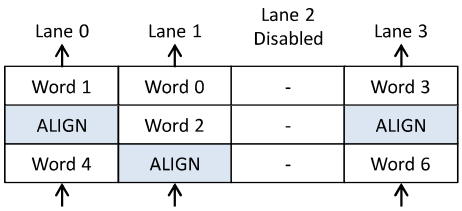

Figure 4. Rows not aligned at reception

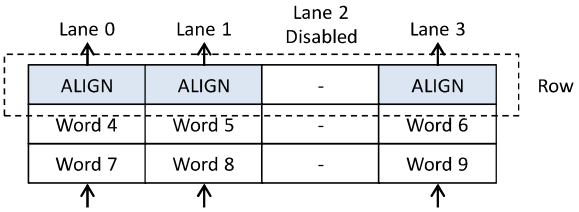

Figure 5. Aligned rows at reception

To avoid data corruption, data words in a row should be processed in the same order than they were placed in the row by the sending side. To help with this requirement, each lane has a lane number associated and it is enforced that words are processed starting with the lowest lane number. Then, this requirement can be fulfilled if the receiver side knows two parameters:

a) The lane number of each lane

b) The total number of lanes used by the sending side.

This information is provided within the ALIGN word itself, so the information is up to date when the receiver side use these words for the alignment procedure.

C. Alignment state machine

It has been stated that data and control words can not be passed to the Data Link layer until the multi-lane layer has completed the alignment process and both sides are aware of the lanes used for sending and receiving. This means that this process has to be performed each time the number of lanes in active state, i.e. working lanes, changes due to a lane error or a lane being enabled or disabled by the user.

The alignment state machine ensures that no data is being transferred to the Data Link layer when the lanes are not aligned and that ALIGN words are only sent when it is required. Three states are defined:

a) Not Ready: lanes have not been aligned and only ACTIVE and ALIGN words are sent. This state is set when the number of lanes in active state changes or an alignment error is detected.

b) Near-End Ready: lanes are aligned and the Data Link layer is being used to send and receive data and control words. However, ALIGN words are sent as the far-end has still not sent any data word.

c) Both Ends Ready: data words were received indicating that the far-end has aligned the lanes. No ALIGN words are sent. Link is ready.

The ACTIVE control word sent by the multilane layer has three functions:

a) Stop the flow of words from the Data-Link layer when the lanes are not aligned.

b) Indicate to the far-end that the lanes are not aligned.

c) Indicate in the cargo of this word which lanes are in active state, so the other end can synchronize the active state of the lanes. This is especially important in unidirectional transmit lanes, which can not detect if the receiving side has disconnected the lane. In a bidirectional lane, the lane initialisation state machine can detect if the other side exits the active state.

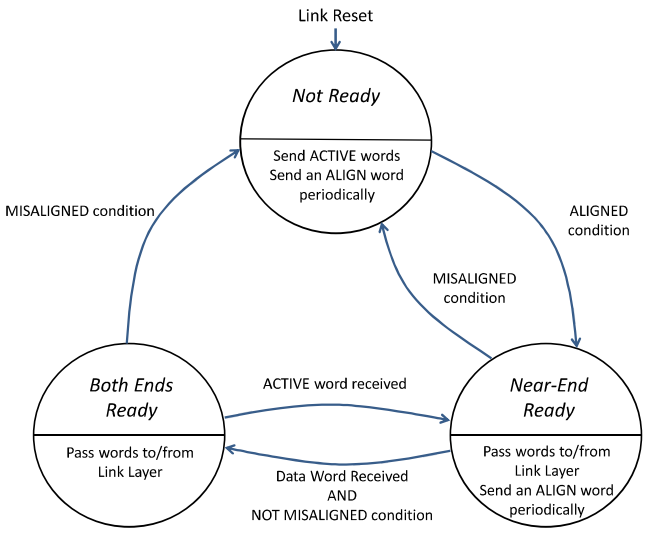

Figure 6. Alignment state machine

Fig. 6 shows the alignment state machine. The ALIGNED condition is asserted when all receiving active lanes are aligned, these lanes are the same active lanes indicated by the received ACTIVE words and they are consistent with the content of the ALIGN words. The MISALIGNED condition is asserted when the active lanes change in the near end or far end or there is an error related with the alignment process.

Fig 7 shows a protocol sequence diagram describing what occurs when one lane fails at the far end and becomes not active. The near end detects this event when it receives the ACTIVE words. It then moves to Not Ready state and starts sending ACTIVE words too. When the alignment process is completed using ALIGN words, the state machines move to Near-End Ready state. Finally when data words are received indicating the other side is in Near-End Ready, the state machines move to Both Ends Ready state and the link is considered to be ready.

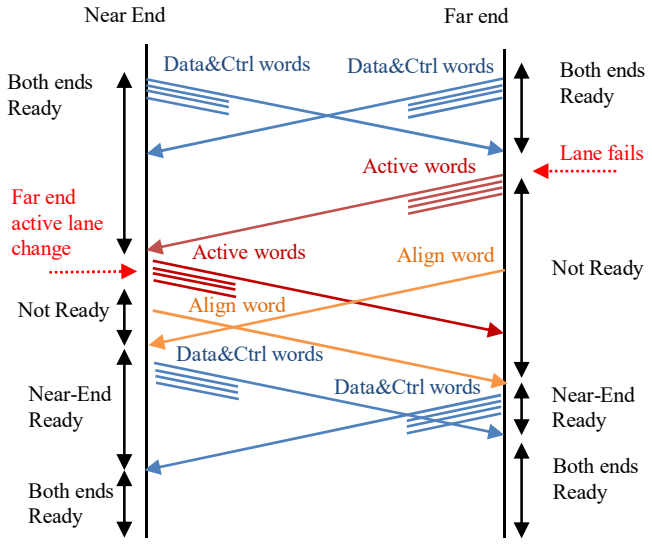

Figure 7. Link ready protocol sequence diagram

	A:TX0	A:TX1	B:RX0	B:RX1	B:Align	B:FIFO0	B:FIFO1	B:RCV0	B:RCV1
269	Active(11)	Active(11)	INIT3	Active(01)		IDLE	Active(01)		RxERROR
270	Active(11)	Active(11)	INIT3	Active(11)		IDLE	Active(11)	RxERROR	
271	Active(11)	Active(11)	IDLE	Align(2,1)		IDLE	Align(2,1)	RxERROR	
272	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Align(2,1)	RxERROR	
273	Active(11)	Active(11)	Align(2,0)	Active(11)		Align(2,0)	Align(2,1)	RxERROR	
274	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
275	Align(2,0)	Align(2,1)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
276	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
277	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
278	Active(11)	Active(11)	Active(11)	Align(2,1)		Active(11)	Active(11)	RxERROR	
279	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
280	Active(11)	Active(11)	Align(2,0)	Active(11)		Align(2,0)	Align(2,1)	RxERROR	
281	Active(11)	Active(11)	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
282	Align(2,0)	Align(2,1)	Active(11)	Active(11)		Active(11)	Active(11)		
283	SDF	SDF	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
284	Data 0	Data 1	Active(11)	Active(11)		Active(11)	Active(11)	RxERROR	
285	EDF 0	EDF 0	Active(11)	Align(2,1)		Active(11)	Active(11)	RxERROR	
286	SDF	SDF	Active(11)	SDF		Active(11)	Active(11)	RxERROR	
287	Data 2	Data 3	Align(2,0)	Data 1		Align(2,0)	Align(2,1)	RxERROR	
288	Data 4	Data 5	SDF	EDF 0		SDF	SDF	RxERROR	
289	Data 6	Data 7	Data 0	SDF		Data 0	Data 1		
290	Data 8	Data 9	EDF 0	Data 3		EDF 0	EDF 0	SDF	
291	Data 10	Data 11	SDF	Data 5		SDF	SDF	Data 0	Data 1
292	Data 12	Data 13	Data 2	Data 7		Data 2	Data 3	EDF 0	
293	Data 14	Data 15	Data 4	Data 9		Data 4	Data 5	SDF	
294	Data 16	Data 17	Data 6	Data 11		Data 6	Data 7	Data 2	Data 3
295	Data 18	Data 19	Data 8	Data 13		Data 8	Data 9	Data 4	Data 5
296	EDF 1	EDF 1	Data 10	Data 15		Data 10	Data 11	Data 6	Data 7
297	SDF	SDF	Data 12	Data 17		Data 12	Data 13	Data 8	Data 9

Figure 8. Simulator tool screenshot of the alignment process

The time it takes for a multi-lane link to resume sending data after a lane has failed or a new lane has been added is less than a few microseconds. More precisely, it is the round trip delay of the ACTIVE words plus the delay between the sending of ALIGN words, which is the time needed to send 8 words.

Fig 8 is a screenshot of the simulator tool used for the validation of the alignment process and the associated state machine. The first set of two columns show the words sent by the sender side using two lanes. The ALIGN word indicates that two lanes are used for sending and the lane numbers are 0 and 1. The ACTIVE word indicates that lane zero and one are active (bits zero and one are set). The second set of two columns show how the row sent is received disaligned at the receiver side. The middle Align column shows in yellow when the state machine is in Not Ready state. The third set of two columns shows how row alignment is achieved using a FIFO and the ALIGN word. Finally the last set of two columns show the words received by the Data Link layer.

D. Unidirectional lanes

Single-lane SpaceFibre implementations must be bidirectional even if the end-user data flow is unidirectional, because feedback from the receiver side is required for the protocol to operate. However in a multi-lane implementation, one lane is enough for the protocol related information and the other lanes can be unidirectional, saving power and mass.

The lane layer initialisation state machine was designed for a bidirectional lane, however some additional rules can be defined to allow a unidirectional lane to reach active state without affecting the operation for bidirectional lanes. The state machine just needs to know if the lane is receive or transmit only and if the far end has the lane in active state.

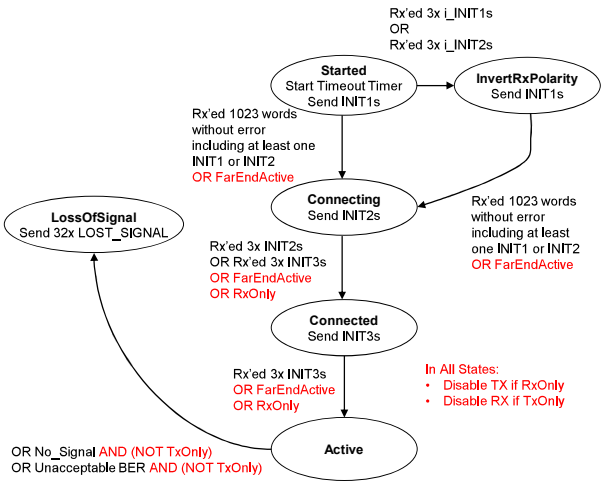

Figure 9. Changes to the lane initialisation state machine

Fig 9 shows in red the required modifications to the lane initialisation state machine. First, if the lane is receive only, the RxOnly condition is set and the state machine immediately moves to Active after reaching the Connecting state. Second, if an ACTIVE word is received after the lane is started indicating that the lane is active at the far end, the FarEndActive condition is set, and the state machine moves to Active. Finally, the LossOfSignal state is not reachable if the lane is TxOnly.

Fig 10 shows how a unidirectional lane is initialised. The side configured as RxOnly, receives INIT1 words until it reaches Connecting state. It then immediately moves to Active. Then, the bidirectional lane(s) send ACTIVE words, which when received, sets the FarEndActive flag. This moves the TxOnly initialisation state machine to ACTIVE.

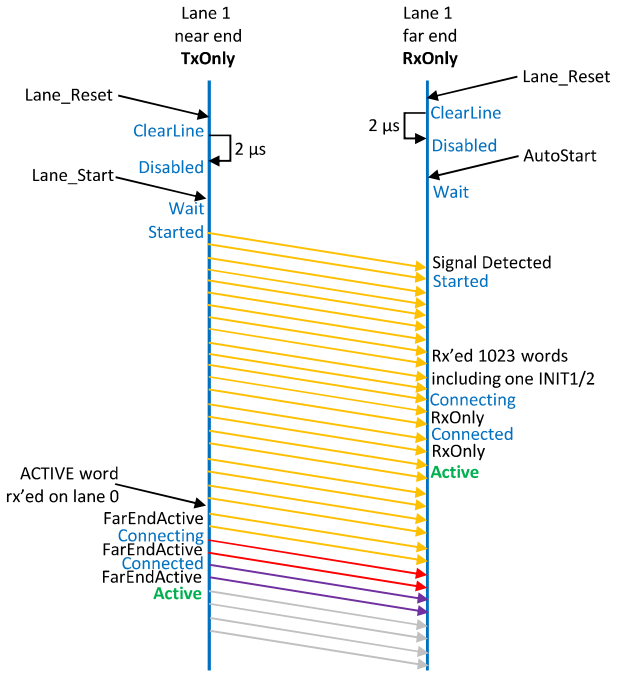

Figure 10. Initialisation of a unidirectional lane

The TxOnly side can not detect loss of signal or receive control words. A mechanism has to be defined in TxOnly lanes to exit the Active state when the RxOnly far end is not anymore in Active state. The solution is to reset the TxOnly lane when ACTIVE words are received in bidirectional lanes indicating that the far end is not anymore active, i.e. the FarEndActive flag is deasserted.

One requirement for unidirectional lanes states that it must be possible for a bidirectional lane to become unidirectional in order to save power. More precisely, a bidirectional lane can be set as RxOnly lane by the user when the data rate sent is reduced. A mechanism is required for the far end to detect this event and change from a bidirectional lane to a TxOnly lane, so it matches the RxOnly setting at the near end.

The solution is for the TxOnly flag to be set when the lane initialisation state machine is in Started state and the FarEndActive flag is set. The FarEndActive flag will move the state machine to Active and the lack of signal at the receiver will be ignored as the TxOnly flag will be set.

E. Hot redundant lanes

An important requirement is the decoupling between the link bandwidth provided by the number of active lanes and the maximum data rate of the end user interface. There are two possibilities:

a) The available link bandwidth is lower than the user interface. This can occur if one or more lanes fails or are disabled. The user interface flow control will limit the data rate of the user.

b) The available link bandwidth is higher than the user interface. This can be useful to provide hot redundancy.

In order to simplify the implementation of the second scenario, the concept of hot redundant lanes is introduced. Hot redundant lanes are lanes that are initialized in the same way than a normal lane, but only send Lane Layer and Multi-Lane layer control words and do not send any Data Link layer word. When no control words must be sent, they send a PRBS sequence that is generated in the same way than the PRBS data words of Idle frames. This mechanism ensures that the word transfer rate between the Multi-Lane layer and the Data Link layer does not exceed the maximum user interface data rate.

Hot redundant lanes must have lane numbers higher than the other lanes. The receiver can identify a hot redundant lane by the content of the ALIGN word received. An ALIGN word sent by a hot redundant lane has the LANES and the iLANES fields both set to zero, which can not occur for non redundant lanes. Hot redundant lanes identified by the receiver are not considered for the reception of Data Link layer words.

F. New control word fields

The addition of multi-lane capabilities requires two new fields in existing control words that do not break compatibility with single-lane implementations:

a) FCT multiplier: Allows to reduce the number of FCTs sent when the data frame size is increased.

b) Multi-Lane capable flag: provided in the INIT3 control word to indicate that the lane is part of a multi-lane link.

IV. Hardware Implementation

After the new specifications that enable multi-lane capabilities to SpaceFibre were successfully simulated in software, a hardware prototype was built using commercial off-the-shelf (COTS) and space-qualified FPGAs. The new multi-lane capable STAR-Dundee SpaceFibre IP Core is an optimised and improved version.

A. Hardware Prototypes

The multi-lane specifications were first evaluated using the STAR-Dundee SpaceFibre PXI board, which has a set of flexible interface connectors that can be used to customise the board, such as SpFi, SpW and external triggers, etc [7].

Fig 11 shows a multi-lane link using two PXI boards with one bidirectional lane and two unidirectional lanes. Each connector has two activity LEDs. If the upper LED is red it indicates that the receiver is disabled. If the lower LED is red it indicates that the transmitter is disabled. Blue colour indicates data transfer. In addition to the SATA laboratory cables used for SpaceFibre, there are two SpaceWire blue cables used for device configuration.

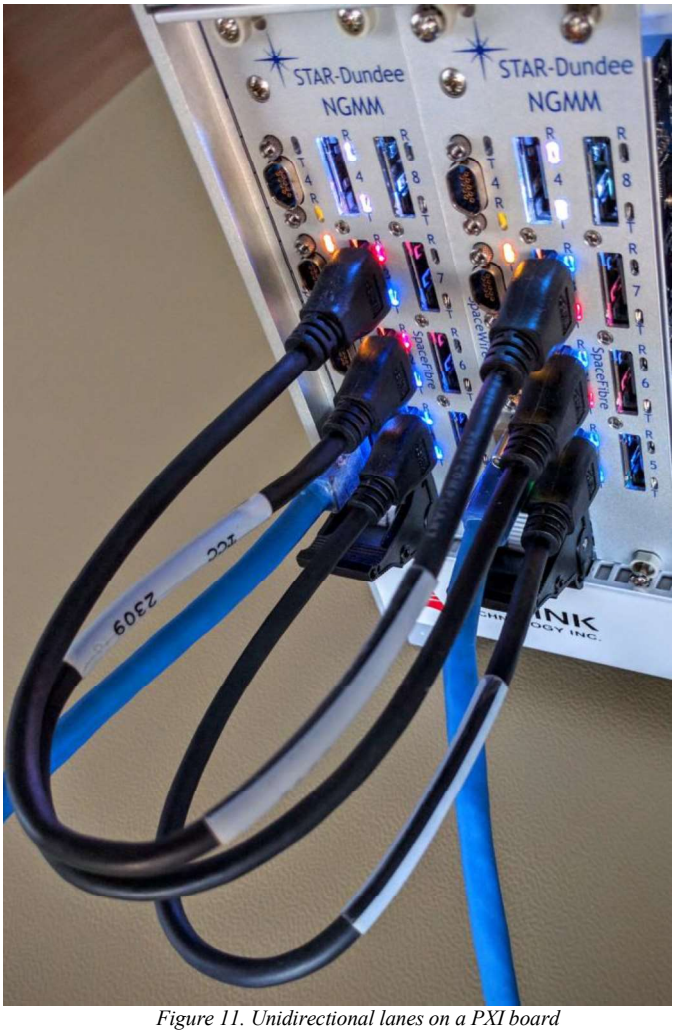

Figure 11. Unidirectional lanes on a PXI board

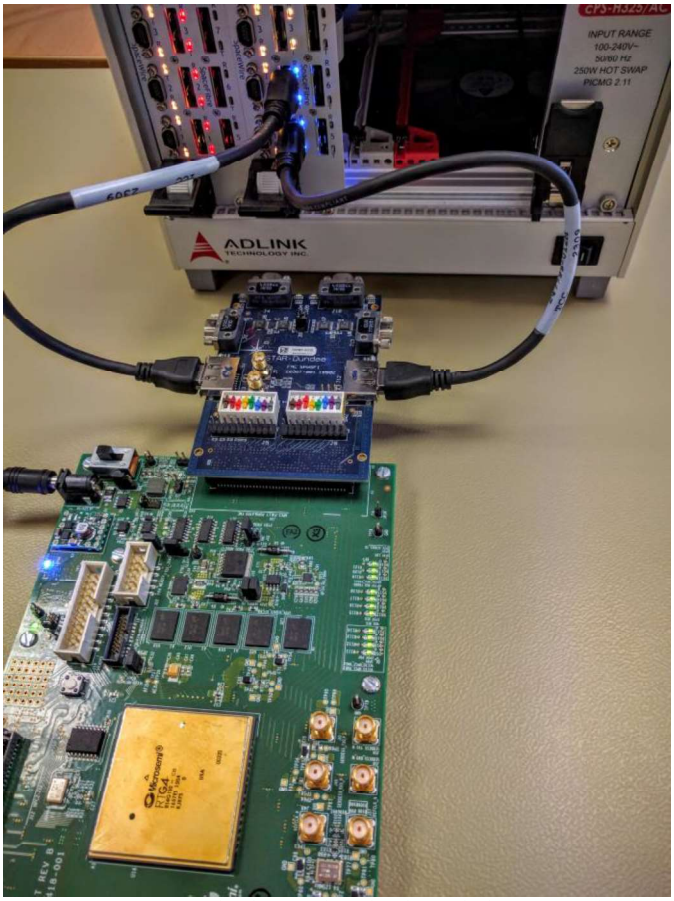

Figure 12. RTG4 development board with a multi-lane SpaceFibre link

The design was then optimised and ported to the radiation hardened RTG4 FPGA. Fig 12 shows the RTG4 development board with a multi-lane link connected to a PXI board using two bidirectional lanes.

For the validation of the new protocol capabilities, the STAR-Fire software was updated to support the new multi-lane capable features. Fig 13 is a screenshot of the the STAR Fire Analyser view. The two middle columns shows the words being sent by lane numbers 1 (left) and 0 (right). At each side the word is decode in its symbol components. The analyser was triggered on the event of the first ACTIVE word sent after the link was started with three lanes. As stated, ACTIVE words are sent when a lane becomes active until alignment is achieved.

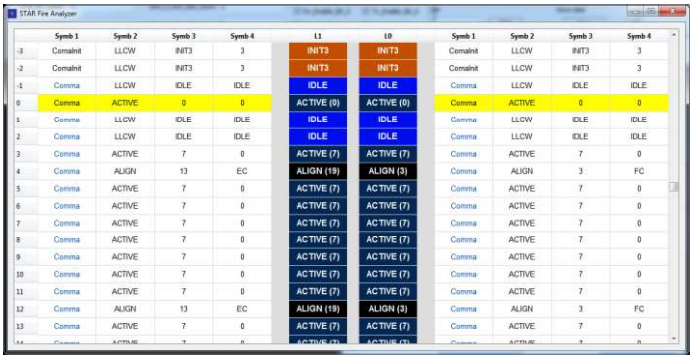

Figure 13. Words sent when lanes 0 and 1 become active

The left side of Fig 14 shows what happens later when alignment is achieved and the alignment state machine moves from Not Ready to Near-End Ready. It is allowed then to send Data Link layer words such as the FCT control words. The ACTIVE words indicate that the first three lanes are active (7_{16}, 111_2). The ALIGN word indicates that three lanes are used for sending and the lane number of each lane (19_{10}, 13_{16}). The right side shows sometime later when data frames are sent with user data from virtual channel 1. At this time, the alignment state machine is in Both Ends Ready as no ALIGN words are being sent.

Figure 14. Words sent for alignment (left) and sending data (right).

Fig 15 shows a link with two lanes in which lane 1 fails and starts sending LOS control words before disabling the SerDes. The alignment state machine moves to Not Ready and ACTIVE words sent indicate lane 1 is not anymore active. After re-alignment, a NACK control word followed by a RETRY control word are sent, so both ends can resume sending data.

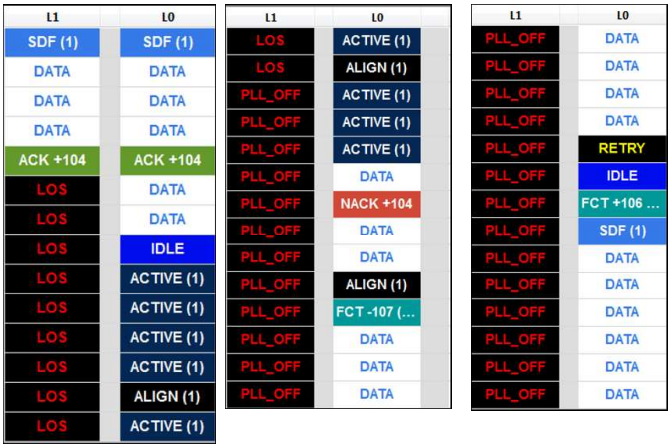

Figure 15. Words sent when a lane fails (from left to right)

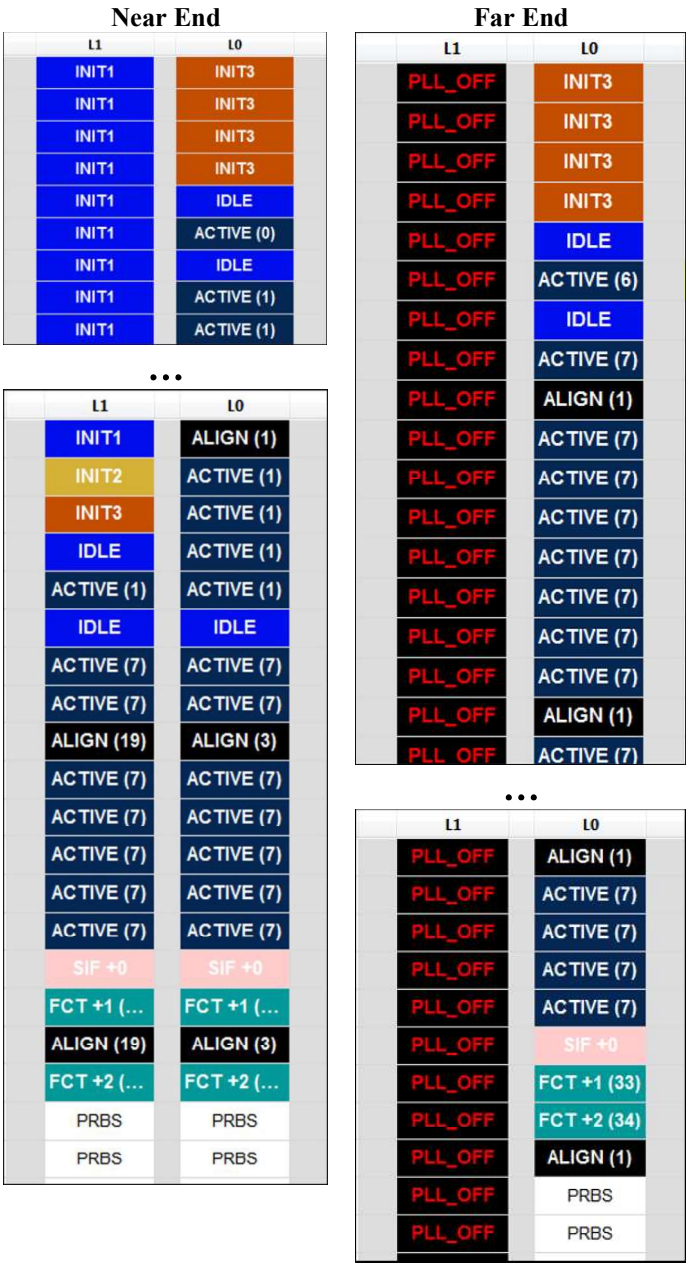

Figure 16. Link starts with lane 1 unidirectional with TxOnly set

Fig 16 shows a multi-lane link with three lanes in which lane 1 is a unidirectional transmit-only lane. That is why the far end has the SerDes transmitter of lane 1 disabled (PLL_OFF). In the near end, this lane starts sending INIT1s until the far end achieves Active state and sends ACTIVE words indicating this lane is active. The near end sends INIT2 and INIT3 and reaches active state on the reception of these ACTIVE words.

B. STAR-Dundee SpaceFibre IP Core

The STAR-Dundee SpaceFibre IP Core was updated to support multi-lane capabilities after the hardware implementations were successfully validated and optimised for low resource usage and easy of use.

Table I provides the resource usage for two radiation hardened FPGAs, Microsemi RTG4 and Xilinx Virtex-5QV,

for different number of lanes and virtual channels. Lanes can operate up to 3.125 Gbps.

TABLE I. RESOURCE USAGE

	RTG4			Virtex-5QV		
	LUT	DFF	RAM Block	LUT	DFF	RAM Block
2 Lanes 1 VC	6494 4.3%	5351 3.5%	8 3.8%	3858 4.7%	3938 4.8%	8 2.7%
2 Lanes 2 VC	7314 4.8%	6088 4.0%	12 5.7%	4503 5.5%	4382 5.3%	12 4.0%
3 Lanes 2 VC	8997 5.9%	7413 4.8%	12 5.7%	5416 6.6%	5226 6.4%	12 4.0%

The IP Core has been designed to fully support the redundancy capabilities of multi-lane. When using hot redundancy, the data flow of the user is not affected when a lane fails, as the data is internally buffered during the time it takes to resume sending data, which is less than 2 µs. When not using a hot redundant lane, there is a graceful degradation of link bandwidth and the QoS mechanism ensures that most important data is sent first. If a redundant lane is available it will be activated in less than 20 µs, providing warm redundancy.

Fig 17 shows the floorplan of a Virtex-5QV with the IP core constrained to be placed in one of the tiles. Using the Xilinx transceiver capabilities, the IP Core can work with a single clock input signal. The user can write and read data to/from the IP Core with the AXI4-Stream interface, using any other clock frequency as the IP includes synchronisation buffers.

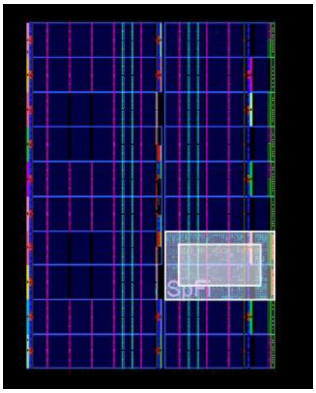

Figure 17. STAR-Dundee multi-lane IP Core in Virtex-5QV

V. CONCLUSION

The new SpaceFibre multi-lane capabilities increase dramatically the data throughput of SpaceFibre links to meet the requirements of next generation of spacecraft payloads.

With the designed multi-lane layer, the additional lanes can also provide hot or warm redundancy, and graceful degradation of the link bandwidth when no redundant lanes are available. In the event of a lane failure, the link is again operative in just a few microseconds, which is close to the round trip delay of the lane, without user intervention and without any data loss.

Furthermore, the flexibility in the number of lanes of a multi-lane link and the support of unidirectional lanes, allows for significant savings in mass and power, which are critical in space applications.

The multi-lane specifications have been validated in simulation and hardware prototypes. These specifications have been designed to be easy to implement in slower radiation hardened FPGAs. The STAR-Dundee SpaceFibre IP Core has been updated to provide all these new multi-lane capabilities in RTG4 and Virtex-5QV FPGAs with low resource usage and high performance.

ACKNOWLEDGMENT

The research leading to these results has received funding the European Space Agency under ESA contract numbers 4000102641 and from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 263148 and 284389.

REFERENCES

[1] Next Generation Processor for On-board Payload Data Processing Application ESA Round Table Synthesis, ESA, TEC-EDP/2007.35/RT, October 2007.

[2] S. Parkes, A. Ferrer, A. Gonzalez and C. McClements, "SpaceFibre Standard", Draft H4, April 2016, available from https://indico.esa.int/indico/event/126/session/0/contribution/1 (last accessed 29th August 2016).

[3] T. Scheckel, "Serial RapidIO: Benefiting system interconnects", Proceedings - IEEE International SOC Conference, pp. 317-318, 2005.

[4] J. Ajanovic, "PCI express 3.0 overview", 2009 IEEE Hot Chips 21 Symposium (HCS), Stanford, CA, 2009, pp. 1-61.

[5] "Interlaken Protocol Definition", A joint specification of Cortina Systems and Cisco System, Revision 1.2, Cortina Systems Inc. and Cisco Systems Inc., 7 October 2008.

[6] Y. Otake, "The study and proposal for improvement the multi-lane operation of SpaceFibre protocol", SpaceWire Conference 2014, Athens.

[7] A. Gonzalez, "A new Generation of SpaceFibre Test and Development", SpaceWire Conference 2016, Yokohama.