Reliable Transport over SpaceWire for James Webb Space Telescope (JWST) Focal Plane Electronics (FPE) Network

Glenn Rakow	Richard Schnurr	Christopher Dailey	Kamdin Shakoorzadeh
NASA GSFC	NASA GSFC	QSS Group, Inc.	QSS Group, Inc.
Code 561	Code 560	GSFC Code 561	Code Code 561
Greenbelt, MD	Greenbelt, MD	Greenbelt, MD	Greenbelt, MD
301-286-5993	301-286-1852	301-286-3302	301-286-9854
Glenn.P.Rakow@nasa.gov	Richard.G.Schnurr@nasa.gov	cdailey@pop500.gsfc.nasa.gov	kshakoor@pop700.gsfc.nasa.gov

Abstract—NASA's James Webb Space Telescope (JWST) faces difficult technical and budgetary challenges to overcome before it is scheduled launch in 2010. The Integrated Science Instrument Module (ISIM), shares these challenges. The major challenge addressed in this paper is the data network used to collect, process, compresses and store Infrared data. A total of 114 Mbps of raw information must be collected from 19 sources and delivered to the two redundant data processing units across a twenty meter deployed thermally restricted interface. Further data must be transferred to the solid-state recorder and the spacecraft.

The JWST detectors are kept at cryogenic temperatures to obtain the sensitivity necessary to measure faint energy sources. The Focal Plane Electronics (FPE) that sample the detector, generate packets from the samples, and transmit these packets to the processing electronics must dissipate little power in order to help keep the detectors at these cold temperatures.

Separating the low powered front-end electronics from the higher-powered processing electronics, and using a simple high-speed protocol to transmit the detector data minimize the power dissipation near the detectors. Low Voltage Differential Signaling (LVDS) drivers were considered an obvious choice for physical layer because of their high speed and low power.

The mechanical restriction on the number cables across the thermal interface force the Image packets to be concentrated upon two high-speed links. These links connect the many image packet sources, Focal Plane Electronics (FPE), located near the cryogenic detectors to the processing electronics on the spacecraft structure.

From 12 to 10,000 seconds of raw data are processed to make up an image, various algorithms integrate the pixel data Loss of commands to configure the detectors as well as the loss of science data itself may cause inefficiency in the use of the telescope that are unacceptable given the high cost of the observatory. This combination of requirements necessitates a redundant/fault tolerant high-speed, low mass, low power network with a low Bit error Rate (1E-9 – 1E-12).

The ISIM systems team performed many studies of the various network architectures that meeting these requirements. The architecture selected uses the SpaceWire protocol, with the addition of a new transport and network layer added to implement end-to-end reliable transport. The network and reliable transport mechanism must be implemented in hardware because of the high average information rate and the restriction on the ability of the detectors to buffer data due to power and size restrictions.

This network and transport mechanism was designed to be compatible with existing SpaceWire links and routers so that existing equipment and designs may be leveraged upon. The transport layer specification is being coordinated with European Space Agency (ESA), SpaceWire Working Group and the Consultative Committee for Space Data System (CCSDS) P1K Standard Onboard Interface (SOIF) panel, with the intent of developing a standard for reliable transport for SpaceWire. Changes to the protocol presented are likely since negotiations are ongoing with these groups.

A block of RTL VHDL that implements a multi-port SpaceWire router with an external user interface will be developed and integrated with an existing SpaceWire Link design. The external user interface will be the local interface that sources and sinks packets onto and off of the network (Figure 3). The external user interface implements the network and transport layer and handles acknowledgements and re-tries of packets for reliable transport over the network. Because the design is written in RTL, it may be ported to any technology but will initially be targeted to the new Actel Accelerator series (AX) part. Each link will run at 160 Mbps and the power will be about 0.165 Watt per link worst case in the Actel AX.

[1] 0-7803-7651-X/03/$17.00 © 2003 IEEE

[2] IEEEAC paper #1264, Updated December 9, 2002

TABLE OF CONTENTS

1. INTRODUCTION

This paper describes the implementation and operation of the high-speed instrument network for the JWST. This network named the Focal Plan Electronics (FPE) network uses the SpaceWire specification and a transport layer that is not part of the SpaceWire Specification. The transport layer is implemented in hardware because of the fast recovery time necessary for the application, and its use is necessary to make SpaceWire robust enough to be used for high reliable applications.

The FPE network provides a high-speed network fabric of scalable point-to-point links (SpaceWire) for the ISIM to communicate. Specifically, it provides the communications between all the Instrument Sensors and the Integrated Science Instrument Module (ISIM) Command & Data Handling (ICDH).

The SpaceWire link design is based upon a previous design used on the NASA Swift Mission's Burst Alert Telescope (BAT). The Transport layer design is a new design, and was undertaken to solve the problem of reliable transport for JWST's ISIM's FPE network. This work is being coordinated with the Europe ESA, the SpaceWire working Group, and the CCSDS P1K SOIF panel with the intent of developing a standard transport layer for SpaceWire. There are minor differences that need to be worked out, mostly consolidating upon a single type (remove type 0, and keep type 1 and one other TBD type, see section 6, SpaceWire Transport Layer). Mostly though the differences have to do with byte order and names (the basic concept is the same).

An overview of the JWST mission will be provided followed by an overview of the instrument for JWST. With this background the architecture of the high-speed network for the ISIM will be described with detailed sections on each topic. This paper will not describe the detailed analysis performed to arrive at the solution for the FPE network operation.

2. JWST MISSION OVERVIEW

The goal of JWST is to observe the Universe in its early stages during the formation of stars and galaxies. It will do this by observing in the Infrared spectrum using a deployable 6 meter aperture telescope, allowing it to see objects 400 times fainter than anything available today. The Telescope will be deployed in a L2 (Lagrange point) orbit to allow for more sensitivity in the detectors. The launch date is scheduled for 2010.

The James Webb Space Telescope is a mission with international collaboration between teams in the United States (NASA), Europe (ESA) and Canada (CSA) and is managed by NASA's Goddard Space Flight Center. The United States is delivering the Spacecraft that was awarded to TRW. NASA's Goddard Space Flight Center is managing the ISIM, and Goddard is delivering the ISIM Command and Data Handling (ICDH) electronics, the Focal Plane Electronics (FPEs), and one of the instruments (NIRCam). The Europeans are delivering two of the instruments (NIRSpec and MIRI), and the Canadians are delivering the Fine Guidance Sensor (FGS).

3. ISIM OVERVIEW

The ISIM sensors may be broken down to three instruments and a Fine Guidance Sensor (FGS) (Figure 1). The three instruments are the Near Infrared Camera (NIRCam), the Near Infrared Spectrometer (NIRSpec) and the Mid-Infrared Instrument (MIRI). Each of the three instruments is comprised of an optical assembly and a Sensor Chip Array (SCA). Each instrument has a different field of view and hence a different number of SCAs that comprise the Focal Plane Assembly (FPAs). Each SCA has a Sensor Chip Electronics Module (SCEM) to sample the SCA. Therefore there are a different number of SCEMs for each instrument based upon the field of view of the instrument. The NIRCam has 12 2Kx2K SCAs and therefore 12 SCEMs. The NIRSpec has 2 2Kx2K SCAs and therefore 2 SCEMs. The MIRI has 3 512x512 CCD arrays and therefore 3 SCEMs. In addition, the FGS has 2 CCD and therefore 2 SCEMs. These add up to a total of 19 separate SCEM modules in the ISIM that collect photon data (the source of image data for the instrument). In addition to the detectors, the ISIM has a Command and Data Handling system, ICDH. The ICDH ingests the image data, processes it depending upon the processing algorithm (Fowler n, Up-the ramp, Multi-accum or Cosmic ray rejection), compresses the data and stores it to a Solid State Recorder (SSR).

The ICDH also issues commands to the detectors. It has two interfaces for issuing commands. . The SCEM commands are over the high-speed SpaceWire network, this is for detector bias, etc. It also has a MIL-STD-1553 data bus to a group of subsystems named the Instrument Control Electronics (ICEs) that exist for each instrument and the FGS. These ICEs are responsible for power distribution, mechanism control, Microelectronics Mechanical Systems (MEMS) control, and calibration for the Optical Assemblies.

The ICDH has a microprocessor that is the controller for the ISIM. The SCEMs do not have microprocessors due to power limitations but rather hardware state machines for local control (implemented in a FPGA).

4. ISIM FPE NETWORK ARCHITECTURE

4.1 Breakdown of Network Components

The components that are interconnected within the ISIM network fabric called the FPE network are the SCEMs and the ICDH. As mentioned earlier each instrument has a different number of SCEMs corresponding to the number of CCDs for its instrument. Among the 3 instruments and the FGS there are 19 different SCEMS from which the ICDH needs to gather data. So the network may be simplified into 19 SCEM communicating to the ICDH. Each SCEM has it's own structure and may be stacked to accommodate changes in detector array size. This simplification makes the design modular as the instrument may be modified by adding or subtracting SCEMs.

4.2 Raw Data Rates

Each SCEM generates image data at 4 pixels approximately every 10us. Each pixel is 16 bits, so this translates to 32 bytes of raw image data every 40us. The 40us value was picked as a time to generate an image packet for network throughput purposes. These image packets are segments of a larger packet. Seventeen SCEM (excluding the FGS) may generate data at this rate. The FGS generates average raw data at an estimated 700Kbps rate. This translates approximately 114 Mps of raw data including telemetry that must be transferred from the 19 SCEMs to the ICDH assuming continuous simultaneous operation of the detectors.

The command data rate to the SCEMs from the ICDH is very low less than 2Mbps. Since the data is sent in the opposite direction to the image data it does not drive the requirements for the full duplex link.

4.3 Architecture Considerations

The network was calculated to have an error every 2.3 minutes based upon the data rate bit error rate over LVDS and the network topology. Because the image packets are segmented across the network and the long integration times, lost packets would cause reconstruction problem and hence loss of images or errors in images. This forced the requirement of reliable transport across the network, which was not addressed by the current SpaceWire Standard.

One of the considerations that drove the network architecture was the requirement for single fault tolerance. If a link fails it is necessary to have another path to transfer information across the FPE network.

The network also had to minimize the number of cables across the thermally restrictive interface yet communicate to 19 separate entities.

In addition, synchronization information is needed to start collection of the image data. This may be done point-to-point but cabling is an issue. It would be favorable to do this over a network with delays in the order of 2us.

The network interface will be one of the largest contributors to this power and therefore must be as low power as possible.

All these consideration drove the selection of selection for bus protocol and network architecture.

4.4 Protocol Trades

A trade study was performed on different data bus protocols. The report is titled "The James Webb Space Telescope Integrated Science Instrument Module Command and Data Handling System Bus Trade Study", revision Draft, June 18, 2002.

The protocol studied were as follows: SpaceWire, 1394a, Ethernet over LVDS, Fiber Optic Data, Fiber Optic Data bus over LVDS, USB 2.0, Custom over LVDS, Quad Speed High Serial (QHSS) from BAE, Custom over RS-422 and 1553/1773.

The criteria considered were as follows: Flight Heritage; Data Rate Sufficiency over 20m; Fault tolerant support; Power, Mass & Volume; Hardware only packet management; Hardware only data assurance; Expansion effort; Mitigation effort; Relative cost; Relative risk.

The protocol with the best score was SpaceWire.

4.5 Topology Trades

Once the SpaceWire was baseline, a number of trades were performed for the network topology.

4.5.1 Point-to-Point

A SpaceWire point-to-point configuration would have required 38 long length (~ 20m) SpaceWire cables for single fault tolerance. This was unacceptable considering the tight harness size requirements imposed by the spacecraft deployment mechanism & Thermal conduction.

4.5.2 Central Routers with Point-to-Point

By using the routing capability in the SpaceWire Specification the number of long cables could be greatly reduced by routing the image packets between the SCEMs and then sending the packets across fewer long cables to the ICDH. The first scheme that used this approach had a board concept for the SCEMs, where each SCEM would be a point-to-point link to a router board for that instruments Focal Plane Electronics (FPE) box. This architecture was reconsidered because of the different number of SCEMs per instrument made the 3 different instruments FPE vastly different in mechanical design.

4.5.2 Router Only

The final topology built upon the previous router scheme design by making each SCEM a router in a stackable enclosure. This way there would be one electrical and mechanical design for an instrument and each instrument would stack the number of SCEMs required for its array size (number of SCAs).

5. Overview Of SpaceWire Protocol

SpaceWire is a European Space Agency (ESA) Specification derived from IEEE-1355-1995. It defines scalable point-to-point links. Each point-to-point link that comprises the network fabric is full duplex using Data Strobe (DS) encoding over a Low Voltage Differential Signaling (LVDS) physical layer. The specification may downloaded from the ESA web-site http://www.estec.esa.nl/tech/spacewire/

5.1 Strengths

The SpaceWire protocol is the most versatile with respect to network topology of the buses considered. Because of the topology flexibility SpaceWire affords (it permits loops in the network) along with its switch fabric and scalability it allowed the best single fault tolerant scheme of the other considered protocols.

Other advantages are its high bandwidth; compact logic design (low power); simple user interface; very small buffer sizes (because of wormhole routing), very quick recovery from errors (20us), deterministic time distribution and protocol flexibility (send any packet structure across it).

5.2 Characters

The SpaceWire protocol defines a set of characters that are used to pass control information and data packets over a link. The two types of characters are control characters and data characters. The control characters are used to pass non-data information for protocol control.

The NULL packet (actually comprised of 2 control characters) is a control character that is used to synchronize the link. It is always sent when no other type of character is ready to be sent to maintain synchronization of the link.

Flow control characters provide control information between two links to prevent data loss by indicating to the peer transmitter the amount of room in the receive buffer. This prevents data from being transmitted when there is no room available to receive data. When this happens the link is "stalled" until receiver room becomes available and flow control indicates that data may be received.

There are two control characters that indicate end of packet. The End-Of-Packet (EOP) and End-Error-Packet (EEP) markers indicate the end of a good packet and end of a partial packet (error packet) respectively. These markers are then used to determine the beginning of the next packet.

Time code is comprised of a special sequence of a logical escape character followed by a data character. This sequence is used to broadcast time or synchronization information over the network fabric. It is designed so the broadcast message will only be received once at each router and terminate if the message loops back upon itself. This message may be interleaved with data characters from a packet, like all control character. This Time code will be used on the ISIM to distribute synchronization information accurate to within 2us.

Characters must be contiguous (no gaps) to maintain synchronization, and when no data or control information is required, NULL packets fill the link to maintain synchronization. All these characters may be interleaved with all other types of characters including data characters that encode a byte of data used for SpaceWire packets. The only restriction is that data characters from different packets may not be interleaved. All characters have an odd parity bit for error detection.

5.3 Overhead

The overhead associated with using SpaceWire link, neglecting the network and transport layer, is one extra byte for a destination address (in the case of logical addressing) and a 4 bits (logical character) for an end-of-packet marker, plus 25% overhead to encode each byte of data. This assumes that flow control information is not interleaved with packet data. This is an accurate assumption for the ISIM because of the little information flow in the reverse direction of the link (one way communication). This is because the image packets to the ICDH require a much higher bandwidth (110Mbps raw compared to 2Mbps to the SCEMs). The overhead for the image packets because of SpaceWire is 30%. This makes the 110Mbps raw data rate 143Mbps. Note: the engineering telemetry from the SCEM is missing from this number for simplicity but would increase the total data rate by 1.6Mbps.

5.4 Initialization

A simplified version of how a link is established is as follows. During start-up and initialization of the link, each side of the link will transmit NULL (synchronization) packets to each other. Upon a link receiving a NULL packet and after sending a NULL packet, each link will then transmit the amount of Flow control characters for their respective receive buffer size (each flow control character represents room for 8 data characters, max size of SpaceWire buffers is 56 bytes or 7 flow control characters). After transmitting and receiving Flow control credit, the link is able to transmit data. Data is transmitted by a SpaceWire packet format that defines the first data character as the destination address followed by 0 or more data characters and terminated by an end-of packer character. In this final state when data characters may be transmitted, all character types may be transmitted (necessary to maintain flow control and synchronization of the link).

5.5 Link Error Recovery

If an error occurs on a SpaceWire link, the link will perform a "silent protocol" and re-initialize. The "silent" protocol refers to the response of a link when it detects an error. It will go silent (stop transmitting). Once this happens the other side of the link will disconnect (stop transmitting) due to a lack of signaling. This begins the re-initialization process. The types of errors that may occur are parity error, character sequence error and disconnect error.

5.6 SpaceWire Router

A SpaceWire Router is a non-blocking cross bar switch that allows connection between a group of ports (3 ports to 30 ports). A port is defined as a SpaceWire Link. Connection may be made between any links input port to any other links output port. As long as any output port is available it may be used. If two or more link input ports are requesting the same link output port than fair arbitration is used, causing some links to stall and wait.

5.7 Wormhole Routing

SpaceWire implements a routing scheme called "wormhole" routing. It is easiest to describe by comparing it to a "store and forward" scheme. In the "store and forward" scheme, the whole packet must be buffered at a receiver before it may be passed to the next router in the network. In the "wormhole" routing scheme only the beginning of a packet must be received before being passed to the next router in the network. This reduces the latency on the network and requires very little buffering in the port receiver. In a "worm hole" scheme a packet may be "strung" across many different routers and may stall due to unavailability of an output port. This is fine because the flow control will prevent data loss (assuming sufficient buffering at source).

5.8 Routing Scheme

There are two routing scheme that SpaceWire may use.

5.9 Hardware Addressing

Hardware Addressing is defined as destination addresses from 1 to 30. These addresses do not use the look-up table in the router but are hard routed to the output port number in the Hardware address. For example, destination address 5 would be routed to port number 5. All hardware addresses are deleted upon leaving the router. This requires a concatenation of hardware addresses if a packet is to transverse multiple routers. Hardware addressing has large overhead.

5.10 Logical Addressing

Logical addressing uses the router's look-up table, and these addresses are defined for addresses in the range of 32 to 254 (255 is reserved). Upon a logical address entering a router it is looked up and mapped to a physical port number. Logical addresses may be deleted or not depending upon the information in the look-up table. This is a more bandwidth efficient method of routing packets if multiple routers are in the path.

5.11 Configuration 0 Space

Destination address 0 is reserved for configuration of the routing table port mapping and control of link speed, etc.

5.12 Packet Recovery from Error

If a packet that is being "wormed" across the network fabric is broken because of an error on a link, the following action is taken. The partial packet that has already been routed (downstream of the resetting link) to a link that is not in reset will be truncated with an Error End-of-Packet (EEP) marker, and continue to route its way across the network, as if nothing had happened (the destination will eventually discard the error packet). The portion of the packet that has not crossed the link because the link is being reset (upstream portion of the packet) will be consumed and not transmitted. Once the link is re-established, the beginning of a new packet will be the first byte sent.

5.13 External Port

Each router has at least one external port which to source and sink data to and from the network. The external port's address is one count higher than the number of ports on the router. For example, a four-port router would have its external port address to be 5 (Figure 3).

6. SPACEWIRE TRANSPORT LAYER

SpaceWire does not define a mechanism for reliable transport across the network, i.e., an acknowledgement/ retry scheme. For many systems reliable transport is necessary to meet mission objectives of observation efficiency, as is the case for JWST and/or to ensure proper receipt of commands. This transport layer will be implemented in hardware and resides between the external interface of the router or point-to-point link and the user interface that sources and sinks data to and from the SpaceWire network. This is an end-to-end mechanism and not implemented along the packets path at each router. This SpaceWire transport layer is being developed at Goddard for the JWST FPE network and it is being proposed and coordinated with ESA, the SpaceWire Working Group and CCSDS P1K SOIF panel. There are some very minor differences between the groups but the final solution is converging. The biggest contention is over the type 0 Packet Format proposed by JWST. JWST may standardize on the type 1 format allowing only one SpaceWire network type to be used by all. This collaboration is an effort to standardize this layer as an addition to the SpaceWire standard. There is a concerted effort to make the information fields added for the transport layer comparable to other transport layers like TCP/IP. It is also important that the connection transactions necessary to control a connection are compatible with existing transport layers, i.e., TCP/IP. This drastically reduces the implementation risk and enables seamless intra-network routing to occur.

6.1 Transparency

The transport layer works with SpaceWire logical addressing. In defining this layer a packet format has been created for packets that require reliable transport. This transport layer must work with designs that do not have this transport layer design. For those situations packets will pass transparently through the transport layer up to the higher layer. This is important for compatibility purposes.

6.2 Packet Format

The first byte must be the SpaceWire Destination address according to the existing SpaceWire specification. This scheme is designed to work for SpaceWire packet formats that have only one logical address (hardware addressing and logical addressing with header deletion are not allowed; these formats may be used but not with reliable transport). The second byte is the Source address. The third byte is the Packet Type that may define three different packet formats. The remaining format depends upon the first byte of the Packet Type byte, as described below.

6.3 Data Packet Format Type 0

When the first bit of the Packet type is 0, a Type 0 packet format is defined (Figure 4). This is the lowest overhead of the three data packet formats and was proposed by the ISIM team to preserve bandwidth. A variable length payload follows the Packet Type byte that is followed by the last byte of the packet, an optional 8-bit CRC. The CRC polynomial has not been determined at the time of this writing.

The Packet Type for Type 0 has following fields: first bit indicates Type 0, the second bit is the Acknowledgement bit, the third bit is the Segmentation bit, the next two bits indicate sequence number, and the last 3 bits indicate channel number.

Figure 4

Type 0 Data Packet Format

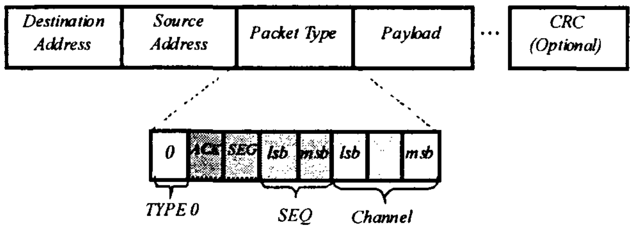

6.4 Data Packet Format Type 1

When the first bit of the Packet type is 1, a Type 1 or Type 3 packet format is defined (Figure 5). Type three's packet format is currently undefined for expandability purposes. Type one's format is defined and gives the user more packet sequence numbers (total of 16) and more channel numbers (total of 256) but incurs one more byte of overhead than a type 0 format (for ISIM this extra byte is undesirable).

Figure 5

Type 1 Data Packet Format

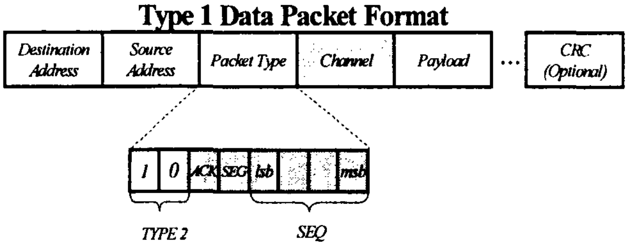

6.5 Packet Type Bit Definitions

6.5.1 Acknowledgment Bit

The Acknowledgment bit indicates if the packet is an acknowledgment packet or a normal packet.

6.5.2 Segmentation Bit (Last Packet Marker)

Segmentation bit indicates whether this packet is the last packet of the segment. Segmentation is not allowed for Channel 0 (described later).

6.5.3 Sequence Number Field

Packet sequence numbers are used to prevent out of sequence packets and to delete duplicate packets.

6.5.4 Channel Field

Indicates the virtual channel target for a packet. This field will be used in conjunction with the Source address to allow larger number of channels.

6.6 Channel Configuration

A channel must be established before any packet (data) transactions may occur. Channel 0 is reserved for establishing (and configuring) a connection. Channel 0 also contains a special pre-defined payload.

6.6.1 Channel 0 Payload

The Channel 0 packet always has a channel 0 payload that is used to configure and establish a channel. This may be thought of as the network control channel (Figure 6).

6.6.1.1 Channel Field

The Channel Field of the Channel 0 payload may not be 0. This byte field is used to identify the channel to be set-up.

6.6.1.2 Protocol Field

The Protocol byte identifies the protocol for the channel to be set-up (defined in the previous byte). This field is not important for the ISIM application but exists for generic purposes.

6.6.1.3 Source Port

This 16-bit Source Port field indicates the source port for the channel specified in the first byte of the payload. This field is not important for the ISIM application but exists for generic purposes.

6.6.1.4 Destination Port

This 16-bit Destination Port field indicates the Destination port for the channel specified in the first byte of the payload. This field is not important for the ISIM application but exists for generic purposes.

6.6.1.5 Options Field

This 16-bit Options field is currently undefined. It will probably contain the timeout value used to retransmit packets. Another possible option may include Issochrony.

6.6.1.6 Configuration Field

The configuration field and channel field are currently the only used fields of the Channel 0 payload. The other fields are either don't care or static information.

6.6.1.6.1 SYN Bit (Open)

Indicates channel setup, sequence bit synchronization and resets all receive state machines.

6.6.1.6.2 FIN Bit (Close)

The FIN bit is used to indicate channel closed. Currently this bit is not planned for use by ISIM because the concept of channel finished does not exist in the current design.

6.6.1.6.3 RST Bit (ERROR)

The RST bit is used to indicate fatal network errors or any errors that occur during connection. It is also used to flag a request for an unsupported connection. The application layer will handle these errors.

6.6.1.6.4 CRC Bit

The CRC bit is used to indicate that the channel specified in the first byte of the Channel 0 payload will have a CRC byte in the Data Packet Format. Note: CRC is always enabled for channel 0.

6.6.1.6.5 Reliability Field

The Reliability Field is a two-bit field to indicate if reliable transport will be enabled for the specified channel, and how many retries will be permitted. A 00 indicates unreliable (i.e., no acknowledges or retries). A 01 indicates reliable transport with 1 retry. A 10 indicates reliable transport with 2 retries. An 11 indicates reliable transport with 4 retries. Note: Configuration channels are fixed to 11 (reliable transport with 4 retries).

6.6.1.7 Maximum Transfer Unit (MTU)

This field is used to indicate to the initiator of the connection the maximum size (bytes) of the packets that may be sent over the specified channel. This defines the size of the retransmission buffer. This value will be fixed for the ISIM (details of this transaction are described later).

6.7 Channel Establishment

The channel establishment is performed by a series of messages transfers between channel Initiator and Destination. This "handshaking" is based upon TCP/IP protocol for channel establishment to facilitate bridging the SpaceWire transport layer to a TCP/IP transport layer, thus allowing use of existing commercial designs for test equipment, etc.

Note: active = '1'; inactive = '0'

To establish a channel the Initiator sends a channel 0 message. The Channel 0 payload has the SYN bit set active and the ACK, FIN and RST bits set inactive. The MTU value is set to the requested value for the channel specified in the Channel 0 payload.

If the destination wants to open the channel, it will acknowledge the initiator by echoing back the original channel 0 message (with the Source and Destination field swapped in the Data Packet Format, and if applicable the Source Port and Destination Port swapped in the Channel 0 payload) with the SYN bit set active, the ACK bit set active and the RST and FIN bit set inactive. The MTU value will be set to the actual value that the Destination can support.

At this point the connection is open from initiator to destination.

If the Initiator acknowledges the receipt of the Destinations acknowledge (by using a channel 0 message), than the channel will also be open in the reverse direction (from Destination to Initiator). This message has the Source and Destination fields swapped again in the Data Packet Format (and perhaps the Source Port and Destination Port swapped in the Channel 0 payload) with the SYN bit set inactive, the ACK bit set active, and the RST and FIN bits set inactive. The MTU value is the negotiated value (same for both Destination and Source).

At this point the channel is active and the connection tables are guaranteed to be consistent.

6.8 Channel Shunt down

The channel shut down does not make sense in the ISIM application (the channel will always be open) so it will not be used, but a description of the mechanism is included.

To shunt down a connection the process is similar to the establishment process. The Initiator sends a channel 0 message. The Channel 0 payload has the FIN bit set active) and the SYN, ACK, and RST bits set inactive. The Channel field in the Channel 0 payload indicates the channel to shunt down. The MTU value does not matter.

The Destination may respond with the FIN bit of the Channel 0 payload either active or inactive. If the Destination acknowledges (ACK bit active) with FIN bit set inactive (other bits; SYN, RST inactive) than the channel is half closed (from Source to Destination). If the Destination acknowledges (ACK bit active) with FIN bit set active (other bits; SYN, RST inactive) than the channel is closed in both directions.

Either way the Destination responses (FIN set active or inactive), if the Source acknowledges the Destination's acknowledge (ACK bit active) with the FIN bit set active (other bits; SYN, RST inactive) than the channel is fully closed.

6.9 Channel Reset

The reset of a channel represents a fatal error. This use for the ISIM is TBD. The Initiator of a Reset will send a Channel 0 message to the Destination with the RST bit set active and the SYN, ACK and FIN bits set inactive.

Upon receipt of the reset, the destination will notify the upper layer (User).

6.10 Sending Packets on Channel

6.10.1Normal Packets

After the channel is setup the initiator may send packets to the destination. This is accomplished by sending a message with the previously opened channel with the ACK bit reset and the CRC byte (if CRC was enabled in the channel establishment). The channel number will map to a destination source pair that was determined at channel establishment from the Data Packet Format (not from the Channel 0 payload).

6.10.2 Acknowledgement Packets

The Destination will acknowledge the message with a Data Packet format (type 0 for the ISIM). This acknowledgement packet will have the Destination address, the Source address, the Packet Type and an optional CRC byte, depending upon the establishment configuration. The ACK bit will be set in the Packet Type byte, and the sequence number will correspond to the sequence number of the packet being acknowledged. The Segment bit always indicates no segmentation.

6.10.3 Timeout and Retransmission

After a normal packet is sent (non-acknowledgement packet), a timeout counter is started. If an acknowledgement is not received before the counter expires than the normal packet will be retransmitted. The timeout counter will be specified in the optional field in the Channel 0 payload upon channel establishment. The timeout counter for the ISIM will be less than 40us (the packet generation rate). The number of retries for the ISIM will be one making it a simple send and wait approach, i.e., every packet that uses this service must be acknowledged before a new different packet may be sent. If the timeout happens a second time then the User interface will be signaled that a failure occurred and the retransmission buffer will be cleared waiting for the channel to be reset

The ISIM reliable transport logic will likely have a 128 byte retransmission buffer (defined by the MTU field in channel 0 payload), based on Actel AX memory block size that defines the maximum size of a packet that may use the service. Packets larger than 128 bytes, which want to use reliable transport must be segmented at the User Layer and reconstructed at the destinations User Layer.

6.10 Overhead

The overhead to the 32-byte raw image packet format by SpaceWire was 30%. The transport layer will add an additional 2 bytes if CRC is not used (over the one byte that SpaceWire requires). This increases the overhead for the image packets to 38% for no CRC (increases the link rate to 152Mbps). If CRC were used the total overhead would be 42% (increases the link rate to 156Mbps).

7. USER LAYER

7.1Different Types of Users

As mentioned earlier, there are two different types of users in the FPE Network. The SCEM and the ICDH: Each has a different User interface.

7.2 ICDH Reliable Transport

The ICDH User interface performs segmentation of larger packet into smaller SpaceWire packets (64 bytes or less; size of retransmission buffer) to be transmitted over a reliable channel. It uses the segmentation bit to indicate how to reconstruct the larger packets at the destination SCEM. These packets are command packets to the SCEM for configuring image and telemetry collection.

7.3 ICDH Unreliable Transport

The ICDH also is responsible for configuring the SpaceWire router tables that are contained in the SCEMs. These packets do not use the reliable transport mechanism, because they use SpaceWire hardware addressing that is not supported by the transport layer. The SpaceWire routing table status information is passed through the transport layer and verified correct at the application layer.

7.4 SCEM Reliable Transport

The SCEM generates image data and telemetry data. They will use reliable transport and the packets will be segmented into SpaceWire packets 128 bytes or less and reconstructed back at the ICDH. ISIM image packets will probably not use the CRC because of the overhead.

8. IMPLEMENATION

The target implementation for the SpaceWire design including the reliable transport is an Actel AX1000 FPGA. It is expected that this design will be embedded with the user interface logic of the SCEM or ICDH Bus interface Card (BIC). The design will have 2 clock domains, one for the highest transmit frequency and one one-quarter less for the core logic. The goal is to run the high frequency at 160MHz.

Figure 1: ISIM Block Diagram

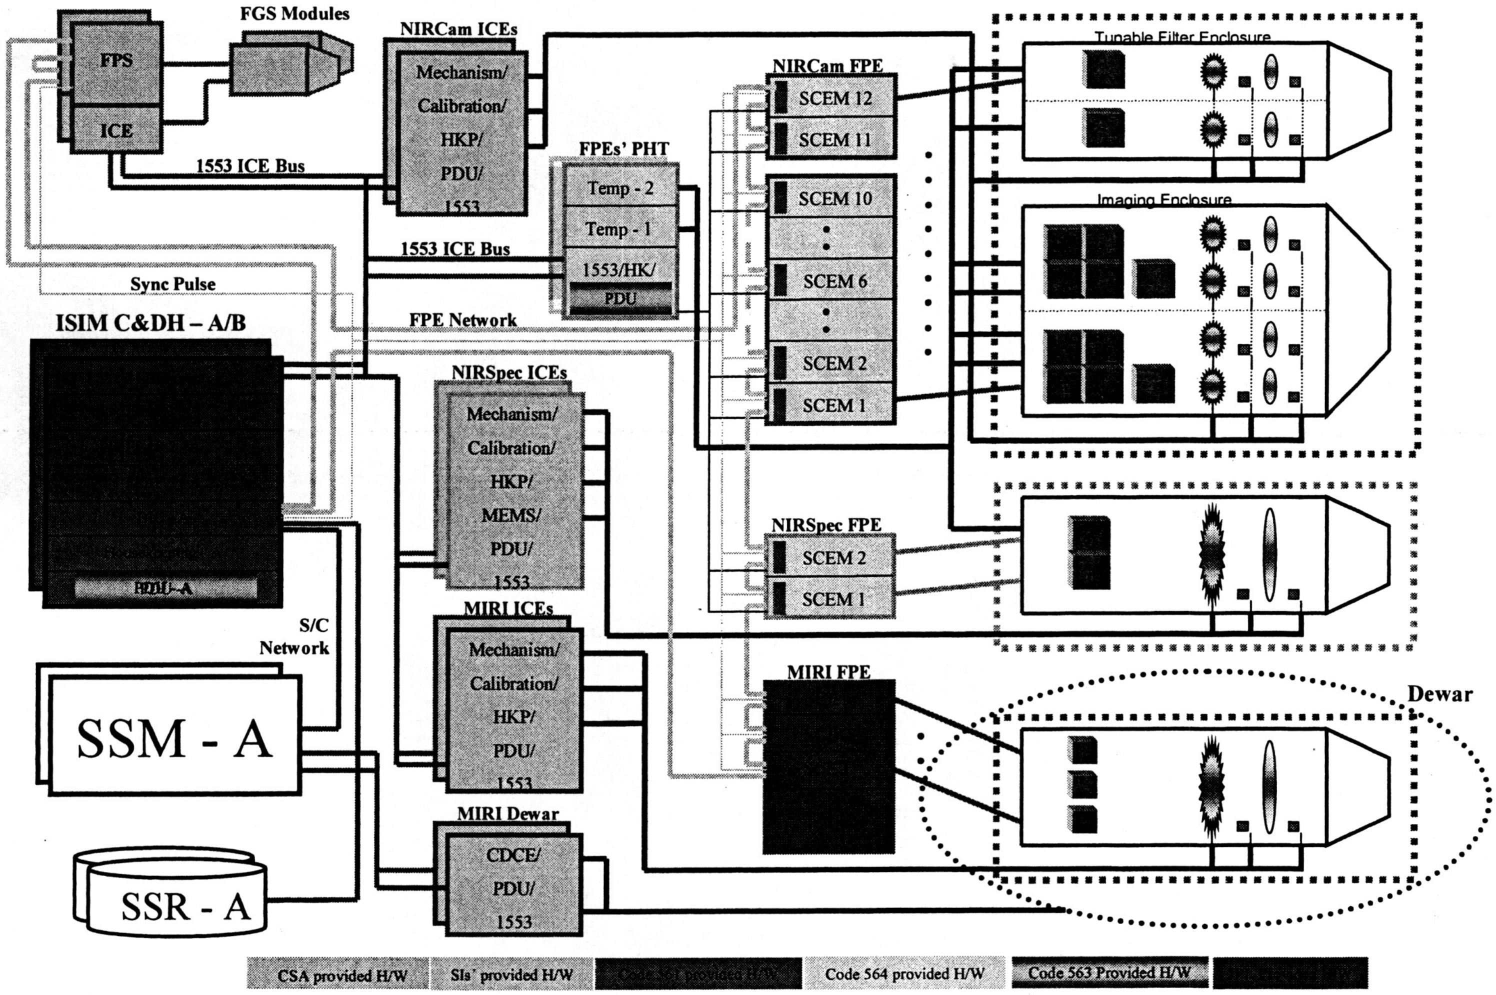

Figure 2: JWST ISIM Network Topology

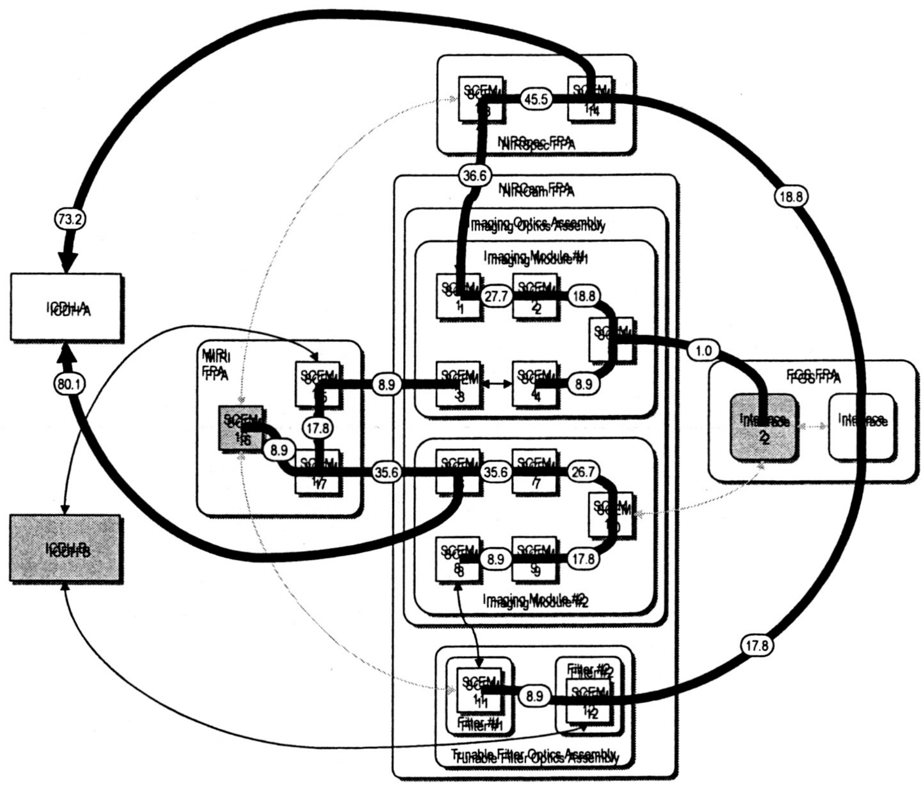

Figure 3: SCEM 3 Port Router

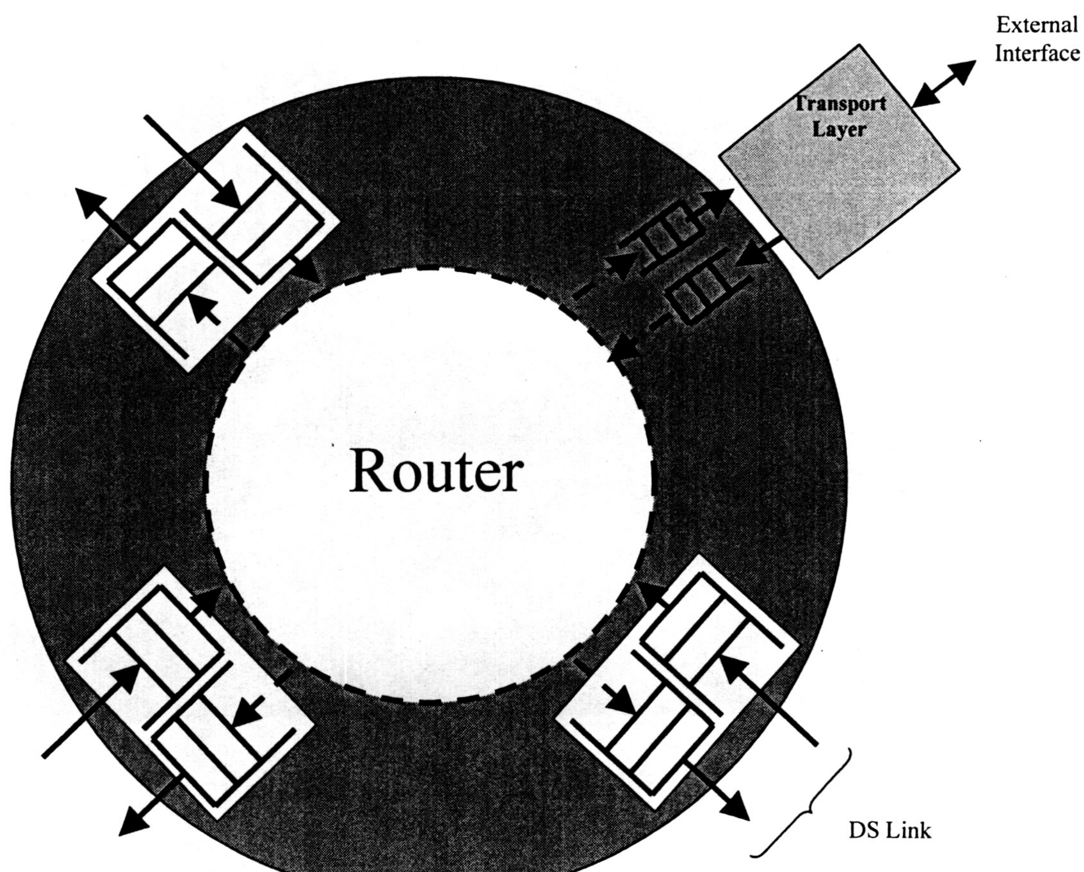

Figure 6: Channel 0 Payload

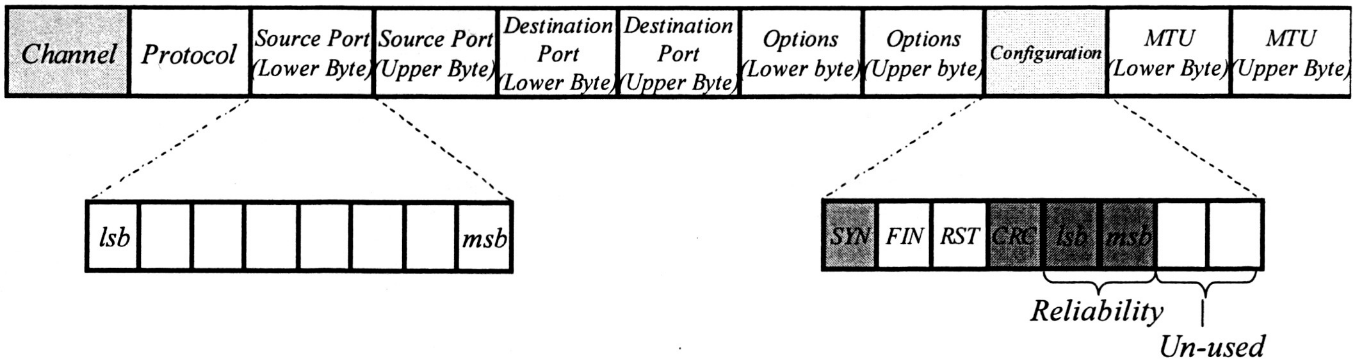

9. Conclusion

The SpaceWire network is a crucial component of the ISIM instrument and future success of JWST. It is a simple elegant solution for high bandwidth distributed senor systems, allowing small compact low power electronics at the sensors. The topologies possible with SpaceWire network fabrics makes designing fault tolerant system easier than most other protocols, while minimizing the cost of adding this redundancy. The reliable transport logic will make the SpaceWire network robust without compromising the SpaceWire specification. The Transport Layer of SpaceWire is still being refined (at time of writing), but the basic concept has been agreed upon by the major entities, ESA, SpaceWire Working Group, CCSDS P1K SOIF. At the writing of this paper the JWST project was undergoing a review to cut cost of the development effort. Broad studies are in process to find ways to cut several hundred million dollars from the development effort of the entire observatory. The ISIM as it was presented may be radically different in the spring of 2003.

10. Acronym List

AX	Accelerator Series
BIC	Bus Interface Card
C	Celsius
CCSDS	Consultative Committee for Space Data Systems
C&DH	Command & Data Handling
CRC	Cyclic Redundancy Code
DS	Data Strobe
EOP	End-of-Packet
EEP	End Error Packet
ESA	European Space Agency
FGS	Fine Guidance Sensor
FIFO	First In First Out
FPA	Focal Plane Assembly
FPAP	Focal Plan Assembly Processor
FPE	Focal Plan Electronics
FPGA	Field Programmable Gate Array
GSFC	Goddard Space Flight Center
HK	Housekeeping Card
HKP	Housekeeping
ICDH	ISIM Command & Data Handling
ICE	Instrument Control Electronics
IEEE	Institute of Electrical and Electronics Engineers
IP	Intellectual Property
ISIM	Integrated Science Instrument Module
JPL	Jet Propulsion Laboratory
JWST	James Webb Space Telescope
LVDS	Low Voltage Differential Signaling
m	meter
Mbps	Megabit per second
MEMS	Microelectronic Mechanical Systems
MHz	Megahertz
MIRI	Mid Infrared Instrument
NASA	National Aeronautics and Space Administration
NGST	Next Generation Space Telescope
NIRCam	Near Infrared Camera
NIRSpec	Near Infrared Spectrometer
PDU	Power Distribution Unit
SBC	Single Board Computer
SCE	Sensor Chip Electronics
SCEM	Sensor Chip Electronics Module
SSM	Spacecraft Support Module
SSR	Solid State Recorder
um	micrometer
UTMC	United Technologies Microelectronics Center
VHDL	VHSIC Hardware Description Language

9. BIOGRAPHY

Glenn Rakow is an electrical engineer at NASA/Goddard Space Flight Center's Flight Data Systems & Radiation Effects Branch. He is currently the lead SpaceWire Network Development Engineer for the JWST ISIM. Since 1998 he has been an advocate for SpaceWire, promoting SpaceWire in the US satellite community, and designing to the standard. He designed the first US SpaceWire Link RTL VHDL IP Core that was used on the NASA Swift Mission in the Burst Alert Telescope (BAT). This design was targeted to two different UTMC 0.6um Rad Hard Gate Arrays. In his current job he is leading the effort to realize the ISIM SpaceWire network that includes the router and a transport layer not defined in the SpaceWire Specification. Before this Mr. Rakow worked as a FPGA design engineer for TRIANA and as a lead Power Electronics Engineer on a series of SMEX mission SWAS, TRACE and WIRE. When he first came to NASA in 1988 he worked on the first SMEX mission SAMPEX designing a power distribution box. He earned a Bachelor of Science in Electrical Engineering from the University of Maryland, Collage Park, in 1988, and a Master of Science in Electrical Engineering from George Washington University, Washington D.C., in 1997.

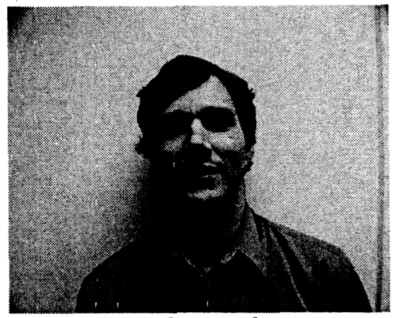

Richard Schnurr is the GSFC representative to the CCSDS Standard On board Interfaces (SOIF) panel as well as the Chief Architect of the Electrical Engineering division. In these roles Mr. Schnurr is working with many missions in early formulation to design standard on-board networks meeting specific mission requirements. Mr. Schnurr has worked at GSFC 19 year. He graduated from the University of Maryland with a BS in Electrical engineering in 1984. Rick worked on JWST as the IC&DH study lead during formulation.

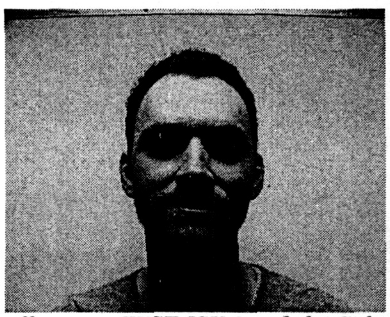

Christopher Dailey is a computer engineer at NASA Goddard Space Flight Center (GSFC). He is currently working the SpaceWire Network development for JWST ISIM, as will as other digital design efforts on JWST ISIM and the Solar Dynamics Observatory (SDO), another satellite program being developed at NASA/GSFC. He has been the lead, co-lead and/or designer of several boards, FPGAs and ASICs for NASA/GSFC's space flight and ground systems including the EOS Terra, Aqua & Aura and EO-1 satellites as well as the department of Defense and private industry. He earned a Bachelor of Science in Computer Engineering with a Minor in Computer Science from the Pennsylvania State University in 1991 and was an intern at NASA/GSFC in 1988.

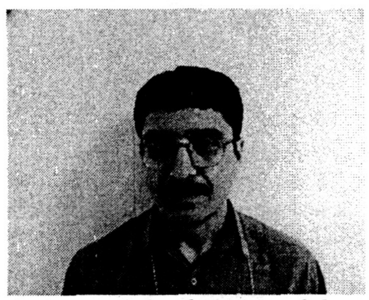

Kamdin Shakoorzadeh is a senior engineer at QSS Group Inc. supporting NASA/Goddard Space Flight Center's Flight Data Systems & Microelectronics Branches. He is currently supporting the design and development of the ICDH and FPE electronics for the JWST ISIM project. For the past 19 years, he has been supporting various GSFC programs such as the Infra-Red Array Camera (IRAC) Instrument, the Advance Hitchhiker Avionics (ACE), The Microwave Anisotropy Probe (MAP) Instrument, the X-Ray Timing Explorer (XTE), the Extreme Ultraviolet Explorer (EUVE), and the Shuttle payloads, at various levels of the development from conceptual design and systems engineering, through electronics design, simulation, flight software development, and integration and test. He has a BSME from Lake Superior State University, and a BSEE and MSCS from George Mason University, and has completed various courses on Spacecraft Systems Design and Engineering, at Applied Technology Institute