



SpWRT

SpaceWire-RT

EU-Russia Cooperation for
Strengthening Space Foundations

Final Publishable Summary Report

June 2013



1 EXECUTIVE SUMMARY

The SpaceWire-RT project has created a spacecraft on-board networking technology that is able to fulfil many on-board communication requirements including payload data-handling and avionic applications. In addition the project has substantially strengthened collaborative links between Russian and European academia and aerospace industry, building on the strengths and capabilities of the project partners.

The requirements for spacecraft on-board communications were researched by the industrial partners covering both Russian and European requirements and a comprehensive set of use cases were defined. The SpaceFibre network technology being developed by the European Space Agency was identified as a good baseline for SpaceWire-RT. Enhancements to SpaceFibre necessary to meet the SpaceWire-RT requirements were identified and research undertaken in the areas of quality of service (QoS), fault detection, isolation and recovery (FDIR) and network layer concepts. The design for a simple, powerful, and comprehensive quality of service mechanism resulted. This QoS mechanism was then extended to include specific classes of fault detection, “babbling idiot” and quiet unit detection, in support of FDIR. Specifications for the SpaceFibre Quality layer, covering QoS and FDIR, and the Network layer were written.

The entire SpaceFibre protocol stack was validated through simulated using SDL and SystemC. The results of these detailed simulations were used to help remove errors and inconsistencies in the specification and to make necessary clarifications and improvements to the SpaceFibre standard. Several problems with the SpaceFibre specification were identified and potential solutions were explored and traded-off. The specification was updated accordingly.

A VHDL implementation of the SpaceFibre interface was designed and implemented in a field programmable gate array (FPGA). This provided a means of further testing and validating the SpaceFibre specification including its QoS and FDIR features. Results of this implementation fed back into the specification and further inconsistencies and problems were resolved, with the experimental FPGA implementation being updated and re-tested. The feasibility of ASIC implementation of the SpaceFibre was explored and a complete SpaceFibre interface was shown to be only three or four times the chip area of a SpaceWire interface.

The SpaceFibre specification was updated several times throughout the SpaceWire-RT project and a Russian translation made. It is planned for this specification to be taken through formal standardisation with the European Cooperation for Space Standardization (ECSS) once other aspects of SpaceFibre have been completed, including multi-lane operation. Research and development related to SpaceWire-RT is being carried forward by the project partners with the aim of spaceflight ready components and system within the next two to three years.

Table of Contents

1	EXECUTIVE SUMMARY	2
2	Project Context and Objectives.....	4
2.1	Context	4
2.2	Project Objectives.....	6
2.3	Project Team	7
3	Main Scientific and Technological Results	8
3.1	Project Overview.....	8
3.2	WP1 Spacecraft Avionics and Payload Use Cases	9
3.3	WP2 Concepts and Specification Development.....	10
3.4	WP3 SpaceWire-RT Validation and Simulation	22
3.5	WP4 SpaceWire-RT VHDL IP Core Development	24
3.6	SpaceWire-RT ASIC Feasibility and Virtual Prototyping	28
3.7	SpaceWire-RT Standard Draft.....	29
3.8	SpaceWire-RT Objectives Addressed	31
4	Potential Impact, Main Dissemination Activities and Exploitation of Results	32
4.1	Potential impact	32
4.2	Main dissemination activities	33
4.3	Exploitation of Results	34
4.4	Education and Training.....	35
5	Project Details	36
6	References.....	37

2 Project Context and Objectives


2.1 Context

2.1.1 The Need for Standard Network Technology for Spacecraft Applications


The trend towards “Operationally Responsive Space”, where spacecraft can be rapidly assembled, configured and deployed, to meet specific mission needs, e.g. disaster support, requires flexible, high-performance, on board communication networks with plug-and-play capability. Earth observation missions employing synthetic aperture radar or hyperspectral imaging have very high data rates requiring multi-gigabit/s onboard network technology. The growing autonomy of scientific missions to remote planets requires networks that are robust and durable, able to recover from transitory errors and faults automatically. The importance of spacecraft mass reduction motivates the sharing of networks for payload data-handling and avionics. Avionics and robotics applications impose requirements on network responsiveness and determinism. Increasing international collaboration on scientific and Earth observation spacecraft requires standard network technology where a component developed by one nation will interoperate effectively with equipment developed by another. SpaceWire-RT fulfils these demanding requirements with a flexible, robust, responsive, deterministic and durable standard network technology that is able to support both avionics and payload data-handling applications.

2.1.2 SpaceWire

The existing SpaceWire technology [1],[2],[3] was a very successful first step in this direction, providing networking technology for payload data-handling on over 40 major space missions, including the missions listed in Table 1. It falls short, however, of the deterministic and galvanic isolation requirements for avionics systems. SpaceWire is also of insufficient data-rate for some space applications.

Table 1 Example SpaceWire Missions**James Webb Space Telescope**


A large infrared telescope being built by NASA with a 6.5 m diameter mirror.

**Astro-H**


A Japanese X-ray telescope due for launch in 2015, which has the largest SpaceWire network of any spacecraft design to date.

**GAIA**

ESA mission to map the stars to very high resolution due to be launched in December 2013.

**ExoMars Rover**

European Mars rover using SpaceWire extensively.

**Bepicolombo**

Joint ESA, JAXA mission to Mercury comprising a Polar Orbiter and Magnetospheric Orbiter.

**Lunar Reconnaissance Orbiter**

A NASA mission currently in orbit around the Moon.

**Sentinel 1**

An Earth observation spacecraft using radar imaging being built by ESA.

**Sentinel 2**

ESA oceanographic mission.

2.1.3 SpaceFibre

SpaceFibre [4],[5],[6],[7] is a very high-speed serial data-link currently being developed by ESA which is intended for use in data-handling networks for high data-rate payloads. SpaceFibre is able to operate over fibre optic and electrical cable and support data rates of 2 Gbit/s in the near future and up to 20 Gbit/s long-term. It aims to complement the capabilities of the widely used SpaceWire onboard networking standard: improving the data rate by a factor of 10, reducing the cable mass by a factor of four and providing galvanic isolation. SpaceFibre aims to support Quality of Service (QoS) along with Fault Detection, Isolation and Recovery (FDIR). An important feature of SpaceFibre is that it has the same packet format as SpaceWire, so that several SpaceWire links can be easily multiplexed over a single SpaceFibre link, and application software designed for SpaceWire can operate over a SpaceFibre network. SpaceFibre is currently missing the important QoS and FDIR capabilities and the network layer of SpaceFibre has yet to be defined.

2.2 Project Objectives

The SpaceWire-RT research programme aims to conceive and create communications network technology, suitable for a wide range of demanding space applications where responsiveness, determinism, robustness and durability are fundamental requirements. This is a critical component technology for future spacecraft avionics and payload data-handling. The SpaceWire-RT project substantially strengthened collaborative bonds between the Russian and European organisations involved in the research, and led to technology of vital importance for future space missions.

The SpaceWire-RT research programme had the following key objectives:

1. A justified set of requirements and use cases for SpaceWire-RT covering both spacecraft payload and avionics, which takes into account requirements from European and Russian spacecraft primes and equipment manufacturers.
2. A conceptual design and outline specification for SpaceWire-RT that fulfils the requirements and use cases.
3. Validation of the SpaceWire-RT specification, with important, novel features of SpaceWire-RT networks tested using SDL and System-C models.
4. A SpaceWire-RT VHDL IP Core tested and validated in an FPGA implementation.
5. An assessment of the feasibility of implementing SpaceWire-RT as an ASIC core considering available ASIC technologies suitable for space.
6. A draft standard document for SpaceWire-RT in both English and Russian.
7. Dissemination of the results of the SpaceWire-RT study to the European and Russian space industries, and to the international space community.
8. An exploitation plan covering the availability of flight grade chips, test and development equipment, IP cores, and a programme for a test flight.

2.3 Project Team

A highly experience team of Russian and European academic and industrial organisations was assembled for the SpaceWire-RT project. The project team is introduced below:



University of Dundee (UNIVDUN) has long experience in spacecraft onboard network technology, writing the SpaceWire standard with support from the European Space Agency (ESA) and input from engineers across Europe. UNIVDUN has extensive experience in related IP core design having designed SpaceWire interface, remote memory access protocol (RMAP) and router cores for ESA, JAXA and other organisations.

The principal responsibilities of UNIVDUN were overall project management, conceptual design and specification, VHDL IP core design, and SpaceWire-RT standard specification.



Astrium GmbH is a major international spacecraft manufacturer that undertakes system integration and also designs and manufactures spacecraft payloads.

Astrium was responsible with SubMicron for the requirements and use cases for SpaceWire-RT based on their extensive spacecraft design experience.



SubMicron is a Russian organisation that specialises in the design and production of spacecraft avionics.

SubMicron was responsible with Astrium for the requirements and use cases for SpaceWire-RT based on their extensive experience of spaceflight electronic equipment.



St Petersburg University of Aerospace Instrumentation (SUAI) are experienced in the simulation and testing of network technologies using SDL particularly in the field of aerospace. SUAI designed their own SpaceWire IP cores and devised an innovative interrupt mechanism for SpaceWire.

SUAI was primarily responsible for the simulation and validation of the SpaceWire-RT protocols and for translating the SpaceWire-RT standard specification in to Russian. SUAI also contributed to the evolution of the SpaceWire-RT specification.



ELVEES is a leading Russian ASIC design house with expertise in space grade electronics and chip design.

ELVEES was responsible for reviewing the SpaceWire-RT specification and considering the feasibility of ASIC implementation.

3 Main Scientific and Technological Results

3.1 Project Overview

An overview of the project is illustrated in Figure 1.




Figure 1 Overall strategy of the work plan

The technical work began with WP1 and WP2. WP1 (Spacecraft Avionics & Payload Use Cases) focused on the requirements and use cases for spacecraft onboard communication networks for both payload applications and avionics. Requirements for avionics networks were gathered, covering reliability, fault tolerance, fault isolation, performance, responsiveness and determinism. A series of use cases was explored covering a diverse set of avionics and payload applications. These requirements formed the basis for the technical capabilities of the SpaceWire-RT network technology.

WP2 (Concept and Specification) reviewed existing networking technology, including commercial technologies, SpaceWire, SpaceFibre, and other onboard technologies. From the requirements of WP1 a set of evaluation criteria were derived and the candidate networking concepts traded-off against one another. SpaceFibre was selected as the baseline technology for SpaceWire-RT. Several areas of SpaceFibre which required substantial further research and development were identified for investigation in the SpaceWire-RT project, including QoS, FDIR and the network layer. Concepts were devised that met the requirements from WP1 and a specification produced.

In WP3 (Simulation and Validation) simulation models covering key aspects of SpaceWire-RT were designed and used to evaluate the proposed SpaceWire-RT solution. The simulation models were developed in SDL and used to run various validation scenarios. The results of the simulation were used to update the SpaceWire-RT specification and to inform the VHDL IP Core Development (WP4) and ASIC Feasibility (WP5) activities.

WP4 (VHDL IP Core Prototype) began with the architectural design of the SpaceWire-RT IP Core based on information from the outline SpaceWire-RT specification (WP2). Once the architectural design was ready the design and development of the VHDL IP core and associated test bench was

carried out, followed by testing. The IP core was implemented in an FPGA and tested extensively. The results were presented to the SpaceWire Working Group and the SpaceWire-RT specification updated.

WP5 (ASIC Feasibility) covered the review and evaluation of suitable ASIC technologies for SpaceWire-RT implementation, the procurement of suitable ASIC libraries, and initial design and simulation work. Owing to the expense of ASICs, device manufacture was beyond the scope of the present activity, however, the principal risk areas with an ASIC development were investigated and the feasibility of ASIC implementation assessed.

The results from all the previous work packages were fed into WP6 (Standard Draft) where a draft standard for SpaceWire-RT was written by UNIVDUN with all partners involved in review of the document. The SpaceWire-RT standard built on the SpaceFibre specification, adding in the quality and network layers. The draft standard was presented to the SpaceWire Working Group and feedback used to improve the document. SUAI translated the standard into Russian.

3.2 WP1 Spacecraft Avionics and Payload Use Cases

Astrium GmbH and Submicron gathered requirements from spacecraft manufacturers and spacecraft equipment suppliers across Europe and Russia respectively. These requirements covered many aspects of payload data-handling and avionics networks. A qualitative summary of the major communication requirements is provided in Table 2.

Table 2 Qualitative Communication Requirements					
	Distance	Rate	Latency	Packet size	QoS
Data-handling network	Short to long	Low to very high	Not important	Short to long	Reserved bandwidth
Control bus	Short to long	Low	Low	Short to long	Deterministic delivery
Telemetry bus	Short to long	Low	Low	Short	Reserved bandwidth
Computer bus	Short	Very high	Low	Short to long	Reserved bandwidth
Time-sync bus	Short to long	Low	Very low	Short	High priority
Side-band	Short	Low to high	Very low	Short	High priority

An important aspect of this work on requirements was the need for a range of QoS depending on the particular application. These include the capability to reserve network bandwidth to prevent different

data flows interfering with one another, deterministic data delivery for control applications, and high priority, very low-latency communication for time-synchronisation and side-band signalling.

The requirements were analysed by University of Dundee and compared against the characteristics of the planned SpaceFibre standard. It was clear that SpaceFibre would meet the majority of the requirements for SpaceWire-RT and so was adopted as the baseline network technology for very high data-rate applications. The outcome of this analysis was used to inform the work on QoS within SpaceFibre and resulted in a coherent precedence concept being proposed for QoS in SpaceFibre.

3.3 WP2 Concepts and Specification Development

3.3.1 SpaceWire-RT and SpaceFibre

The SpaceWire-RT project built on the emerging SpaceFibre standard which is being designed by the University of Dundee for ESA with inputs from international spacecraft engineers. The aim is to publish the final standard through the European Cooperation for Space Standardization (ECSS). The SpaceWire-RT project has contributed to the following parts of the SpaceFibre standard:

- QoS Mechanisms in the Quality layer,
- FDIR Mechanisms in the Quality layer,
- Network level concepts in the Network layer.

The QoS and FDIR work has been fully integrated in the SpaceFibre standard specification. A baseline network layer specification has been added to the SpaceFibre standard specification. The overall SpaceFibre protocol stack has been rationalised and the SpaceFibre standard restructured accordingly.

Several versions of the SpaceFibre specification were written as the SpaceWire-RT project progressed, incorporating QoS, FDIR and network layer specifications. Extensive simulation of each of these specification versions was carried out in the SpaceWire-RT project to help improve and validate the SpaceFibre draft standard. Draft E1 of the standard [7] was used for the critical simulation of the Virtual Channel and Retry Layers. Draft D [6] was used for the lower layers because this simulation work was carried out before Draft E1 of the SpaceFibre standard was written. The results of the simulations were used to improve the SpaceFibre specification.

3.3.2 SpaceFibre Protocol Stack

The SpaceFibre protocol stack incorporating the SpaceWire-RT results is illustrated in Figure 2.




Figure 2 SpaceFibre Protocol Stack

The **Network layer** is responsible for the transfer of application information over a SpaceFibre network. It provides two services: Packet Transfer Service and Broadcast Message Service. The Packet Transfer Service transfers SpaceFibre packets over the SpaceFibre network, using the same packet format and routing concepts as SpaceWire uses. SpaceFibre supports both path and logical addressing. The broadcast message service is responsible for broadcasting short messages (8 bytes) to all nodes on the network. These messages can carry time and synchronisation signals and be used to signal the occurrence of various events on the network.

The **Management layer** is responsible for configuring, controlling and monitoring the status of all the layers in the SpaceFibre protocol stack. For example it can configure the QoS settings of the virtual channels in the QoS and FDIR layer.

The **Quality layer** is responsible for providing quality of service and managing the flow of information over a SpaceFibre link. It frames the information to be sent over the link to support QoS and scrambles the packet data to reduce electromagnetic emissions. The QoS and FDIR layer also provides a retry capability, detecting any frames or control codes that go missing or arrive containing errors and resending them. With this inbuilt retry mechanism SpaceFibre is very resilient to transient errors.

The **Multi-Lane layer** is responsible for operating several SpaceFibre lanes in parallel to provide higher data throughput. In the event of a lane failing the Multi-Lane layer provides support for graceful degradation, automatically spreading the traffic over the remaining working links.

The **Lane layer** is responsible for lane initialisation and error detection. In the event of an error the lane is automatically re-initialised. The Lane layer encodes data into symbols for transmission using 8B/10B encoding and decodes these symbols in the receiver. 8B/10B codes are DC balanced supporting AC coupling of SpaceFibre interfaces.

The **Physical layer** is responsible for serialising the 8B/10B symbols and for sending them over the physical medium. In the receiver the Physical layer recovers the clock and data from the serial bit stream, determines the symbol boundaries and recovers the 8B/10B symbols. Both electrical cables and fibre-optic cables are supported by SpaceFibre.

The specific results of the SpaceWire-RT project that contributed to the SpaceFibre standard specification will now be described.

3.3.3 SpaceFibre Quality of Service

The specification, validation, prototype implementation and testing of the SpaceFibre Quality layer was a major output of the SpaceWire-RT project. In this section a detailed overview of this technology is provided.

3.3.3.1 Frames and Virtual Channels

To provide quality of service, it is necessary to be able to interleave different data flows over a data link or network. If a large packet is being sent with low priority and a higher priority one requests to be sent, it must be possible to suspend sending the low priority one and start sending the higher priority packet. To facilitate this SpaceWire packets are chopped up into smaller data units called frames. When the high priority packet requests to be sent, the current frame of the low priority packet is allowed to complete transmission, and then the frames of the high priority packet are sent. When all the frames of the high priority packet have been sent, the remaining frames of the low priority packet can be sent.

Each frame has to be identified as belonging to a particular data flow so that the stream of packets can be reconstructed at the other end of the link. Low priority packets belong to one data stream and high priority packets belong to another data stream.

Each independent data stream allowed to flow over a data link is referred to as a virtual channel (VC). Virtual channels are unidirectional and have a QoS attribute, e.g. priority. At each end of a virtual channel is a virtual channel buffer (VCB), which buffers the data from and to the application. An output VCB takes data from the application and buffers it prior to sending it across the data link. An input VCB receives data from the data link and buffers it prior to passing it to the receiving application.

There can be several output virtual channels connected to a single data link, which compete for sending information over the link. A medium access controller determines which output virtual channel is allowed to send the next data frame. When an output VCB has a frame of data ready to send and the corresponding input VCB at the other end of the link has room for a full data frame, the output VCB requests the medium access controller to send a frame. The medium access controller arbitrates between all the output VCBs requesting to send a frame. It uses the QoS attribute of each of the requesting VCBs to determine which one will be allowed to send the next data frame.

Priority is one example of a QoS attribute. Other types of QoS are considered in the subsequent sections.

3.3.3.2 Precedence

For the medium access controller to be able to compare QoS attributes from different output VCBs, it is essential that they are all using a common measure that can be compared. The name given to this measure is precedence. The competing output VCB with the highest precedence will be allowed to send the next frame.

3.3.3.3 Bandwidth Reservation

When connecting an instrument via a network to a mass memory, what the systems engineer needs to know is “how much bandwidth do I have to transfer data from the instrument to the mass memory?” Once the network bandwidth allocated to a particular instrument has been specified, it should not be possible for another instrument to impose on the bandwidth allocated to that instrument. A priority mechanism is not suitable for this application. If an instrument with high priority has data to send it will hog the network until all its data has been sent. What is needed is a mechanism that allows bandwidth to be reserved for a particular instrument.

Bandwidth reservation calculates the bandwidth used by a particular virtual channel, and compares this to the bandwidth reserved for that virtual channel to calculate the precedence for that virtual channel. If the virtual channel has not used much reserved bandwidth recently, it will have a high precedence. When a data frame is sent by this virtual channel, its precedence will drop. Its precedence will increase again over a period of time. If a virtual channel has used more than its reserved bandwidth recently, it will have a low precedence.

A virtual channel specifies a portion of overall Link Bandwidth that it wishes to reserve and expects to use, i.e. its Expected Bandwidth.

When a frame of data is sent by any virtual channel, each virtual channel computes the amount of bandwidth that it would have been permitted to send in the time interval that the last frame was sent. This is known as the Bandwidth Allocation.

Each virtual channel can use this to determine its Bandwidth Credit, which is effectively the amount of data it can send and still remain within its Expected Bandwidth. Bandwidth Credit is the Bandwidth Allowance less the Bandwidth Used accumulated over time.

The Bandwidth Credit is updated every time a data frame for any virtual channel has been sent. A Bandwidth Credit value close to zero indicates nominal use of bandwidth by the virtual channel. A negative value indicates that the virtual channel is using more than its expected amount of link bandwidth. A positive value indicates that the virtual channel is using less than its expected amount of link bandwidth.

To simplify the hardware required to calculate the Bandwidth Credit it is allowed to saturate at plus or minus a Bandwidth Credit Limit, i.e. if the Bandwidth Credit reaches a Bandwidth Credit Limit it is set to the value of the Bandwidth Credit Limit.

When the Bandwidth Credit for a virtual channel reaches the negative Bandwidth Credit Limit it indicates that the virtual channel is using more bandwidth than expected. This may be recorded in a status register and used to indicate a possible error condition. A network management application is able to use this information to check correct utilisation of link bandwidth by its various virtual channels.

For a virtual channel supporting bandwidth reserved QoS, the value of the bandwidth counter provides the precedence value for that virtual channel.

The operation of a bandwidth credit counter is illustrated in Figure 3.




Figure 3 Bandwidth Credit Counter

The bandwidth credit for a particular VC increments gradually. At point (1) a frame is sent from by this VC, resulting in a sudden drop in credit. The size of the drop is amount of data sent in the frame divided by the percentage bandwidth reserved for the VC. This means that the smaller the percentage bandwidth the larger the drop, and hence the longer it takes to regain bandwidth credit.

After the drop at point (1) the bandwidth credit gradually increments until point (2) when another frame is sent by the VC. Further frames are sent at points (3), (4), (5) etc. If the frames sent are full frames then the drop in bandwidth credit every time a frame is sent, will be the same size.

The bandwidth credit counter for another VC is illustrated in Figure 4. This VC has about half the bandwidth of the VC in Figure 3 allocated to it. This means that the drops in bandwidth credit when frames are sent by this VC are about twice the size, as can be seen Figure 4 at points (1), (2) and (3).




Figure 4 Bandwidth Credit Counter with Smaller Reserved Bandwidth

The bandwidth credit counter of another VC is shown in Figure 5. In this case the bandwidth credit slowly increments and although some frames are sent at points (1), (2) and (3), the bandwidth credit eventually saturates, reaching its maximum permitted value at point (4). Although more bandwidth should be accumulated after point (4) this is effectively ignored since the maximum possible bandwidth credit has been reached. At point (5) a frame is sent once more, resulting in a drop from the maximum bandwidth credit value.




Figure 5 Bandwidth Credit Counter Reaching Saturation

All three VCs are shown together in Figure 6. When a VC has a data frame ready to send and room for a full data frame at the other end of the link, it competes with any other VCs in a similar state, the one with the highest bandwidth credit being allowed to send the next data frame. At points (1), (2) and (3) the red VC has data to send and sends frames. At points (4), (5) and (6) the green VC has data to send and sends a data frame. At point (7) both the blue and the red VCs have data to send. The blue VC wins since it has the highest bandwidth credit count. After this the red VC is allowed to send a further data frame at point (8).




Figure 6 Bandwidth Credit of Competing VCs

If the bandwidth credit counter reaches the minimum possible bandwidth credit value, it indicates that it is using more bandwidth than expected and a possible error may be flagged. This condition may be used to stop the VC sending any more data until it recovers some bandwidth credit, to help with “babbling idiot” protection.

Similarly if the bandwidth credit counter stays at the maximum possible bandwidth credit value for a relatively long period of time, the VC is using less bandwidth than expected and this condition can be flagged to indicate a possible error.

The bandwidth credit value is the precedence used by the medium access controller to determine which VC is permitted to send the next data frame.

3.3.3.4 Priority

The second type of QoS provided by VCs is priority. Each VC is assigned a priority value and the VC with the highest priority (lowest priority number) is allowed to send the next data frame as soon as it is ready. Figure 7 shows three priority levels. SpaceFibre has 16 priority levels.




Figure 7 Multi-Layered Precedence Priority QoS

Within any level there can be any number of VCs which compete amongst themselves based on their bandwidth credit. A higher priority VC will always have precedence over a lower priority VC unless its Bandwidth Credit has reached the minimum credit limit in which case it is no longer allowed to send

any more data frames. This prevents a high priority VC from consuming all the link bandwidth if it fails and starts babbling. More than one VC can be set to the same priority level in which case those VC's will compete for medium access using bandwidth reservation.

3.3.3.5 Scheduled

To provide fully deterministic data delivery it is necessary for the QoS mechanism to ensure that data from specific virtual channels can be sent (and delivered) at particular times. This can be done by chopping time into a series of time-slots, during which a particular VC is permitted to send data frames. This is illustrated in Figure 8.

Time-slot	1	2	3	4	5	6	7	8
VC 1								
VC 2								
VC 3								
VC 4								
VC 5								
VC 6								
VC 7								
VC 8								

Figure 8 Scheduled Quality of Service

Each VC is allocated one or more time-slots in which it is permitted to send data frames. VC1 is scheduled to send in time-slot 1 and VC2 is scheduled to send in time-slots 2 and 6. The time-slot duration is a system level parameter, typically 100 μ s, and there are 256 time-slots.

During a time-slot, if the VC is scheduled to send in that time-slot, it will compete with other VCs also scheduled to send in that time-slot based on precedence (priority and bandwidth credit). A fully deterministic system would have one VC allowed to send in each time-slot.

The schedule is always operating. If a user does not want to use scheduling the schedule table is simply filled completely, allowing any VC to send in any time-slot, competing with precedence.

Scheduling can waste bandwidth if only one VC is allowed to send in a time-slot and that VC is not ready. To avoid this situation, the critical VC can be allocated a time-slot and given high priority and another VC can be allocated the same time-slot with lower priority. In this way when that time-slot arrives the high priority VC will be allowed to send its data, but if it is not ready the VC with lower priority can send some data. This configuration is illustrated in Figure 8 time-slot 3 and VCs 6 and 8.

Time-slots can be defined by broadcasting start of time-slot signals, or by broadcasting time and having a local time counter which determines the start and end of each time-slot. The SpaceFibre broadcast message mechanism supports both synchronisation and time distribution.

The SpaceFibre QoS mechanism is simple and efficient to implement and it provides bandwidth reservation, priority and scheduling integrated together, not as separate options. Furthermore SpaceFibre QoS provides a means for detecting “babbling idiots” and for detecting nodes that have ceased sending data when they are expected to be sending information.

3.3.4 SpaceFibre Fault Detection, Isolation and Recovery

The FDIR mechanism for SpaceFibre is another significant output of the SpaceWire-RT project. A brief overview of the FDIR for SpaceFibre is provided in this section.

SpaceFibre provides automatic fault detection, isolation and recovery. When a fault occurs on a SpaceFibre link, it is detected and the erroneous or missing information resent. SpaceFibre recovers from intermittent faults very rapidly, detecting faults, recovering and resending data faster than SpaceWire disconnects and reconnects a link. The retry mechanism does not depend on time-outs, naturally adapting to different cable delays.

Fault detection is provided by checking each 8B/10B symbol for disparity errors and invalid 8B/10B codes. SpaceFibre has selected the 8B/10B K-codes it uses to have enhanced Hamming distance from data-codes. This means that a single bit error occurring in a data-code cannot result in a valid K-code used by SpaceFibre. In addition each data frame, broadcast frame, FCT, ACK and NACK are protected by a CRC.

Fault isolation is provided at various levels in SpaceFibre. AC coupling is used in the physical layer to prevent damage from faults that cause DC voltages exceeding the maximum permitted to appear on the transmitter outputs or receiver inputs. This feature also enables galvanic isolation to be implemented readily. At the Quality level SpaceFibre provides time containment, containing errors in the data frame in which they occur, and bandwidth containment, containing errors to the virtual channel in which they occur; an error in one VC does not affect data flowing in another VC. Babbling idiots are contained using the QoS mechanism described above.

Fault recovery is provided at the link level using a retry mechanism that resends data frames, broadcast frames and FCTs. The retry is very fast, uses a minimum amount of buffer memory, and adapts automatically to different link lengths. In addition to the retry mechanism the multi-lane functionality includes graceful degradation on lane failure. If a lane fails permanently, so that a retry or re-initialisation does not recover lane operation, a multi-lane system will continue using the remaining lanes available. This reduces the bandwidth available but does not stop the link operating. For critical operations an extra lane can be included and the graceful degradation will then provide automatic replacement of a faulty lane. The bit error rate (BER) of a lane is monitored and a lane reported as faulty if the (BER) is above a level which results in the effective link bandwidth being unusable. This feature allows lanes that can re-initialise successfully but which will not run for very long before having to re-initialise again, to be detected, isolated and replaced by a fully functional lane.

3.3.5 SpaceFibre Networks

The third substantial contribution of the SpaceWire-RT project is a baseline concept for the SpaceFibre network layer. Initial validation of the network layer has been carried out by SUAI using SystemC.

A SpaceFibre network uses similar packet formats, packet addressing and routing concepts as SpaceWire. The main difference is that SpaceFibre includes virtual channels.

A SpaceFibre router is illustrated in Figure 9.




Figure 9 SpaceFibre Router

The SpaceFibre router comprises a number of SpaceFibre interfaces and a routing switch matrix. Each SpaceFibre interface has several virtual channels. The VC number for each virtual channel can be configured, except for VC0 which is a virtual channel used for configuration, control and monitoring of the SpaceFibre network. When a packet arrives on a SpaceFibre interface it is placed in the appropriate virtual channel, i.e. the one with the same VC number as it was transmitted on. The leading data character of the packet determines which port of the routing switch the packet is to be forwarded through using either path or logical addressing. The port that it is to be switched to must have a VC configured with the same number as the VC that the packet arrived on. The packet is then passed through the routing switch matrix and placed frame by frame in the VC of the output port. The packet is then transferred across the SpaceFibre link, competing with other VCs in that port for access to the link medium according to their precedence.

If a packet arrives and the output port that the packet is to be switched to does not have a VC with the same number as that on which it arrived, the packet is spilt and an error recorded.

Virtual channels can be used to construct virtual networks where a single VC number is used for connecting to all or several of the nodes attached to the network. This is illustrated in Figure 10 where VC6 (blue) is used to connect all the nodes on the network. Using VC6 the Control Processor can send commands to Instrument 1 or 2 or the Mass Memory unit, setting their operating mode or reading housekeeping information, etc. This virtual network acts like a SpaceWire network.




Figure 10 A Simple SpaceFibre Network

Virtual channels can also be used to construct virtual point to point links from one node to another. VC2 and VC4, in Figure 10, are providing virtual point to point links. VC2 provides a virtual point to point link between Instrument 2 and the Mass Memory Unit and VC4 between Instrument 1 and the Mass Memory. These virtual channels can be each allocated the bandwidth they need to send their data to the Mass Memory Unit. Once this bandwidth is allocated other virtual channels or virtual networks will not interfere with their operation.

Figure 11 shows a more realistic onboard network using SpaceFibre which includes a SpaceWire to SpaceFibre Bridge. Two high data-rate instruments (Instruments 1 and 2) have SpaceFibre connections. Four less demanding instruments have SpaceWire connections to the SpaceWire to SpaceFibre Bridge. Each instrument has a virtual point to point connection to the Mass Memory Unit and there is a virtual point to point connection between the Mass Memory and the Downlink Telemetry Unit. The Control Processor has a virtual network for configuring and controlling all devices on the network.




Figure 11 Realistic SpaceFibre Network

Figure 11 is solving a complex communication task with many separate, isolated virtual channels providing point to point links, and a virtual network being used to control the entire system. Figure 12 shows this same network with the virtual channels removed, revealing the simplicity of implementation of a complex communication task when using SpaceFibre.




Figure 12 Simple System Architecture with SpaceFibre

3.3.6 Physical Layer

SpaceFibre is a very high-speed communication protocol. The QoS and FDIR protocols developed for SpaceFibre can potentially be applied to a range of slower speed network technologies. If a consistent set of QoS and FDIR protocols can be applied across network technologies like SpaceWire then the concepts will be common and application software can be readily reused regardless of the performance required from the network. SpaceWire-RT aims to use the QoS and FDIR protocols being developed for SpaceFibre across a range of different lower level protocol layers to cover a range of data rates and applications.

At present two physical layer protocols are being considered:

SpaceFibre CML: The SpaceFibre protocol is designed to run at multiple Gbits/s using current mode logic (CML) over copper cables for distances up to 5 m or fibre optic cables for distances of 100 m or more. Multiple lanes can be operated in parallel to increase the data rate. SpaceFibre uses 8B/10B encoding and a phase locked loop (PLL), or similar technology, to recover a clock from the received data stream and to perform bit synchronisation.

SpaceFibre LVDS: The SpaceFibre protocol could also be run at slower speed using Low Voltage Differential Signalling (LVDS), which is capable of 600 Mbits/s or possibly up to 1 Gbits/s data rates.

LVDS is used in SpaceWire and therefore has space heritage. SpaceFibre LVDS would use 8B/10B encoding and a PLL for bit synchronisation.

It is proposed that both the CML and LVDS options for SpaceFibre are included in the SpaceFibre standard specification.

3.3.7 Oversampled SpaceFibre

SpaceFibre requires clock-data recovery circuitry in the receiver to recover the clock from the received data bit-stream. Normally this is accomplished using some form of phase-locked loop (PLL) which requires specialised circuitry. Current space qualified FPGAs do not provide suitable PLL clock recovery circuitry, so an alternative clock recovery scheme is required if SpaceWire-RT is to be implemented in space qualified FPGAs. The receive signal can be over sampled (sampled with a clock of significantly higher frequency than the receive bit rate) and bit synchronisation achieved [8]. Data rates of up to 100 Mbits/s should be achievable, which is adequate for many spacecraft applications. This oversampling technique would simply be an alternative implementation for the clock-data recovery for SpaceFibre and be suitable at relatively low data-rates (e.g. 100 Mbits/s or slower).

3.4 WP3 SpaceWire-RT Validation and Simulation

One of the main objectives of the WP3 was to check the correctness and consistency of the SpaceWire-RT specification. This was done by specification and simulation using SDL for the lower layers of SpaceFibre and by simulation using SystemC for the network layer. Simulation and investigation was split into two activities: checking correctness and consistency of SpaceWire-RT and testing against requirements. From the start of WP3 three drafts of the SpaceFibre specification were released which were taken as the basis for the SpaceWire-RT standard. During specification and simulation several documents with proposals, change requests and issues for discussion were produced and provided to the University of Dundee. Overall, SUAI produced 139 change requests, 87 of which resulted in changes to the SpaceFibre standard. In this way the results of specification and simulation in SDL and SystemC were applied during the SpaceWire-RT development.

In the scope of testing SpaceWire-RT against the requirements of industry a number of tests were performed, which evaluated and checked SpaceWire-RT functionality in point-to-point connection (primarily SDL model) and in networks with different topologies (SystemC model). To test basic mechanisms of the SpaceWire-RT two different network configurations were used: point-to-point configuration and mixed configuration (tree and circular configurations combined in one network). Overall this testing showed that SpaceWire-RT satisfies the requirements provided by the industry.

3.4.1 SDL Point-to-Point Model

The SDL model was primarily focused on checking internal functionality and how all mechanisms of all layers work in common in one node. The SDL model was layered according to the SpaceFibre

specification and service access points (SAPs) were specified as the interfaces between layers, as shown in Figure 13. Subsequently the layering in SpaceFibre was reduced as illustrated in Figure 2.




Figure 13 Service access points for the SpaceFibre protocol stack

3.4.2 SystemC Network Model

SystemC network models were primarily focused on simulation and checking of network aspects of the SpaceWire-RT, performance parameters (e.g. latency) and some features that cannot be covered by SDL point-to-point simulation (e.g. broadcast distribution, QoS, etc.). The network structure depicted in Figure 14 was used to test the Network layer aspects of SpaceFibre using SystemC. This structure is a mixture of tree and circular configurations giving the ability to test all possible situations described in the test plan using one common test system.




Figure 14 Testing network structure

3.5 WP4 SpaceWire-RT VHDL IP Core Development

3.5.1 SpaceFibre VHDL IP Core

The QoS and FDIR capabilities for SpaceFibre were implemented in an experimental SpaceFibre VHDL IP core so that the concepts could be tested in hardware. The design was implemented in a STAR Fire unit from STAR-Dundee, which contained a suitable FPGA with SpaceWire and SpaceFibre connectors attached to the FPGA. The design was extensively tested using the inbuilt SpaceFibre link analysis capabilities of the STAR Fire unit. Figure 15 shows the first successful implementation of SpaceFibre with data flowing over SpaceFibre being recorded on a logic analyser.




Figure 15 First Implementation of SpaceFibre

Following testing of the hardware implementation, improvements were made to the SpaceFibre specification and the SpaceFibre VHDL IP core was updated and re-tested. The design, simulation, implementation, test, and update to the specification cycle were iterated several times as problems were discovered, solutions devised, and the specification polished. This cycle of improvement ran concurrently with the improvements resulting from the simulation of the standard specification in WP3. Figure 16 shows three STAR Fire units being used to test SpaceFibre. The black cables plugged into the front panel are the SpaceFibre electrical cables. The FPGA is the largest chip on the right hand side of the unit.




Figure 16 Testing SpaceFibre using the STAR Fire Units

3.5.2 Oversample SpaceFibre VHDL IP Core

An oversampling technique suitable for SpaceWire-RT based on [8] was designed, simulated, implemented and tested in an FPGA. SpaceFibre ran successfully in the FPGA using oversampling for clock-data recovery. The simulation results are shown in Figure 17.




Figure 17 Behavioural Simulation Results

The design was implemented in a pair of STAR-Dundee STAR Fire units. The test arrangement is illustrated in Figure 18. The two STAR Fire units were modified so that each unit had one normal SpaceWire interface (Port 1) and one Oversampled SpaceFibre interface (Port 2). A host PC provided a stream of SpaceWire packets via a SpaceWire USB Brick which were sent to the SpaceWire interface on one STAR Fire unit. The SpaceWire packets were then sent out of the Oversampled SpaceFibre interface to the other STAR Fire unit. The SpaceWire packets were received and looped-back using a SpaceWire router inside the STAR Fire unit so that they were sent back out of the Oversampled SpaceFibre interface. The SpaceWire packets were received back at the first STAR Fire unit and passed up to the host PC for checking, via the SpaceWire USB Brick.




Figure 18 Hardware Testing - Data flow

Software running on the host PC checked the SpaceWire packets received noting any errors. The “Sink Statistics” dialog box, shown in Figure 19, displays the total data packets received, data rate statistics and any detected errors. This shows that packets are successfully being received after passing through the STAR Fire units using the oversampling technique. The “Sink Statistics Graph” display in Figure 19 shows the data rate and link utilisation against time. This shows a disconnection event when the SpaceWire cable used for the SpaceFibre links was removed and replaced. During this event, around time -10 seconds, no errors were detected by the SpaceWire validation software

(as can be seen in the “Sink Statistics” window), showing that the link disconnection was handled as expected by SpaceFibre’s FDIR capabilities and all errors that occur are recovered transparently.




Figure 19 SpaceWire Validation Software Output

Figure 20 provides the STAR Fire analyser output and frame view for one of the SpaceFibre links, showing the data being received as it passes through the system. Expected SpaceFibre operation was observed using the Oversampled SpaceFibre design.




Figure 20 STAR Fire Analyser Output: Word and Frame Views

3.6 WP5 SpaceWire-RT ASIC Feasibility and Virtual Prototyping

From the SpaceWire-RT Outline Specification (WP2) and VHDL IP Core Design (WP4) the key issues for implementation of SpaceWire-RT in space qualifiable ASIC technology were identified:

- Size of memory required for virtual channel buffering which results in increasing ASIC size and power consumption, unless memory blocks are used.
- Data rate that can be handled by the physical layer drivers (CML or CML compatible and LVDS) and is limited by the space qualifiable ASIC technologies.
- Power consumption is less than 200 mW (175 mW for the 180nm SpaceWire-RT ASIC IP core) at increasing data rates in work range.

The approach taken in the SpaceWire-RT ASIC feasibility is summarised in Figure 21




Figure 21 SpaceWire-RT ASIC IP Core Design Flow

Available space qualifiable (Radiation Tolerant) ASIC technologies were reviewed, and the area and timing parameters estimations for the ASIC technologies from 250 nm up to 65 nm were derived. The main results of the ASIC feasibility analysis are as follows:

Area: The area (6.87- 8.73 mm²) of the SpaceWire-RT ASIC IP-core on the base 180nm Radiation Tolerant Libraries is approximately 3-4 times larger than a SpaceWire ASIC IP-core (1-8 virtual

channels). Of course, the SpaceWire-RT IP core includes all advantages of QoS, FDIR and other SpW-RT benefits as well as operating at a much higher data-rate than SpaceWire. The area for the memory for eight virtual channels is about 1.78 mm² for the 180 nm technology and is expected to be reduced with the new Memory compiler for the ELVEES RT Library. The estimated area of the SpaceWire-RT ASIC IP Core for various ASIC technologies is summarised in Figure 22.




Figure 22 SpaceWire-RT ASIC IP-Core Area Estimation

Data-Rate: The data-rate of the SpaceWire-RT ASIC IP-core that can be achieved on the base 180 nm Radiation Tolerant Libraries is from the lower level (e.g., 5 Mbp/s) up to 1.25 Gbits/s. Up to four lanes could be supported with this technology increasing the data rate to 5Gbits/s with four lanes. A data rate of 2.5 Gbits/s is expected to be achievable on a 130 nm CMOS process with 10Gbps achievable using four lanes.

The implementation of SpaceWire-RT IP Core in space qualifiable ASIC technology has been shown to be feasible with data rates of 2.5 Gbit/s per lane being expected using 130 nm technology.

The analysis of Work Package 5 results shows that the SpaceWire-RT RT ASIC IP Core is effectively realizable on the basis of existing Russian and European Microelectronics Technologies.

3.7 WP6 SpaceWire-RT Standard Draft

During the SpaceWire-RT project the SpaceFibre standard was extended with the Quality and Network layers being added. Much of the SpaceFibre standard was simulation in SDL and SystemC and the results of the simulation used to help improve all layers of the standard specification. SpaceFibre was implemented as an IP core suitable for implementation and testing in an FPGA. The results of this hardware implementation were very useful in finding problems with and improving the SpaceFibre specification.

The SpaceFibre standard specification was updated to draft F and has been translated into Russian. It is expected that only relatively minor changes will be made to the Quality and Lane layers that form

the core of the SpaceFibre specification. The Network layer provides a baseline for further development and discussion of suitable network layer approaches.

There are three forms of SpaceWire-RT being incorporated into the SpaceFibre standard specification, the key characteristics of which are summarised in Table 3.

Table 3 SpaceWire-RT Protocol Characteristics			
Characteristic	SpFi-FO	SpFi-CML	SpFi- LVDS
Media	Fibre Optic	Copper CML	Copper LVDS
Encoding	8B/10B	8B/10B	8B/10B
Speed Range	0.1 to 20 Gbits/s 50 Gbits/s in future	0.1 to 20 Gbits/s 50 Gbits/s in future	1 to 600 Mbits/s 1 to 100 Mbits/s OS
Distance	100 m	5 m	10 m
Galvanic Isolation	Yes	Yes	Yes
Packet Size	Arbitrary	Arbitrary	Arbitrary
SpaceWire Packet Level	Yes	Yes	Yes
Latency (TBC)	<0.5 μ s	<0.5 μ s	1 μ s
Cable Mass	< 30g/m	< 30g/m	< 30g/m
Power (TBC)	< 200 mW	< 200 mW	< 200 mW
QoS BW Reserved	Yes	Yes	Yes
QoS Priority	Yes	Yes	Yes
QoS Scheduled	Yes	Yes	Yes
QoS Best Effort	Yes	Yes	Yes
Broadcast Message	Yes	Yes	Yes
Determinism	Yes	Yes	Yes
Reliability	Yes	Yes	Yes
Fault Detection	Yes	Yes	Yes
Fault Isolation	Yes	Yes	Yes
Retry	Yes	Yes	Yes
SpaceWire compatible	Packet Level	Packet Level	Packet Level

3.8 SpaceWire-RT Objectives Addressed

The SpaceWire-RT objectives have been fully addressed:

- A comprehensive set of requirements and use cases for spacecraft onboard networking has been produced.
- The SpaceFibre networking technology has been selected as the basis for SpaceWire-RT and essential, innovative QoS and FDIR capabilities designed to complement SpaceFibre.
- A baseline network layer for SpaceFibre has been devised.
- A validated SpaceWire-RT specification has been written with important, novel features of SpaceWire-RT networks tested using SDL and SystemC models.
- The feasibility of implementation in space qualifiable ASIC technologies has been assessed and demonstrated through simulation.
- The results of the SpaceWire-RT study have been disseminated to the European and Russian space industries, and to the international space community.
- The SpaceWire-RT technology has been reviewed at various points during the two years of the SpaceWire-RT project by the International SpaceWire Working Group.
- A draft standard document for SpaceWire-RT has been produced.

4 Potential Impact, Main Dissemination Activities and Exploitation of Results

4.1 Potential impact

SpaceWire is a European technology used by the world's space agencies and space industry on many spacecraft. The main objective of the SpaceWire-RT project was to take this technology to the next level, by providing an enhanced SpaceWire network technology that provides quality of service capabilities suitable for spacecraft data-handling and control applications, enabling it to support rapid spacecraft assembly and also making it applicable to other applications. The creation of this technology has substantially strengthen collaborative bonds between the Russian and European organisations involved in the research, and led to technology of vital importance for future space missions.

Building on the SpaceFibre technology being designed by University of Dundee for ESA, the SpaceWire-RT project has developed a high-performance network technology that fulfils the needs of many spacecraft on-board communication requirements. Specifically the SpaceWire-RT project designed Quality of Service (QoS) and Fault Detection, Isolation and Recovery (FDIR) techniques for SpaceFibre and specified a baseline Network layer. This research resulted in the specification of the Quality and Network layers for SpaceFibre. In addition the entire SpaceFibre protocol stack has been simulated and the results of this simulation used to improve the existing SpaceFibre protocols. A VHDL implementation of the Quality layer has been developed and used to test this layer together with the rest of the SpaceFibre protocol stack, again resulting in improvements to the specification. The feasibility of implementation as flight qualified chips has been explored and shown to be achievable. The resulting inputs to the SpaceFibre specification have been presented to the international SpaceWire working group and incorporated into the draft SpaceFibre standard. It is expected that this will become a formal ECSS specification in 2014/15.

The research on the Quality layer has been very important for SpaceFibre, providing a network technology that is robust against failure, while retaining high-performance, and which can support different types of traffic flowing over a single link. The QoS technique developed combines priority, bandwidth-reservation and scheduling techniques to support just about all spacecraft communication applications with a single network technology. Research on over-sampling has demonstrated that a SpaceFibre interface operating at 100 Mbits/s can be supported without requiring analogue phase-locked loop clock/data recovery technology, allowing it to be implemented in any FPGA technology.

SpaceFibre is likely to be adopted widely with implementations by European, Japanese and Russian organisations already underway. It will support very high-data rate missions like synthetic aperture radar and multi-spectral imaging. It will support integrated avionics applications with control and housekeeping information flowing over the same network as payload data, reducing overall power and mass budgets for the on-board communication network. SpaceFibre is backwards compatible

with the packet level of SpaceWire and will support mixed SpaceWire/SpaceFibre systems allowing ready migration to the improved SpaceFibre network technology.

SpaceFibre is designed primarily for spacecraft applications, but its key characteristics of high-performance, low-latency, integrated QoS, integrated FDIR and implementation simplicity are likely to make it attractive for terrestrial applications too. For example, robotics applications will benefit significantly from SpaceFibre's capabilities.

SpaceFibre will underpin many future science, exploration, Earth observation and commercial missions. These missions relying on SpaceFibre technology will support new scientific discoveries, will help monitor and protect our environment, and will enhance our quality of life through improved global telecommunications, positioning, disaster monitoring and weather prediction satellites.

The SpaceWire-RT project will have a positive economic impact on the SME and industrial partners who have been involved in the project. The partners are already planning to commercially exploit the results of the SpaceWire-RT project through new or improved products which will be available within the next 1 to 7 years. The SME partners also expect this will have a positive impact on employment resulting in approximately 10 new jobs.

Finally, the SpaceWire-RT project will also have an educational impact with the academic partners planning to incorporate elements of the project results into their undergraduate and postgraduate courses.

4.2 Main dissemination activities

Dissemination took the following forms:

- The SpaceWire-RT project WEB site;
- Participation in SpaceWire Working Group meetings and scientific conferences;
- Project Newsletters;
- Publications.

The SpaceWire-RT website (<http://spacewire-rt.org>) provides information on the project for both the general public and the research community. The **News** page is aimed at a wide audience and contains links to the project newsletters and news articles related to the project partners. The **Project** and **Publications** pages are aimed more towards the research community. These pages contain an overview of the project objectives, downloadable copies of the non-confidential project deliverable reports and details of project-related papers presented at scientific conferences and SpaceWire Working Group meetings.

Over the course of the project, all partners have been involved in presenting the key features and concepts of SpaceWire-RT at a wide variety of forums. This included four SpaceWire Working Group meetings, regular meetings of the Russian National SpaceWire Working Group, six scientific conferences in Europe and Russia and three educational conferences in Russia. The Working Group

meetings provided the main vehicle for gaining feedback and validating the future plans at each stage of the project and enabling early engagement with the end user community.

Dissemination activities have continued beyond the end of the project with a number of project-related papers being published in the Radioelectronics, Engineering Technology. – 2013 journal published in Russian by ELVEES in June 2013 and a SpaceFibre tutorial is planned for the SpaceWire 2013 conference being held in June 2013.

4.3 Exploitation of Results

The results of the SpaceWire-RT project will be exploited by the SpaceWire-RT partners in a number of ways:

SpaceFibre standard specification

The primary results of the SpaceWire-RT project, including the QoS and FDIR technology designed for SpaceFibre, has been included in the SpaceFibre standard specification. The University of Dundee will take this specification through the ESA ECSS standardization process starting in 2014, to become a formal ECSS standard. SpaceFibre, incorporating the results of the SpaceWire-RT activity, is expected to become an important technology for future spacecraft on-board data handling applications.

Development of new or enhanced products based on the SpaceWire-RT technology

The SME and industrial partners plan to develop new products and enhance existing products based on the SpaceWire-RT technology.

STAR-Dundee Ltd, a successful spin-out company from the University of Dundee, will commercialise the SpaceWire-RT technology, including the SpaceFibre VHDL IP Core. It is expected that several chip designs incorporating the SpaceFibre VHDL IP Core will be developed and used in many future space missions including multi-spectral imagers and synthetic aperture radar.

STAR-Dundee Ltd will also develop test equipment for SpaceFibre including the QoS and FDIR capabilities from SpaceWire-RT. The STAR-Dundee “STAR Fire” unit already provides interface and analysis capabilities for the current draft of SpaceFibre. Additional test equipment will be designed to support SpaceFibre adoption by the international aerospace industry.

ELVEES, in collaboration with the industrial partners and Universities, plan to design a Radiation Tolerant SpaceWire-RT IP-library, SpaceWire-RT ASIC IP-cores and SpaceWire-RT based chipsets (Multicore DSP microprocessor with the integrated on the silicon SpaceWire-RT links router, microcontroller for Mass Storage with the integrated on the silicon SpaceWire-RT links router, multiports SpaceWire-RT router and multi-channel adapter) for a variety of microelectronics processes of the Russian-European factories.

SUBMICRON will integrate the results of SpaceWire-RT into the design of onboard data systems and control equipment for spacecraft and failure-tolerant structure of onboard computing systems of prospective satellites and piloted transport spacecraft with high-speed information interfaces.

ASTRIUM GmbH will design boards, boxes and systems containing SpaceWire-RT and SpaceWire networks for space applications.

Education and Training

SpaceFibre will be used as an example of a high-performance network technology within the “Research Frontiers” final year undergraduate computing course at the University of Dundee. Postgraduates and Internees within the University of Dundee Space Technology Centre and aerospace industry engineering staff will also receive training on SpaceFibre.

St. Petersburg State University of Aerospace Instrumentation (SUAI) will make use of the SpaceWire-RT on-board networking technology in its education dedicated to aerospace technologies, space data systems and instruments engineering. SUAI will also provide training and maintenance during the uptake of the SpaceWire-RT technology by Russian space industry engineering staff.

5 Project Details

Title SpaceWire-RT

Coordinator



Prof Steve Parkes
Space Technology Centre
University of Dundee
United Kingdom

Consortium



St Petersburg State University of Aerospace Instrumentation,
Russian Federation
www.suai.ru



SubMicron,
Russian Federation
www.submicron.ru



Electronic VLSI Engineering and Embedded Systems,
Russian Federation
<http://multicore.ru/>



Astrium GmbH,
Germany
www.astrium.eads.net

Duration 1st June 2011 - 31st May 2013 (24 months)

Funding Scheme FP7 SPACE, topic SPA.2010.3.2-04 "EU-Russia Cooperation for Strengthening Space Foundations"

Budget EU contribution: €499,997

Website <http://spacewire-rt.org>

For more information spacewire@dundee.ac.uk

6 References

- [1] European Cooperation for Space Standardization, Standard ECSS-E-50-12A, "SpaceWire, Links, Nodes, Routers and Networks", Issue 1, European Cooperation for Space Data Standardization, January 2003.
- [2] European Space Agency, "SpaceWire Web Page", European Space Agency, <http://spacewire.esa.int/>.
- [3] S. Parkes, "SpaceWire User's Guide", available from www.star-dundee.com.
- [4] S. Parkes, C. McClements, M. Dunstan and M. Suess, "SpaceFibre: Gbits/s Link For Use On Board Spacecraft", International Astronautical Congress, Daejeon, South Korea, October 2009.
- [5] S. Parkes and M. Suess, "Virtual Channels, Broadcast Channels and SpaceFibre", International SpaceWire Conference, November 2011, San Antonio, USA, available from <http://2011.spacewire-conference.org/proceedings/>.
- [6] S. Parkes, A. Ferrer, A. Gonzalez-Villafranca, C. McClements, "SpaceFibre Standard", Draft D, University of Dundee, February 2012, available from <http://spacewire.esa.int/WG/SpaceWire/SpW-WG-Mtg18-Proceedings/>
- [7] S. Parkes, A. Ferrer, A. Gonzalez-Villafranca, C. McClements, "SpaceFibre Standard", Draft E1, University of Dundee, September 2012, available from <http://spacewire.esa.int/WG/SpaceWire/SpW-WG-Mtg19-Proceedings/>
- [8] N. Sawyer, "Data Recovery", Xilinx application note XAPP224 (v2.5), July 2005, available from www.xilinx.com.