High-Performance, Multi-Lane, SpaceFibre Routing Switch for ADHA

Steve Parkes
STAR-Dundee Ltd.
Dundee, United Kingdom
steve.parkes@star-dundee.com

Alberto Gonzalez Villafranca
STAR-Barcelona S.L.
Barcelona, Spain
alberto.gonzalez@star-dundee.com

Albert Ferrer Florit
STAR-Barcelona S.L.
Barcelona, Spain
albert.ferrer@star-dundee.com

Marti Farras Casas
STAR-Barcelona S.L.
Barcelona, Spain
marti.farras@star-dundee.com

David Steenari
ESTEC, ESA
Noordwijk, The Netherlands
David.Steenari@esa.int

***Abstract*—The Advanced Data Handling Architecture (ADHA) concept has been developed in collaboration between European industry (integrators, and hardware suppliers) and the European Space Agency (ESA). ADHA defines a set of specifications for electronic data-handing units [1]. which is based on standardised, interchangeable, and interoperable electronics modules. SpaceFibre [2] provides the high data-rate interconnect for ADHA, both inside an ADHA unit, interconnecting the AHDA modules, and externally. This paper describes a SpaceFibre routing switch for AHDA hardware modules which is able to support the demanding data rates and interconnect architecture of ADHA applications.**

Keywords— ADHA, SpaceFibre, SpaceFibre Routing Switch, Avionics, Payload Data-Handling.

1. INTRODUCTION

The Advanced Data-Handling Architecture (ADHA) is a specification for end-to-end data-handling on-board spacecraft. It aims to address three concerns about spacecraft development: time reduction, cost efficiency, and faster development and adoption of innovative technologies. ESA in cooperation with European industry has developed the Advanced Data Handling Architecture (ADHA) specification, which is based on standardised, interchangeable, and interoperable electronics modules. Many ADHA-compliant modules are currently under development, utilising the latest generation of microelectronics components for space. Command and Control (C&C) and data exchange between modules is via CAN-Bus, SpaceWire or SpaceFibre dependent on the needed data rates [1][3].

SpaceFibre [2] provides the high data-rate interconnect for ADHA, both inside an ADHA unit, interconnecting modules, and externally, interconnecting an ADHA data-handling unit to instruments/payloads, downlink transmitters and to other AHDA data-handling units. SpaceFibre routing switches interconnect the modules with a flexible architecture which is able to support single-string, unit-level redundancy and redundancy internal to a unit.

This paper introduces ADHA and SpaceFibre. It then explains why SpaceFibre is particularly suitable for satellite payload data-handling applications. It then describes STAR-Dundee SpaceFibre routing switch technology and how it provides the capabilities required by ADHA. Information on resource utilisation, achievable lane rates, and power dissipation for a range of different routing switch configurations is then given.

2. ADHA ARCHITECTURE OVERVIEW

The ADHA specifications includes the following elements:

- Definition of the electro-mechanical form-factor, and utilisation of the backplane connector.
- Definition of module functional specifications for specific data handling functions, e.g. OBCs (On-Board Computers), MMs (Mass-Memories), Power, DPMs (Data Processing Modules).
- A standardised PA/QA (product assurance/quality assurance) flow, based on ECSS standards and ESA mission classifications.

ADHA units (i.e. electronics boxes), consists of multiple ADHA modules, and a backplane providing connections between the modules.

The ADHA specifications define the use of multiple data interfaces over the backplane, including CAN for C&C, SpaceWire for either C&C or data, and SpaceFibre for high-speed data transfer. The interfaces are provided over different backplane connectors: P1-P6, where P2-P5 are optional, depending on the use case.

ADHA defines a set of module profiles which specify the utilisation of the backplane connectors:

- Power Module Slot
- System Controller Slot

- System Controller Extension Slot
- Peripheral Module Slot
- Extended Peripheral Module Slot

The System Controller Slot is the central controller of the ADHA unit, providing the power on/off and reset control of all peripheral modules, the HMS (health monitoring system), the CAN bus master, and the SpaceWire/SpaceFibre routers (P3, P4 and P6 connectors). The System Controller Extension Slot offers to offload the power on/off reset control function from the System Controller, i.e. in the case of a platform OBC with a hot-redundant reconfiguration module.

The Peripheral Module Slot provides interfaces for an end-node for the CAN bus, SpaceWire and SpaceFibre networks (on the P4 connector).

The Extended Peripheral Module Slot provides the same interfaces as the Peripheral Module, but adds the option to implement a second set of SpaceWire and SpaceFibre routers (on the P3 and P6 connectors).

Through the routers provided in the System Controller and Extended Peripheral Module Slots, an ADHA unit offers both dual-star SpaceWire and SpaceFibre routed networks within an ADHA unit. [4]

On the level of Peripheral Modules, it is possible to implement up to four SpaceWire interfaces, and four dual-lane SpaceFibre interfaces, which are connected to the nominal and redundant System Controller and Extended Peripheral Module Slot.

Note that in most application cases, not all SpaceWire (SpW) and SpaceFibre (SpFi) interfacing options on the backplane are used. E.g. in a typical OBC+MM unit, the SpaceWire router may be implemented in the OBC for C&C, and the SpaceFibre router is implemented in the MM for payload data routing.

3. ADHA EXAMPLE MODULE CONFIGURATIONS

In this section, some example ADHA module configuration and use case examples that integrate SpaceFibre routers are introduced. The following use cases are presented:

- A System Controller Module for ICU (Instrument Controller Unit) / DPU (Data Processing Unit) applications
- A Mass-Memory Master (MMM) Module
- A Data Processing Module (DPM)

A block diagram of an example ADHA System Control Module configuration is shown in Fig. 1. In this example, both SpaceWire and SpaceFibre routers are implemented, providing routing both for payload data and C&C traffic. The target application case for the module is a System Controller in a Data Processing Unit (DPU), interconnecting multiple Peripheral Modules (normally DPM modules) through SpaceFibre. [5]

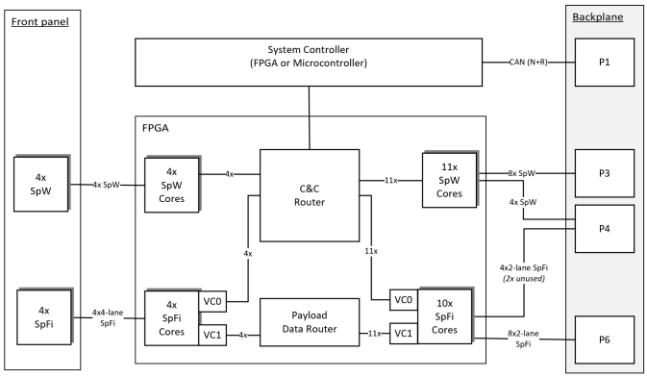

Fig. 1: ADHA System Controller Module example

In the example of Fig. 1, to ensure compatibility with the maximum configuration of Peripheral Modules in the same ADHA unit, and expected external interfaces, the SpaceFibre router provided on the module is required to have the following interfaces:

- Ten dual-lane interfaces to the backplane (total of 20 lanes) – with the possible addition of another dual-lane interface for cross-strapping between System Controller Modules, or Extended Peripheral Modules.
- Four quad-lane, on the front-panel (total of 16 lanes). These lanes are connected to four quad-lane fibre optic transceivers.
- Several SpaceWire interfaces may also be included in the SpaceFibre router or in a separate SpaceWire router.

The SpaceFibre lane rate is to be at least 6.25 Gbit/s per lane (note that the ADHA backplane connectors currently under pre-qualification are targeting to in the future support up to 25Gbit/s/lane [6]). A total of 10×2 + 4×4 = 36 lanes are required.

Another example is the Mass-Memory Master (MMM) Module, shown in Fig. 2. The MMM is responsible for providing external interfaces to instrument/payload units (such as e.g. a payload ICU or DPU), to the payload transmitter (link to ground), and to interconnect Mass-Memory Extension (MEM) modules on the ADHA backplane through SpaceFibre. The MMM manages the overall mass-memory function, while the MEM modules provide the non-volatile storage [7].

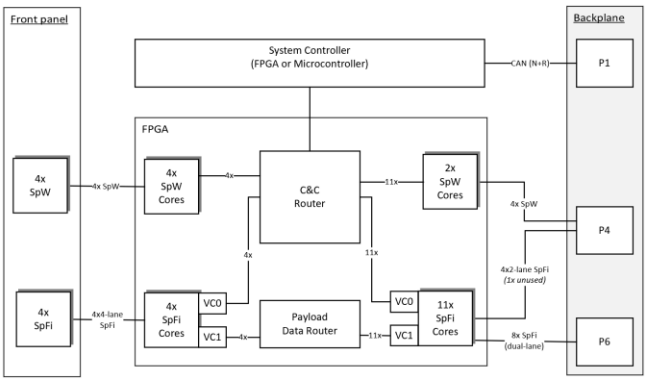

Fig. 2: ADHA Mass-Memory Master Module example.

MMM modules are usually included in either a CDHU (Central Data Handling Unit) together with an OBC, integrated as an Extended Peripheral Module, or in a free-standing mass-memory unit or PDHU (Payload Data Handling Unit), with the MMM placed in the System Controller Slot.

An additional example is the Data Processing Module (DPM, previously referred to as CPM: Co-Processor Module). An example DPM is shown in Fig. 3.

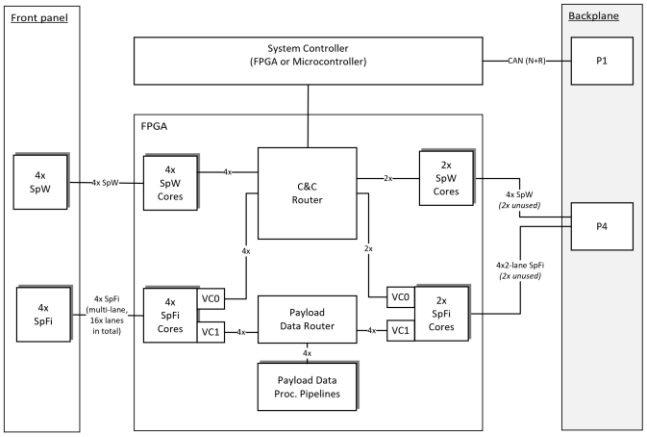

Fig. 3: ADHA Data Processing Module example

The DPM provides SpaceWire and SpaceFibre front-panel interfaces to connect to other equipment, for example, instrument FEEs (Front-End Electronics), or cameras for monitoring or Visual-Based Navigation (VBN) applications. The DPM is specified in the ADHA Data Processing Module Specification document (ADHA-RS-032).

Data in the DPM normally arrives on the front-panel interfaces, enters into the local data processing function, and after processing, is routed towards the backplane interfaces.

The DPM requirement specification recommends to implement at least 4-port, 4-lane (4P4L) SpaceFibre interfaces on the front-panel. Note that for the backplane interfaces, on the DPM the P6 connector is optional, and the P3 (SpaceWire router) is not used, so those connectors are not shown.

In summary, for the three example application cases listed above, the following number of interfaces are needed:

Table 1: Use case examples, summary of SpaceWire, and SpaceFibre interfacing requirements.

Module Type	**DPU System Control.**	**Mass-Memory Master (MMM)**	**Data Proc. Module (DPM)**
Backplane connectors			
P4 (SpW/SpFi end-point) implemented	Yes	Yes	Yes
P3 (SpW router) implemented	Yes	No	No
P6 (SpFi router) implemented	Yes	Yes	No *(optional)*
Backplane interfaces			
Backplane SpW	12	4	2 *+2 unused*
Backplane SpFi (2-lane)	10P2L *+2 unused*	11P2L *+1 unused*	2P2L *+2 unused*
Front-panel interfaces			
External SpW interfaces	4	4	4
External SpFi interfaces	4P4L	4P4L	4P4L
Internal ports			
Internal SpW ports	4	1	0
Internal SpFi ports	0	0	4P1L
Summary			
Number of SpW ports	20	9	6
Number of SpFi interfaces, and lane count	10P2L 4P4L	11P2L 4P4L	2P2L 4P4L 4P1L
Total number of SERDES/MGTs implemented in hardware	36	38	24
Number of VCs per SpFi interface	Min. 2	Min. 2	Min. 2

From the requirements, and for the purpose of demonstration, the following devices for the target application cases have been selected by ESA:

- ICU/DPU System Controller: PolarFire (option: NG-ULTRA)
- Mass-Memory Master (MMM) Module: Versal AI Core (VC1902)
- Data Processing Module (DPM): Kintex Ultrascale KU060

The capabilities of the selected devices, are outlined in the sections below.

In the case of the DPU System Controller, to fit with the number of available MGT (multi-gigabit transceivers) in the PolarFire, we reduce the total number of implemented lanes

to 24, by reducing the number of backplane SpaceFibre interfaces from 10P2L lane, to 4P2L – i.e. supporting connections to eight Peripheral Modules (in the target use case: to DPM modules).

Note that the number of interfaces per board depends heavily on the specific flight-application, and the configurations presented here may differ somewhat to specific flight configurations.

Based on the use cases listed above, we show below, in Fig. 4, a simplified diagram of a resulting example data-handling system configuration of a mass-memory unit, and a data processing unit (DPU).

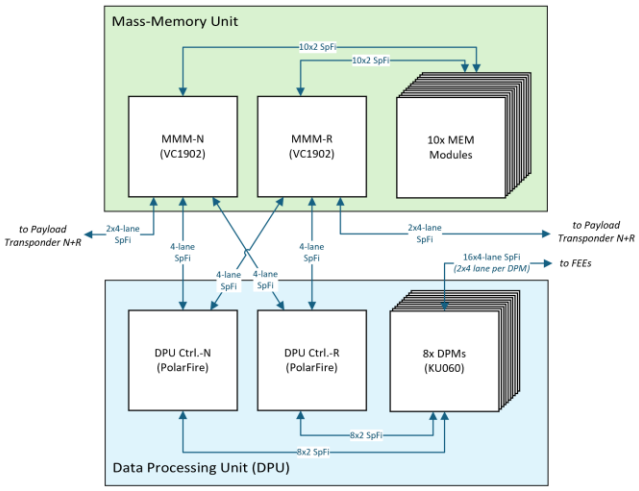

Fig. 4: Simplified Data-Handling System showing only SpaceFibre connections.

4. SPACEFIBRE

SpaceFibre is a data link and network technology developed specifically for spacecraft on-board data-handling. It runs over electrical or fibre-optic physical layers. It uses latent multi-gigabit transceiver (MGT) technology available in current chips (FPGAs and ASICs) to provide data rates of tens and, in the near future, hundreds of Gbit/s. It automatically recovers rapidly from transient errors without loss of data, which has the welcome side-effect of improving the radiation tolerance of the multi-gigabit transceivers. It uses multiple lanes to achieve high data-rates, but also provides graceful-degradation and hot and cold redundancy for those lanes, in the case of a lane failure. Virtual channels with priority, bandwidth reservation and scheduling are used to form virtual networks which isolate different flows of data from one another. SpaceFibre also provides a low-latency broadcast message capability which can be used to distribute system time, to trigger or indicate the occurrence of events, to give notification of errors, etc. SpaceFibre has been at TRL-9 since 2021 and is now flying on more than ten spacecraft and being designed into over 75 more [8].

5. WHY USE SPACEFIBRE IN ADHA?

SpaceFibre is an important part of ADHA. This section reflects on why SpaceFibre is being used.

High data rate: SpaceFibre is able to provide the high data rates needed by ADHA. It is generally able to operate at the lane rates that can be supported by the MGTs in the FPGA. The Microchip RTG4 FPGA is able to support lane rates of 3.125 Gbit/s. Other radiation tolerant FPGAs have MGTs that operate at 10 Gbit/s or higher and the AMD Versal incorporates MGTs that offer lane rates that are higher still. For electrical backplanes, because of the losses in the connectors and PCB tracks, a more modest lane rate has to be targeted. Low-loss PCB materials can help, but still the data rates of copper backplanes are limited to around 10 Gbit/s per lane. For higher data rates multiple lanes per link can be used. ADHA specifies dual-lane links giving 12.5 Gbit/s in each direction at a lane rate of 6.25 Gbit/s. Over fibre optic media, the data rate that can be achieved is potentially high, but is limited by the optical transceiver. Current radiation tolerant fibre optic transceivers operate at either up to10 Gbit/s or up to 25-28 Gbit/s per lane.

MGT availability: SpaceFibre relies on high-speed Serializers/De-serializers (SerDes), otherwise known as Multi-Gigabit Transceivers (MGTs), to transmit data over a physical link and to recover the data at the far end of the link. Many FPGAs now include MGTs making the inclusion of SpaceFibre in an FPGA straightforward. When external SerDes had to be used (such as the TLK2711, Wizard Link device) the power consumption per lane was substantial, of the order of 1 W. In recent FPGAs the MGTs have a power consumption in the region of 100 mW.

Compatibility with SpaceWire at the packet level: SpaceFibre is compatible with SpaceWire at the packet level, using the same packet format and addressing scheme. System software that sends commands and data over SpaceWire can operate over SpaceFibre. If an application running over SpaceWire needs higher data rates, lower latency or higher reliability, it is trivial to move that application to run over a virtual channel/network of SpaceFibre. Mixed SpaceWire and SpaceFibre networks can also be supported.

Transient error containment and recovery: SpaceFibre provides rapid error recovery in the case of a transient error. The fault is contained in the link where it occurred and the data on the line (i.e. data that has left the link layer in the near-end transmitter but not yet been received correctly by the link layer in the far-end receiver) at the time of the fault is retransmitted taking around 3 µs. The state of the link is retained and no error propagates outside the link, substantially simplifying error recovery. Radiation testing of the STAR-Dundee SpaceFibre IP cores on the RTG4, PolarFire and Versal have shown that SpaceFibre is able to recover radiation induced faults in the MGTs, substantially improving the radiation tolerance of the MGTs.

Graceful degradation on lane failure: If a lane fails in a multi-lane link, the faulty lane will be removed from the link and operation will continue with the remaining links. The bandwidth is reduced but the most important information, as determined by the priority of the virtual channel it is being sent over, will get through and the least important information will be held up. If control information is given high priority, it will always be able to get through in the event of a lane failure.

Broadcast messages: Broadcast messages are short, low-latency messages that are broadcast across a SpaceFibre network. They contain a broadcast channel (normally the logical address of the node sending the broadcast message), a broadcast type (that identifies the type of information it contains) and 8-bytes of user information. Broadcast messages can be used to distribute CCSDS unsegmented time format, GNSS synchronisation pulses, equipment triggers, network error messages, event notification, etc. The latency through a single routing switch is better than 1 µs.

SpaceFibre protocols run in hardware: All the capabilities of SpaceFibre run in hardware with software being used only for network configuration, overall control and status monitoring. This improves reliability, increases performance and removes the need for high performance IO processors to handle information transfer.

6. SpaceFibre Routing Switch

SpaceFibre routing switches are the key element in a SpaceFibre network, switching packets between the nodes connected to the routing switch. STAR-Dundee showed its first SpaceFibre router in 2017 [9], followed in 2018 with demonstration in the RTG4 FPGA [10]. In 2022, the first multi-lane router was demonstrated in the Hi-SIDE project [11] and tested to TRL-5. The SpaceFibre Router IP has been ported to the Microchip RTG4, PolarFire and PolarFire SoC, and to the AMD KU060 and Versal [12][8].

The architecture of the STAR-Dundee SpaceFibre Router IP core is illustrated in Fig. 5.

At the heart of the SpaceFibre routing switch is a non-blocking routing switch. To this are connected:

- A configurable number of SpaceFibre ports (port 1 to port P_1), each of which has a configurable number of lanes and a configurable number of virtual channels.
- A configurable number of AXI4-Stream ports (port P_1+1 to P_2), each having a configurable number of virtual channels. There is one AXI4-Stream interface for each virtual channel.
- A configurable number of SpaceWire ports (P_2+1 to P_3), each of which has a single virtual channel and a packet buffer for helping to reduce the packet blocking on the virtual network to which that SpaceWire port is connected.
- A configuration port (port 0) which provides RMAP access to the configuration, control and status registers and routing table via a single virtual channel. These registers can also be accessed locally via an optional AXI4-Lite interface.

The virtual channels on each input can be mapped to virtual networks in a routing switch, with up to 64 virtual networks available. This is very helpful for minimising the number of virtual channels required in an instrument, for example.

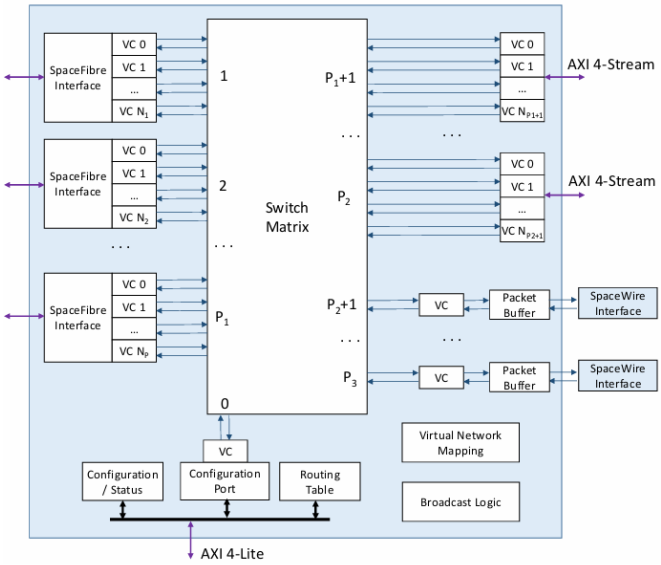

Fig. 5: SpaceFibre Routing Switch IP Core Architecture

7. FPGAs for ADHA SpaceFibre Router

The FPGA that the SpaceFibre routing switch is to be implemented in needs to have MGTs that operate at 10 Gbit/s or more, and to have up to 28 lanes available. Several radiation tolerant FPGAs from Microchip, AMD and NanoXplore are listed in Table 2. There are a few possible FPGAs that could be used for the most demanding ADHA use cases: the KU060 with 32 lanes, and the Versal 1902, with 44 lanes, or the NG-ULTRA with 32 lanes. The STAR-Dundee SpaceFibre Router IP core has been ported to a wide range of FPGAs which are suitable for space applications, including KU060 and Versal. In Table 2, green is good, yellow is possibly acceptable, and red is not acceptable. At present it is not possible to use the MGTs effectively in the NG-ULTRA.

The PolarFire or RTG4 could be used with up to 24 lanes, e.g. reducing some backplane interfaces to single lane to allow more front-panel interfaces, or to use the full backplane interfaces and provide 2-4 lanes on the front-panel.

Table 2: Radiation Tolerant FPGAs and their MGTs

Manufacturer / FPGA	Lanes	Speed (Gbit/s)	Size (mm)
Microchip			
RTG4	24	3.125	42.5 x 42.5
PolarFire RTPF500ZT	24	10.3	40 x 40
PolarFire SoC RTPFS160ZT	8	12.7	23 x 23
PolarFire SoC RTPFS460ZT	20	12.7	40 x 40
AMD			
AMD Kintex XQRKU060	32	12.5	40 x 40
AMD Versal XQRVC1902	44	26.5	45 x 45
AMD Versal XQRVE2302	8	26.5	23 x 23
NanoXplore			
NG-ULTRA NX2H540TSC	32	12.5	45 x 45

The KU060 FPGA was used in the STAR-Dundee STAR-Tiger SpaceFibre routing switch (shown in Fig. 6). This router has eight dual-lane ports and two quad-lane ports (a total of 20 lanes). It operates with lane rates of 6.25 Gbit/s. It was developed and tested to TRL-5. The power consumption is 14.5 W typical at 20 °C, with all links running with lane speeds of 6.25 Gbit/s. The bisection bandwidth of the STAR-Tiger is 60 Gbit/s.

Fig. 6: STAR-Tiger SpaceFibre Routing Switch

The Versal 1902 FPGA may be more appropriate for the most demanding AHDA application use cases especially when more than 28 SpaceFibre lanes are required.

8. Implementation and Performance

Various SpaceFibre routing switch configurations representative of what is required for ADHA were evaluated on the PolarFire, KU060 and Versal 1902 FPGAs. All the designs were tested in hardware using the appropriate development boards. In each case the router configuration is detailed with #P being the number of SpFi ports, #L the number of lanes per port, #VC the number of virtual channels (VCs) per port and #SpW the number of SpaceWire ports. For example, 12P1L4VC + 1SpW has 12 single-lane SpaceFibre ports each with 4 VCs plus 1 SpaceWire port. For each different FPGA and router configuration the specific FPGA used, the router configuration, the achieved lane rate, the resource utilisation, the power estimate and in some cases a placement diagram, are provided. All the designs are without TMR.

The results for the Microchip PolarFire are presented in Table 3, Table 4 and Table 5, for the AMD KU060 in Table 6 and Table 7, and for the AMD Versal in Table 8 and Table 9.

The bisection bandwidth gives an idea of the relative performance of different routing switch designs. The bisection bandwidth is given by

$$P/2 \times L \times 0.8R$$

where

P is the number of ports,

L is the number of lanes per port, and

0.8R is the lane rate less the 8B10B overhead.

Table 3: PolarFire with 12P1L4VC + 1SpW Router

Part used: Microchip PolarFire RTPF500ZT-1CG1509M

Router Configuration: 12P1L4VC + 1SpW

Maximum Lane Rate Achieved: 6.5 Gbit/s

Bisection Bandwidth: 31.2 Gbit/s

Resource Usage: 35.6%

Type	Used	Total	Percentage
4LUT	171117	481272	35.56
DFF	132745	481272	27.58
I/O Register	0	1896	0.00
User I/O	9	632	1.42
-- Single-ended I/O	9	632	1.42
-- Differential I/O Pairs	0	316	0.00
uSRAM	86	4440	1.94
LSRAM	358	1520	23.55

Power Summary: 2.1 W

	Power (mW)	Percentage
Total Power	2088.775	100.0%
Static Power	160.544	7.7%
Dynamic Power	1928.231	92.3%

Placement:

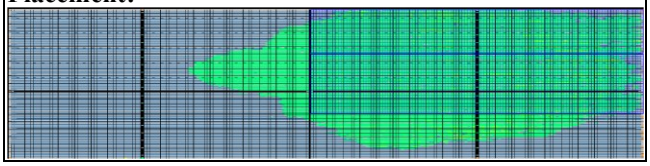

Table 4: PolarFire with 10P1L4VC + 14P1L2VC + 4SpW Router

Part used: Microchip PolarFire RTPF500ZT-1CG1509M

Router Configuration: 10P1L4VC + 14P1L2VC + 4SpW

Maximum Lane Rate Achieved: 4.8 Gbit/s

Bisection Bandwidth: 46.1 Gbit/s

Resource Usage: 75.5%

Type	Used	Total	Percentage
4LUT	363360	481272	75.50
DFF	258806	481272	53.78
I/O Register	0	1896	0.00
User I/O	9	632	1.42
-- Single-ended I/O	9	632	1.42
-- Differential I/O Pairs	0	316	0.00
uSRAM	178	4440	4.01
LSRAM	570	1520	37.50

Power Summary: 4.3 W

	Power (mW)	Percentage
Total Power	4317.173	100.0%
Static Power	162.803	3.8%
Dynamic Power	4154.370	96.2%

The lane rate in Table 4 was achieved without any optimisation. It is expected that this can be improved. At a lane rate of 5.0 Gbit/s, with a dual-lane link the user data rate will be around 7.5 Gbit/s in both directions.

Table 5: PolarFire with 12P2L4VC + 4SpW Router

Part used: Microchip PolarFire RTPF500ZT-1CG1509M
Router Configuration: 12P2L4VC + 4SpW
Maximum Lane Rate Achieved: 5 Gbit/s
Bisection Bandwidth: 48 Gbit/s
Resource Usage: 58.7%

Type	Used	Total	Percentage
4LUT	282711	481272	58.74
DFF	207225	481272	43.06
I/O Register	0	1896	0.00
User I/O	9	632	1.42
-- Single-ended I/O	9	632	1.42
-- Differential I/O Pairs	0	316	0.00
uSRAM	122	4440	2.75
LSRAM	551	1520	36.25

Power Summary: 3.6 W

	Power (mW)	Percentage
Total Power	3628.687	100.0%
Static Power	162.600	4.5%
Dynamic Power	3466.087	95.5%

Table 6: KU060 with 24P1L4VC + 4SpW Router

Part used: AMD Kintex Ultrascale XQRKU060-CNA1509-1M-m
Router Configuration: 24P1L4VC + 4SpW
Maximum Lane Rate Achieved: 6.25 Gbit/s
Bisection Bandwidth: 60 Gbit/s
Resource Usage: 78.2%

Resource	Utilization	Available	Utilization ...
LUT	259208	331680	78.15
LUTRAM	336	146880	0.23
FF	270852	663360	40.83
BRAM	390	1080	36.11
IO	3	620	0.48
GT	24	32	75.00

Power Summary: 9.0 W

	Power (mW)	Percentage
Total Power	25582	100.0%
Static Power	14452	56.5%
Dynamic Power	11130	43.5%

Placement:

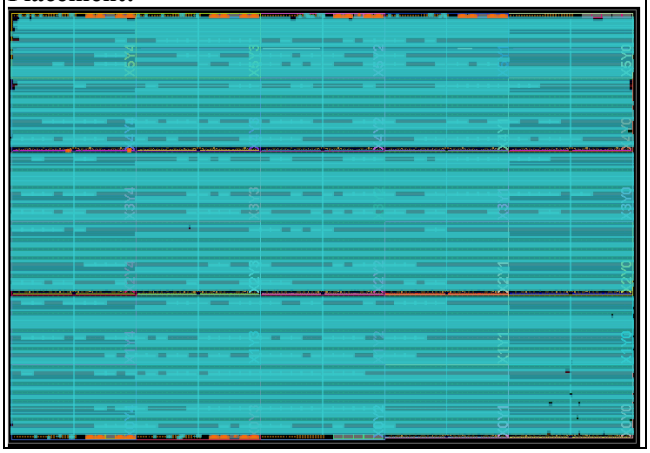

Table 7: KU060 with 16P2L4VC

Part used: AMD Kintex Ultrascale XQRKU060-CNA1509-1M-m
Router Configuration: 16P2L4VC
Maximum Lane Rate Achieved: 6.25 Gbit/s
Bisection Bandwidth: 80 Gbit/s
Resource Usage: 59.1%

Resource	Utilization	Available	Utilization...
LUT	195855	331680	59.05
LUTRAM	801	146880	0.55
FF	222741	663360	33.58
BRAM	403	1080	37.31
IO	3	620	0.48
GT	32	32	100.00

Power Summary: 10.7 W

	Power (mW)	Percentage
Total Power	10684	100.0%
Static Power	1443	13.5%
Dynamic Power	9241	86.5%

Table 8: Versal 1902 with 28P1L4VC + 3SpW Router

Part used: AMD Versal XCVC1902-VSVA2197-2MP-e-S
Router Configuration: 28P1L4VC + 3SpW
Maximum Lane Rate Achieved: 10 Gbit/s
Bisection Bandwidth: 112 Gbit/s
Resource Usage: 46%%

Utilization | Post-Synthesis | Post-Implementation

Graph | Table

Resource	Utilization	Available	Utilization...
LUT	282398	899840	31.38
LUTRAM	392	449920	0.09
FF	343090	1799680	19.06
BRAM	443	967	45.81

Power Summary: 26.6 W

	Power (mW)	Percentage
Total Power	25582	100.0%
Static Power	14452	56.5%
Dynamic Power	11130	43.5%

Placement:

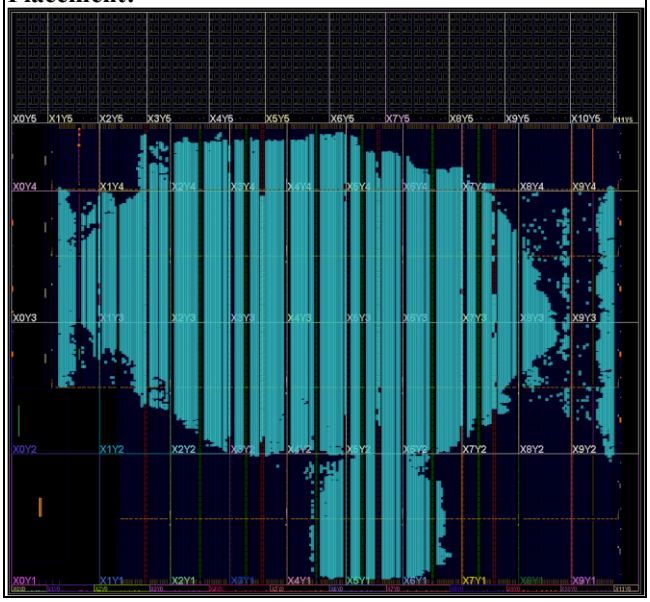

Table 9: Versal 1902 with 19P2L4VC + 4SpW Router

Part used: AMD Versal XCVC1902-VSVA2197-2MP-e-S

Router Configuration: 19P2L4VC + 4SpW

Maximum Lane Rate Achieved: 10 Gbit/s

Bisection Bandwidth: 152 Gbit/s

Resource Usage: 54.8%%

Resource	Utilization	Available	Utilization %
LUT	253143	899840	28.13
LUTRAM	950	449920	0.21
FF	304381	1799680	16.91
BRAM	530	967	54.81

Power Summary: 27.9 W

	Power (mW)	Percentage
Total Power	27947	100.0%
Static Power	14505	51.9%
Dynamic Power	13442	48.1%

Placement:

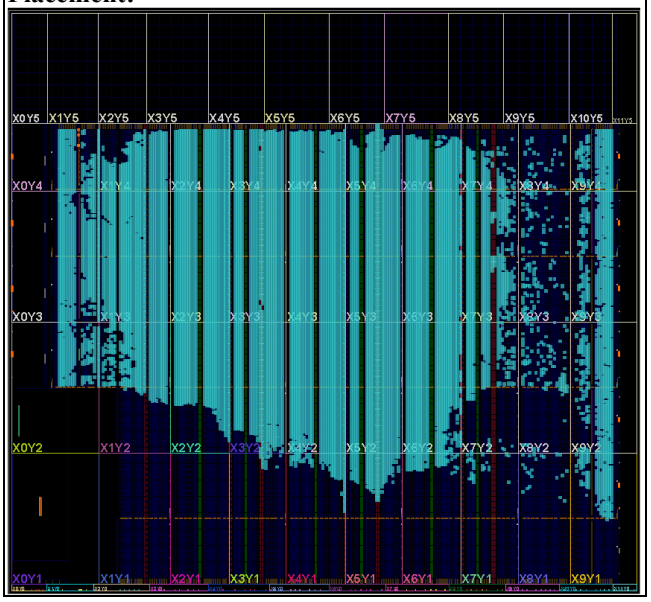

9. CONCLUSIONS

ADHA requires SpaceFibre routing switches for the System Controller, Mass-Memory Master and Data Processing modules. Various levels of performance are required depending on the application. Several routing switch architectures have been tested on the PolarFire, KU060 and Versal 1902 FPGAs. The achieved lane rates, resource utilisation and power estimates have been presented, along with the bisection bandwidth of each router implementation for use as a quick comparison of relative router performance.

A comparison of the performance of the various routing switches is presented in Fig. 7. The Versal 1902 provides superior bisectional bandwidth, whereas the PolarFire offers substantially better performance per Watt. The KU060 is in between.

SpaceFibre router technology is available to support a wide range of ADHA based systems and is currently being integrated into ADHA boards.

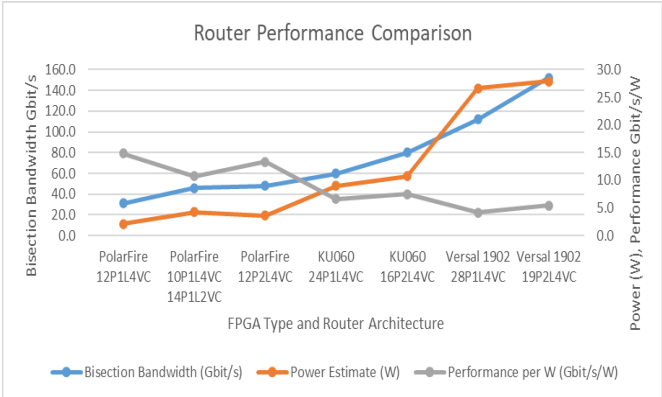

Fig. 7: Router Performance Comparison

10. REFERENCES

[1] ESA, "Advanced Data Handling Architecture (ADHA)", ADHA-2 Workshop, ESTEC, November 2022.

[2] ECSS Standard ECSS-E-ST-50-11C, "SpaceFibre – Very High-Speed Serial Link", Issue 1, European Cooperation for Space Data Standardization, May 2019, available from http://www.ecss.nl.

[3] Olivier Mourra, Kostas Marinis, Felix Siegle, Hadrien Carbonnier, David Steenari, Laurent Hili, Dario Pascucci, Julian Bozler, and Jan Johansson. "Advanced Data Handling Architecture (ADHA): System Architecture and Design Description." In *2023 European Data Handling & Data Processing Conference (EDHPC)*, pp. 1-8. IEEE, 2023.

[4] Franz, Robin, Julian Bozler, Michael Stähle, and Felix Siegle. "A Space Wire/SpaceFibre architecture based on ADHA." In *2022 International SpaceWire & SpaceFibre Conference (ISC)*, pp. 01-05. IEEE, 2022.

[5] Olivier Mourra, Felix Siegle, David Steenari, Kostas Marinis, Laurent Hili, Dario Pascucci, and Julian Bozler. "ADHA, an Agile Platform Enhancing New Satellite On-Board Data Processing Systems." In *2023 European Data Handling & Data Processing Conference (EDHPC)*, pp. 1-9. IEEE, 2023.

[6] Léo Farhat, Joaquin Jimenez, Adria Escoda, and Olivier Mourra. "Reliability Assessment of High Data Rate cPCI Serial Space Connectors." In *2023 European Data Handling & Data Processing Conference (EDHPC)*, pp. 1-3. IEEE, 2023.

[7] Christian Spindeldreier, "Scalable High-Speed Mass Memory Units for ADHA and beyond", 18th ESA Workshop on Avionics, Data, Control and Software Systems (ADCSS2024), ESA ESTEC, Oct 22-24, 2024,

[8] Steve Parkes, Albert Ferrer, Alberto Gonzalez, David Gibson, "SpaceFibre Onboard Interconnect: from Standard, through Demonstration, to Space Flight", IEEE Aerospace, Big Sky, Montana, USA, Mar 1-8, 2025.

[9] Steve Parkes, Albert Ferrer Alberto Villafranca, Chris McClements and David McLaren, "SpaceFibre Network and Routing Switch", IEEE Aerospace, Big Sky, Montana, USA, Mar 4-11, 2017.

[10] Steve Parkes, Ashish Shrivastava, Chris McClements, Pete Scott, David Dillon, Albert Ferrer and Alberto Gonzalez, "SpaceFibre Camera", 8th International SpaceWire Conference, Los Angeles, California, USA, May 14-18, 2018.

[11] Hi-SIDE Consortium, https://www.HiSIDE.space/.

[12] Steve Parkes, Albert Ferrer, Alberto Gonzalez, Marti Farras, Blair Winton, Pete Scott, Keith Carruthers and David Gibson, "A SpaceFibre Routing Switch", DASIA 2023, Sitges, Spain, June 2023.