SpaceFibre Networks

SpaceFibre, Long Paper

Steve Parkes, Chris McClements, David McLaren,
Space Technology Centre, University of Dundee,
166 Nethergate, Dundee, DD1 4EE, UK
smparkes@dundee.ac.uk

Albert Ferrer Florit, Alberto Gonzalez Villafranca,
STAR-Dundee Ltd.,
STAR House, 166 Nethergate, Dundee, DD1 4EE, UK

***Abstract*— SpaceFibre [1][2][3] is the next generation of SpaceWire [4] on-board data-handling network technology for spaceflight operations, which runs over both electrical and fibre optic media. SpaceFibre has many benefits compared to SpaceWire, including much higher data-rates, integrated quality of service, fault recovery capabilities, multi-laning with graceful degradation and hot and cold redundancy, and low-latency broadcast messages that can carry 8-bytes of user information. Importantly SpaceFibre is backwards compatible with SpaceWire at the network level, allowing existing SpaceWire equipment to be incorporated into a SpaceFibre network without modification. SpaceFibre networks have been defined by the University of Dundee and STAR-Dundee, and incorporated in the network layer definition of the current draft SpaceFibre standard. STAR-Dundee has designed a SpaceFibre routing switch to evaluate various routing concepts, validate the standard specification and demonstrate a complete SpaceFibre network. A demonstration system has been built and key parts of the SpaceFibre network technology have been demonstrated.**

***Index Terms* — SpaceFibre, SpaceWire, Networking, Spacecraft Electronics.**

I. Introduction

SpaceFibre is the next generation of SpaceWire technology for spacecraft on-board data-handling. It is able to operate at multi-Gbits/s over distances of up to 5 m using electrical cable and 100 m using fibre optic cable. It is galvanically isolated, includes quality of service and fault detection, isolation and recovery capabilities. SpaceFibre is backwards compatible with SpaceWire at the Network level, which enables existing SpaceWire equipment to be connected into a SpaceFibre network without modification. Furthermore SpaceFibre has been designed to have a small footprint, enabling its implementation in flight qualified FPGAs and ASIC devices without using a large part of the device.

This paper outlines the operation of SpaceFibre networks, describes the SUNRISE SpaceFibre routing switch, and summarises the results of tests with this routing switch.

II. SpaceFibre Links

A. Links and Lanes

A SpaceFibre link is made up of one or more lanes, which carry information from one end of the link to the other. SpaceFibre lanes can run over an electrical or fibre optic physical layer. In a multi-lane link, some of the lanes can be unidirectional provided that at least one lane is bi-directional [5]. The SpaceFibre link provides quality of service and error recovery [3].

B. SpaceFibre Virtual Channels

SpaceFibre links carry traffic (application information) through one or more virtual channels. There is a maximum of 32 virtual channels on a link, which are numbered consecutively starting at 0. Traffic entering virtual channel N comes out of virtual channel N at the other end of the link.

Each virtual channel is provided with a quality of service (QoS) which has three components: bandwidth reservation, priority and scheduling. Bandwidth reservation, reserves a portion of the link bandwidth for the virtual channel. Priority assigns a priority-level to the virtual channel so that higher priority virtual channels 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 virtual channel is permitted to send information. These three different QoS components are not alternatives, they work together. [3]

III. SpaceFibre Networks

In this section the operation of a SpaceFibre network is described.

A. SpaceFibre Packets

SpaceFibre packets are identical to SpaceWire packets. They are formed from data characters, end of packet markers, and error end of packet markers, as illustrated in Figure 1.

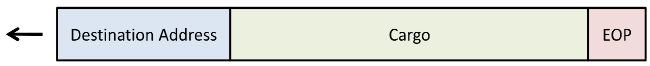

Figure 1 SpaceWire Packet Format

The "Destination Address" is the first part of the packet to be sent and is a list of data characters that represents either the identity of the destination node or the path that the packet has to take through a SpaceFibre network to reach the destination node. In the case of a point-to-point link directly between two nodes (no routers in between) the destination address is not necessary.

The "Cargo" is the data to be transferred from source to destination. Any number of data bytes can be transferred in the cargo of a SpaceFibre packet.

The "End_of_Packet" (EOP) is used to indicate the end of a packet. The data character following an End_of_Packet is the start of the next packet. There is no limit on the size of a SpaceFibre packet. “Error End of Packet” (EEP) is a form of

EOP which is used to indicate the premature end of a packet due to the occurrence an error.

B. SpaceFibre Virtual Networks

A SpaceFibre network is effectively a set of independent parallel SpaceWire networks. These parallel, independent networks are called “SpaceFibre virtual networks”. Each virtual network runs over its own, distinct set of SpaceFibre virtual channels, comprising a virtual channel across each link used by the virtual network. Several virtual networks can then operate concurrently over a single physical SpaceFibre network. The overall physical network and the collection of virtual networks that run over that physical network is called the “SpaceFibre network”.

The traffic running over each virtual network is constrained by the SpaceFibre quality of service mechanism to remain within its allocated bandwidth and to observer the priority and schedule allocated to it. A virtual network is able to opportunistically use more bandwidth than it has been allocated, when no other virtual network has traffic to send over the links of the SpaceFibre network that the particular virtual network wants to use.

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

C. Packet Addressing

SpaceFibre uses both path and logical addressing, which operate in the same way as SpaceWire. It is not possible to route a packet between two different virtual networks in a routing switch. As already stated virtual networks on a SpaceFibre network are like a set of parallel, independent SpaceWire networks. The packet routing is within one virtual network.

Path addressing uses the leading data character of a packet to determine how the packet should be routed at the next routing switch. If the value of the leading data character is in the range 0 to 31, it determines which port of the routing switch the packet will be forwarded through. For example, if the leading data character is 2, the packet will be forwarded through port 2 of the routing switch. If the leading data character is 0, it will be routed to port 0, the internal configuration port of the routing switch. If the leading data character is 31 and there are only 9 ports in the router, the packet will be discarded. Note that the ports of a router are number consecutively, starting at 0 for the internal configuration port.

If the leading data character is in the range 32-255, it is a logical address. The value of the leading data character is then used as the index into a routing table, which once configured, determines which port the packet is to be forwarded through. For example, if the leading data character is 40 and the entry in the routing table for index 40 contains the value 3, the packet will be routed to port 3 of the router. The routing table is configured using RMAP commands sent to the router configuration port [6]. Before configuration of the routing table has been done, any logical address will result in the packet being discarded. Path addressing operates at all times, before and after the routing table has been configured.

D. Fills

SpaceFibre runs much faster than SpaceWire, so requires an interface to the application which is wider than that of SpaceWire to carry the extra data. The interface to a SpaceFibre port is typically 32-bit wide or a multiple of 32-bits, whereas SpaceWire is 8-bits wide. If SpaceFibre is to send a packet which is not a multiple of 32-bits, the start of the packet or its tail end can be filled with Fill characters to make it 32-bit aligned. Therefore, a SpaceFibre data word contains four data characters, EOPs, EEPs or Fills. The use of Fills is illustrated in Figure 2 and Figure 3, where P represents a path-address data character, D represents a data character, E an EOP or EEP, and F a Fill.

Filling the start allows for a 32-bit aligned cargo, when path addressing is being used, as illustrated in Figure 2.

F	F	F	P
P	P	P	P
D	D	D	D
D	D	D	D
E	F	F	F

Figure 2 Fills at the start of a SpaceFibre packet

Fill characters are added at the beginning of a packet, to align a path address which is not a multiple of four data characters in length or to fill spaces that were previously occupied by a path address. This allows the leading SpaceFibre path address bytes to be removed by a router and replaced by Fill characters in order to keep the word-alignment of the SpaceWire cargo when it arrives at the destination. It also allows some fills to be added to the start of a packet to ensure that the cargo of the packet is 32-bit aligned when there is a path address that is not a multiple of four data characters.

Filling the end allows for the cargo to be any number of N-Chars, not a multiple of four N-Chars, as illustrated in Figure 3.

D	D	D	D
D	E	F	F
F	F	P	P
D	D	D	D
D	D	D	D
E	F	F	F

Figure 3 Fills at the end of a SpaceFibre packet

The Fill character is used in a data word containing an EOP or EEP to fill otherwise empty characters that follow the EOP or EEP. The above example shows two small packets in part of a frame being aligned to 32-bits.

E. Virtual Network Masters

A “network master” is a node on a SpaceFibre virtual network which is a source of SpaceFibre packets able to send packets autonomously, i.e. without first receiving a request from another node. Note that a network master is different to a network manager, the latter is a network master that configures, controls and monitors the status of the entire SpaceFibre network.

If there is one network master on a virtual network then that virtual network can be deterministic. For example, the network master might be a control processor sending Remote Memory Access Protocol (RMAP) packets to other instrument nodes to control them and collect data from them, using RMAP. The traffic on the virtual network is controlled by the one network master node. The set of virtual channels that the specific virtual network runs over is allocated the bandwidth and priority according to its needs. If the virtual network is to provide time-bound determinism, its virtual channel will also be scheduled by the SpaceFibre QoS mechanism.

Within a single SpaceFibre virtual network, if there are two independent network masters, it is possible that they both send a packet to the same node, or through the same link to a router and then on to different nodes. Whenever these two network masters want to send a packet over the same link at the same time, there is a "collision" and one packet will have to wait for the other one to be sent. This is the same as the temporary "packet blocking" that can occur in a SpaceWire network. Each SpaceFibre virtual network operates just like a separate SpaceWire network, including temporary packet blocking.

Now, in some applications the temporary network blocking was a real pain in a SpaceWire network, especially if long packets were being used. Traffic from one application could delay traffic from another one, which could be difficult to handle under some circumstances. SpaceFibre solves this problem, by having multiple, independent virtual networks. If there is a single network master on each of these virtual networks, the packet blocking is avoided completely. It is still possible to have multiple network masters on the same virtual network, provided that packet blocking is not an issue for the traffic flowing over that virtual network, or provided that another mechanism is used to control the flow of traffic over that network.

This approach maintains full backwards compatibility with SpaceWire at the network level, which is essential if the large legacy of existing SpaceWire equipment is not to be squandered. Reuse of existing, proven equipment, reflected by the Technology Readiness Level (TRL), is an important way of improving reliability and reducing the cost of space missions. SpaceFibre offers a path for substantially upgrading the capabilities and performance of an onboard network without losing that valuable legacy.

IV. Reference Architecture

It is worth considering an example of how the virtual networks might be used in a typical space mission. First, a reference architecture is described.

A. Earth Observation Reference Architecture

A reference architecture has been devised which is representative of a typical high data-rate Earth Observation mission. This architecture is illustrated in Figure 4.

Instruments 1 and 2 are high data-rate instruments each with a SpaceFibre interface. They are connected via two SpaceFibre routers to the mass-memory unit which has two SpaceFibre interfaces. Each instrument is able to transfer data at up to 2 Gbits/s using a 2.5 Gbit/s SpaceFibre link.

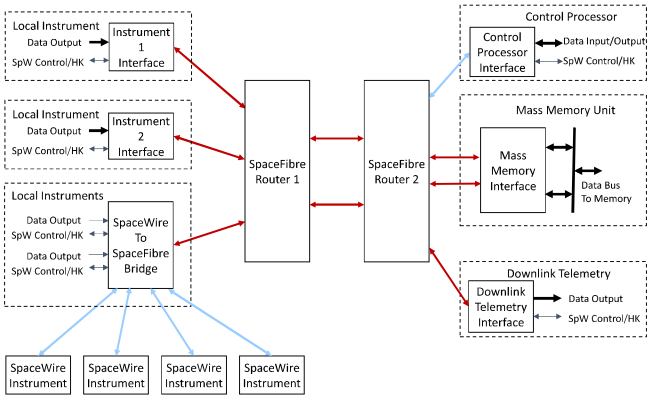

Figure 4 SpaceFibre Earth Observation Mission Reference Architecture

Four existing SpaceWire instruments are attached to a SpaceWire to SpaceFibre bridge device, each via a separate virtual channel of the SpaceFibre interface. Data from these SpaceWire devices is sent over the SpaceFibre network to the mass-memory unit.

Data from the mass-memory unit is passed to the downlink telemetry unit.

A control processor is able to access all of the instruments, the mass-memory unit and the downlink telemetry unit along with the SpaceFibre routing switches to configure and control the devices and to read housekeeping information from them.

The architecture in Figure 4 does not really need two routing switches, but two are included in the reference architecture to make it more generic.

B. Example Allocation of Virtual Networks

There are several functions that need to be carried out by the reference architecture of Figure 4. These functions are listed below:

1. SpaceFibre network management: configuring, monitoring and reconfiguring the SpaceFibre network;
2. Payload management; instrument control and status monitoring (housekeeping);
3. Data-handling system management; control and status monitoring (housekeeping) of the mass-memory unit and the downlink telemetry unit;
4. Sending data from the high data-rate Instrument 1 to the mass-memory unit
5. Sending data from the high data-rate Instrument 2 to the mass-memory unit;
6. Sending data from the four SpaceWire instruments to the mass-memory unit;
7. Sending data from the mass-memory unit to the downlink telemetry unit.

Each of these functions could be allocated a separate virtual network, requiring a total of seven virtual networks in the routing switches and mass-memory unit. Since there is only one control processor (ignoring a possible redundant unit), it is necessary to run the SpaceFibre network management, the payload management and the data-handling functions on the same processor. These functions can then share a virtual network since there will always only be the one control processor using that virtual network. This reduces the number of virtual networks required to five. The five parallel virtual networks are illustrated in Figure 5.

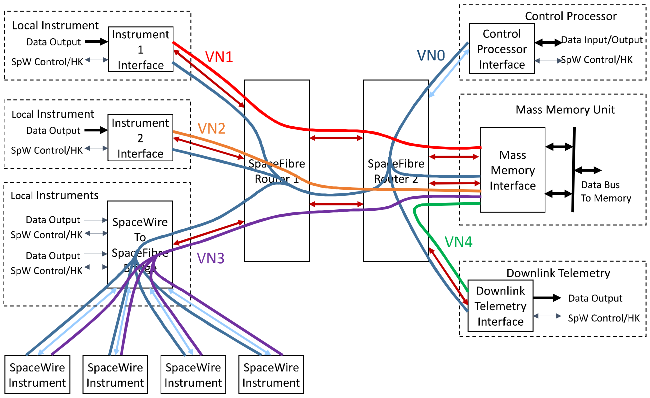

Figure 5 Parallel Virtual Networks

The control processor performing the network management, payload management and data-handling system management, uses one virtual network (VN0) and is able to access all of the instruments, routing switches, mass-memory and downlink telemetry units, over that one virtual network.

Instrument 1 uses another virtual network (VN1) to send data to the mass-memory unit. Similarly instrument 2 uses VN2.

The SpaceWire instruments all share one virtual network (VN3) for sending data to the mass-memory unit. This means that they will compete for access to the virtual network, as if they were running over a SpaceWire network.

C. Networks with Large Number of Nodes

When there are a large number of nodes in a network, it is possible to handle them in several different ways.

Firstly, a single virtual network could be used for several nodes which all act as network masters. It is simply accepted that within this virtual network temporary packet blocking will occur and will not be a problem for the applications related to those nodes. This virtual network operates the same as a SpaceWire network

Secondly, it is possible to increase the number of virtual channels so that there is one for each SpaceWire instrument. This depends on the number of virtual channels available in the SpaceFibre routers and mass-memory unit. In any case there is a limit to the maximum number of virtual networks that can be used. There is actually a maximum of 32 virtual channels over a link and 64 virtual networks across a SpaceFibre network.

A third alternative is to use one network master on a virtual network to handle all the communication for the nodes on that network. The configuration, control and housekeeping network is an example of this where there is one master node that uses RMAP commands to request information to all the nodes on the network including the configuration nodes within the routing switches.

A similar approach could be used for sending data from several instruments to the mass-memory unit. The mass interface controller could send out RMAP commands to request data from each of the SpaceWire instruments on a single virtual network in turn. For example the mass-memory unit could use VN3 to send RMAP commands to the SpaceWire instruments which respond with the requested data, which is then placed in memory.

Another possibility is to schedule the sending of information from the various equipment over a virtual network using time-slots, which are delimited by broadcast messages over the SpaceFibre network or time-codes on the SpaceWire network. Each equipment then sends its data in its allocated time-slot or time-slots.

D. Virtual Network to Virtual Channel Mapping

Virtual networks are mapped on to a set of virtual channels, one virtual channel for each link used by the virtual network. Each virtual channel on a link is mapped to one and only one virtual network. The virtual channel number used by a virtual network over one link does not need to be the same as the virtual channel number used on another link.

The simplest way of mapping a virtual network to a virtual channel is to use a one to one mapping, so that virtual network VN0 uses virtual channel VC0 on all of the links in the network. Similarly VN1 uses VC1 and so on. The problem with this simple approach is that it complicates the instrument nodes of the network. For example, a typical instrument will require two virtual networks; VN0 which is used for control and monitoring and another virtual network which is used for data transfer to a mass-memory unit. This is the case with instruments 1 and 2 in Figure 5, which use VN1 and VN2 respectively. If a mapping is done from the virtual network to the virtual channels, the hardware required in the instrument interfaces is simplified. For example, instrument 1 VN1 is mapped to VC1 and instrument 2 VN2 is mapped to VC1. This mapping needs to be done at both ends of the respective links. The routing switch then uses this mapping to route a packet to an output port on the same virtual network number as that on which the packet arrived. The virtual channel numbers may be different on the link over which the packet arrived and the link over which the packet is being forwarded, but the virtual network numbers mapped to these virtual channels are the same.

Using the example of Figure 5, the links running from Router 1 to Router 2 will carry instrument data from instrument 1 over VN1 and from instrument 2 over VN2. This data can go over either of the links between the two routing switches depending on the packet address. So over these links the following mapping applies:

- VN0 -> VC0, this is always the case
- VN1 -> VC1
- VN2 -> VC2
- VN3 -> VC3

VN4 does not use the links between the routers.

So VN2 is mapped to VC1 for the link from instrument 2 to router 1, because there are only two virtual channels available in the instrument interface. VN2 is then mapped to VC2 over the links from router 1 to router 2.

The virtual network to virtual channel mapping makes the routing switches more complex, because it has to handle the mapping, but makes the instrument interfaces simpler, because they normally only need two virtual networks, which can be supported by two virtual channels. The virtual network mapping also permits more virtual networks on a SpaceFibre networks than there are virtual channels on a SpaceFibre link, i.e. there are up to 64 virtual networks allowed in a network but only 32 virtual channels over a link. This is possible because some virtual networks may use completely separate parts of the network.

V. SUNRISE SPACEFIBRE ROUTING SWITCH

A SpaceFibre router has been designed and implemented in the SUNRISE project funded by the UK Space Agency and STAR-Dundee. The architecture of this router is shown in Figure 6.

The SUNRISE router has eight SpaceFibre ports, numbered 1 to 8, each with four virtual channels. There is a configuration port (port 0) which is used for device configuration and which can be accessed using virtual channel 0 of any of the other ports. Another port (port 9) provides an interface to four SpaceWire ports using four virtual channels, one for each SpaceWire port. SpaceWire and SpaceFibre packets are switched by the routing switch in the same way, using the leading data character of a packet to determine the output port that the packet is to be switched to. Both path and logical addressing can be used with the SUNRISE router.

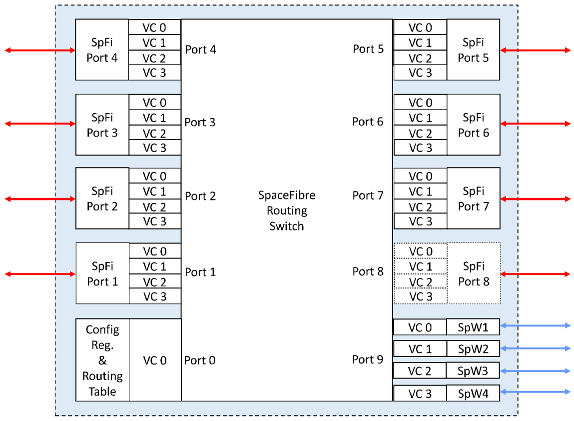

Figure 6 SUNRISE SpaceFibre Router Architecture

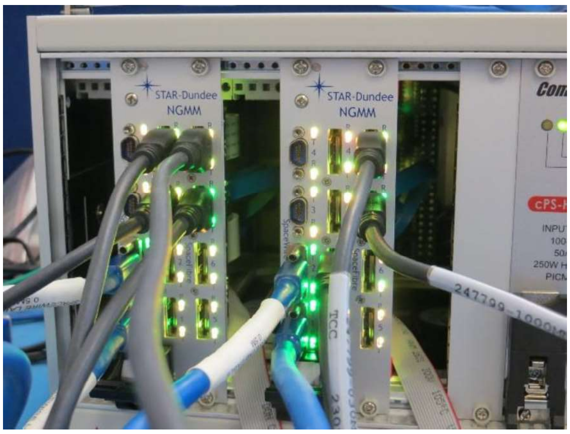

Figure 7 SUNRISE SpaceFibre Routers Under Test

The SUNRISE router was implemented initially in a Xilinx Spartan 6 FPGA. Two of the SUNRISE routers are shown under test in Figure 7. The SUNRISE routers are implemented on 3U cPCI/PXI boards. Power is taken from the backplane and the eight SpaceFibre and four SpaceWire ports are available on the 40mm wide front panel.

The SUNRISE router is now being implemented in a Microsemi RTG4 FPGA as shown in Figure 8 [7][8].

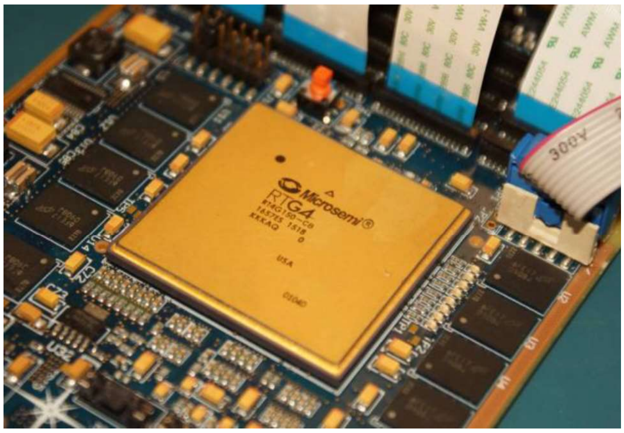

Figure 8 Prototype Microsemi RTG4 board for SUNRISE SpaceFibre Router

VI. DEMONSTRATION OF SPACEFIBRE NETWORK

The reference architecture has been implemented using a combination of radiation tolerant FPGAs and commercial FPGAs. This is illustrated in Figure 9.

Figure 9 SpaceFibre Network Demonstration

The equipment used in the demonstration system is detailed in Figure 10 [9].

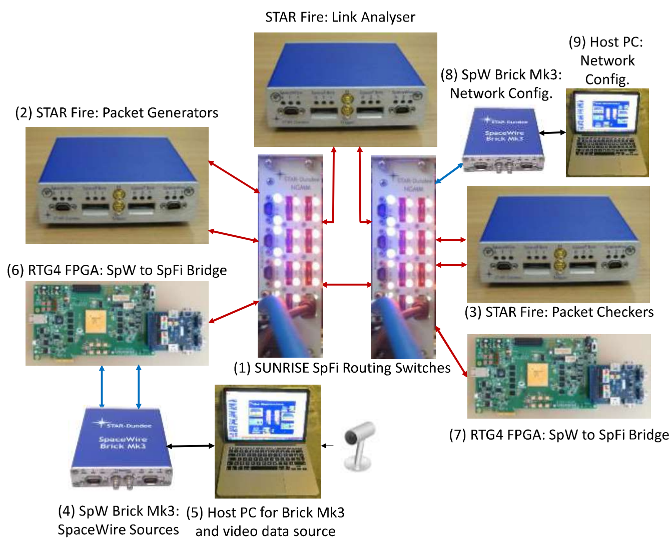

Figure 10 SpaceFibre Network Demonstration System

The following functions were demonstrated and validated:

- **SpaceFibre network operation:** using the two SUNRISE routing switches (1). Packets were successfully routed across the routing switches, remaining in their virtual networks.
- **High data-rate:** The STAR Fire unit has two SpaceFibre interfaces and incorporates packet generators that are able to generate SpaceFibre packets at the full data rate (2.5 Gbits/s) over each SpaceFibre link. STAR Fire unit (2) was used to simulate the two high data-rate instruments of the reference architecture. Data from these two "instruments" was sent across the network to the STAR Fire unit (3), which accepts and checks the high-data rate packets. STAR Fire unit (3) is acting like a mass memory accepting data from the high data-rate instruments.
- **SpaceWire to SpaceFibre bridging:** A SpaceWire Brick Mk3 (4) was used to generate two streams of SpaceWire packets under control of the host PC (5). The SpaceWire links are attached to a Microsemi RTG4 development board via a STAR-Dundee FMC board [7]. The RTG4 is programmed with a SpaceWire to SpaceFibre bridge design connecting four SpaceWire interfaces to four virtual channels of a SpaceFibre interface. One SpaceWire link is sending video data from a webcam attached to the host PC (5). The other SpaceWire link is sending packets from a SpaceWire packet generator running on host PC (5) to another PC (9) so that they can be checked for errors. The SpaceWire packets are converted to SpaceFibre packets, which is trivial as they have the same format, and sent across the SpaceFibre network. The video data is sent to another RTG4 board (7). The other data is sent via port 9 of a SpaceFibre router (1) which is a port where the virtual channels are connected to SpaceWire interfaces. This SpaceWire data goes across a SpaceWire link to another Brick Mk3 (8) and on to a host PC (9) where it is checked.
- **Quality of Service:** The STAR Fire packet generators (2) provide a total data rate of 2 x 2.5 Gbits/s, using all the network bandwidth between the two routing switches (1). The virtual channels they are using are assigned relatively low priority. The SpaceWire to SpaceFibre bridge (6) uses a virtual channel with higher priority. Whenever it wants to send data, it is able to do so, within the constraints of its allocated bandwidth. This is demonstrated by the real-time video data stream being transferred across the network.
- **Fault detection, isolation and recovery:** the link being used to transfer the traffic from the SpaceWire to SpaceFibre Bridge (6) between the two routing switches (1) can be unplugged. The video traffic then stops and the SpaceWire packet generated data stops. When the link is plugged back in the SpaceWire packet generated data continues and there is no loss of packets. The packets are checked for errors including missing packets in the host computer (9). While the link was disconnected no packets could be transferred but the packet being transferred when the link was disconnected was not lost. Clearly with the video data stream, data is lost once the buffers in the system are filled. The key point is that packets in transit across the network are not lost.
- **Network configuration:** the network is configured using the host computer (9) via a SpaceWire connection to the right hand routing switch (1).

For debugging and analysis purposes, a third STAR Fire unit (10) operating as a link analyser is included on one of the links between the two routing switches (1) [10].

VII. Conclusions

SpaceFibre networks have been defined by the University of Dundee and STAR-Dundee, and incorporated in the network layer definition of the current draft SpaceFibre standard. STAR-Dundee has designed the SUNRISE SpaceFibre routing switch to evaluate various routing concepts, validate the standard specification and demonstrate a complete SpaceFibre network. A reference architecture for a SpaceFibre network targeted at Earth Observation applications has been defined. A demonstration system has been built reflecting this reference architecture and key parts of the SpaceFibre network technology have been demonstrated.

Acknowledgments

The research leading to these results has received funding from the European Space Agency under ESA contract numbers 4000102641 and 17938/03/NL/LvH, from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement numbers 263148 and 284389 and from the UK Space Agency and CEOI-ST under University of Leicester contract numbers: RP10G0348A02, RP10G0348B206 and RP10G0348A207.

References

[1] S. Parkes, A. Ferrer Florit and A. Gonzalez Villafranca, "SpaceFibre Standard", Draft H5, University of Dundee, July 2016.

[2] S. Parkes, C. McClements and M. Suess, "SpaceFibre", International SpaceWire Conference, St Petersburg, Russia, 2010, ISBN 978-0-9557196-2-2, pp 41-45.

[3] S. Parkes et al, "SpaceFibre: Multi-Gigabit/s Interconnect for Spacecraft On-board Data Handling", IEEE Aerospace Conference, Big Sky, Montana, 2015.

[4] 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.

[5] A. Ferrer Florit, A. Gonzalez Villafranca and S. Parkes, "SpaceFibre Multi-Lane", International SpaceWire Conference, Yokohama, Japan, 2016, ISBN 978-0-9954530-0-5.

[6] ECSS Standard ECSS-E-ST-50-52C, "SpaceWire – Remote memory access protocol", Issue 1, European Cooperation for Space Data Standardization, 5 February 2010, available from http://www.ecss.nl.

[7] S. Parkes et al, "SpaceWire and SpaceFibre on the Microsemi RTG4 FPGA", IEEE Aerospace Conference, Big Sky, Montana, 2016.

[8] S. Parkes, A. Ferrer Florit, A. Gonzalez Villafranca, C. McClements, B. Yu, P. Scott, J. Logan, D. Dillon and D. McLaren and "SpaceFibre Flight Equipment", International SpaceWire Conference, Yokohama, Japan, 2016, ISBN 978-0-9954530-0-5.

[9] A. Gonzalez Villafranca, S. Parkes, C. McClements, B. Yu, P. Scott and A. Ferrer Florit, "A New Generation of SpaceFibre Test and Development Equipment", International SpaceWire Conference, Yokohama, Japan, 2016, ISBN 978-0-9954530-0-5.

[10] S. Mudie and S. Parkes, "SpaceFibre Link Analysis", International SpaceWire Conference, Yokohama, Japan, 2016, ISBN 978-0-9954530-0-5.