

International Conference on Space Optics—ICSO 2014

La Caleta, Tenerife, Canary Islands

7–10 October 2014

Edited by Zoran Sodnik, Bruno Cugny, and Nikos Karafolas



Satellite based optical backplanes and the next generation space interconnect standard (NGSIS): a modular open standards approach for high performance interconnects for space

C. P. Collier



International Conference on Space Optics — ICSO 2014, edited by Zoran Sodnik, Nikos Karafolas, Bruno Cugny, Proc. of SPIE Vol. 10563, 1056334 · © 2014 ESA and CNES
CCC code: 0277-786X/17/\$18 · doi: 10.1117/12.2304084

SATELLITE BASED OPTICAL BACKPLANES AND THE NEXT GENERATION SPACE INTERCONNECT STANDARD (NGSIS): A MODULAR OPEN STANDARDS APPROACH FOR HIGH PERFORMANCE INTERCONNECTS FOR SPACE

C. P. Collier

Space Vehicles Directorate, Air Force Research Laboratories, Kirtland AFB, USA

I. INTRODUCTION:

This paper describes the Next Generation Space Interconnect Standard (NGSIS) effort for the Reinventing Space community. This description covers the goals and objectives of the NGSIS effort, the composition of the NGSIS participants, and the approach and roadmap being used for the NGSIS effort and the current status. NGSIS is a Government-Industry collaboration effort to define a set of standards for interconnects between space system components with the goal of cost effectively removing bandwidth as a constraint for future space systems. Initial emphasis is on the standardization of internal connectivity at the electronic chassis level. This includes the needs for high reliability but limited rate data needs typical for spacecraft command and data handling (C&DH) as well as high rate data needs expected for next generation, high performance mission payloads.

The architectural approach adopted for NGSIS development was to select several appropriate and proven industry standards as "points of departure" and then develop a set of extensions to these standards to address space industry specific needs. The intent of this approach is reduced cost, risk and effort by using proven technologies while providing a set of common extensions that can be adopted across the space industry to enable interoperability of board level components from different sources and vendors. The NGSIS team has selected the VITA OpenVPX standard family for the physical baseline. VPX supports both 3U and 6U form factors with ruggedized and conduction cooled features suitable for use in extreme environments. Current efforts to develop the space specific extensions are being worked by establishment of an NGSIS Working Group (VITA 78) under the auspices of the VITA Standards Organization (VSO), an accredited standards development organization.

The Serial RapidIO protocol has been selected as the basis for the digital data transport. The Serial RapidIO protocol uses an efficient, high performance packet switching architecture to provide an interconnect capability suitable for chip-to-chip and board-to-board communications. Data rates are scalable per lane up to a current maximum of 6.25 Gbps (gross bit rate) and allow aggregations of lanes into channels with capacities in excess of 60 Gbps, with a roadmap to higher rates in the future. Current efforts to develop the space specific extensions for improved reliability, robustness and additional feature desirable for space systems are being worked by establishment of the "Part S" task group by the RapidIO Trade Association as a collaboration effort with the NGSIS participants.

II. HISTORY OF NGSIS:

The NGSIS effort initiated in 2010 when a Technologist at the AFRL Space Vehicles Directorate performed an outreach effort within the US Government space community to determine interest in higher speed interconnects for space with principal interest in photonics and optical technologies. Several organizations in NASA, the Air Force, and other parts of DoD indicated interest and expected future needs for data transport at rates significantly greater (and with a higher degree of scalability and flexibility) than that provided by the state of industry.

Based on these results, an initial organizing meeting to garner additional interest and participation was held at the GOMACTech 2010 conference. Several other Government and Industry organizations had also recognized similar internal needs for higher speed data transport and recognizing the potential benefits from a standardized industry approach, expressed interest in participating in the fledgling standards efforts. Patrick Collier from the AFRL Space Vehicles Directorate and Raphael Some from JPL Autonomous Systems Group volunteered to organize the initial effort and recruited several senior executives from within DOD and NASA space community to serve as an Executive Steering Committee.

Several working groups/committees were organized to address Requirements, Specification, Protocol, and Fault Management.

A regularly scheduled set of teleconferences were arranged and several additional technical outreach and interchange meetings were held at other venues such as The Sandia High Performance Computing for Space Conference and the Reconfiguring Space/MAPLD conferences to refine the objectives and recruit additional

participation. Sufficient interest materialized across the community to enable volunteer staffing of the established working groups.

Several members of the participating organization presented results from recent internal studies that addressed high speed data transport architectures. There was a high degree of commonality in the nature of the studies and associated results in comparing existing high speed interconnect architectures, such as Serial Rapid IO, 1 Gigabit Ethernet, 10 Gigabit Ethernet, PCI Express and others..

Based on these overlapping results, the consensus decision from the participants were to:

- Pursue a structured and transparent system architecture and engineering process using IEEE-147/ISO-42010 as a roadmap.
- Pursue an electrical solution rather than photonic for the physical layer
- Pursue a VPX like architecture for the physical architecture
- Pursue Serial Rapid IO as the major protocol of interest

In order to better facilitate the development of the emerging standards and to help garner additional advocacy and support, the parent bodies for the standards of interest were engaged on the topic of supporting space specific extensions of their respective standards. As a result, the VMEbus Industry Trade Association (VITA) Standards Organization (VSO), the parent body for the VPX family of standards created the VITA 78 SpaceVPX working group. And the RapidIO Trade Association, the parent body for the SRIO standard created the Part-S working group.

III. SPACEVPX

SpaceVPX or VITA 78 is a new standard being developed under the VITA Standards Organization. It builds on the VITA 65, OpenVPX™, and extends it for space applications. An example use case for SpaceVPX is shown in Figure 1. Data from sensors are represented by data in blocks. Data is switched between inputs, mass storage, processing and output blocks. Downlinks represent data being sent to the ground. Note the full redundancy present among the elements; this in part satisfies the need for additional fault tolerance in space systems. A second topology fully supported by SpaceVPX involves a peer to peer mesh between all elements instead of a switch. Combinations of the two topologies are also easily assembled using the standard.

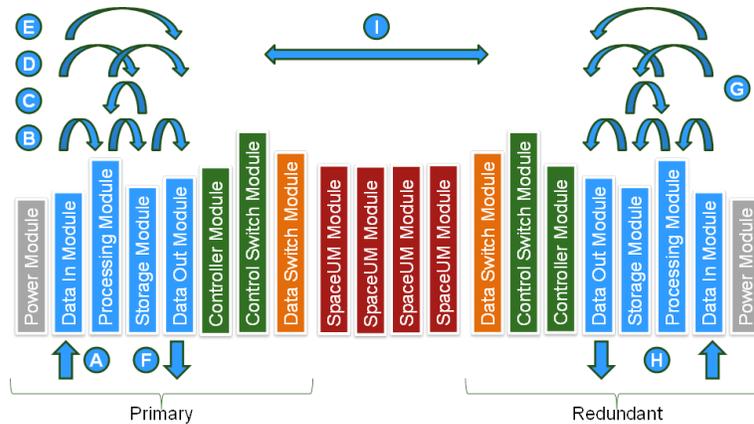


Fig.1. SpaceVPX System Use Case

A. OpenVPX Standards

SpaceVPX builds upon several standards that are part of the ANSI/VITA OpenVPX family. These include the base VITA 46 VPX standard and its ANSI/VITA 65 OpenVPX derivative. SpaceVPX also allows other compatible connectors to be used, including ANSI/VITA 60 and 63. ANSI/VITA 48 forms the base of the mechanical extensions in SpaceVPX. ANSI/VITA 66 and 67 also may be applied to replace electrical segments of the connector with RF or optical. ANSI/VITA 46.11, still under development, provides a base of the management protocol that SpaceVPX builds upon for fault tolerant management of the SpaceVPX system.

Five major interconnect planes organize the connections in OpenVPX. The data plane provides high speed multi-gigabit fabric connections between modules and typically carries payload and mission data. The control plane, also a fabric, typically has less capacity and is used for configuration, setup, diagnostics and other operational control functions within the payload. The management plane's function is providing setup and control of the basic modules functions typically having to do with power sequencing and low level diagnostics. The utility plane contains the power, clocks and other base signals needed for system operation. The expansion plane may be used as a separate connection between modules utilizing similar or heritage interfaces in a more

limited topology such as a bus or ring. Pins not defined as part of any of these planes are typically user defined or available for pass through from daughter or mezzanine cards or to rear transition modules (RTM). ANSI/VITA 65 should be consulted for more detail on these structures.

B. Major Changes in SpaceVPX

In evaluating the use of OpenVPX for potential space usage, several shortcomings were noticed. The biggest one was the lack of features that could support a full single fault tolerant and highly reliable configuration. Utility and Management signals were bussed and in most cases only supported on a set of signals, via signal pins to a module. A pure OpenVPX system has opportunities for multiple failures as a result. A full management control mechanism was also not fully defined with VITA 46.11 still in development. The fact that the typical OpenVPX control planes are PCI Express or Ethernet, was another shortcoming of concern since SpaceWire is the dominant medium speed data and control plane interface for most spacecraft. A third area was the desire to reuse the infrastructure of OpenVPX for prototyping and testing SpaceVPX on the ground.

C. Fault Tolerance

To provide the fault tolerance needed, utility and management signals needed to be dual-redundant and then switched to each SpaceVPX card function. A trade study was performed early to compare between various implementations including adding the switching to each card in various ways as well as creating a unique switching card. The latter approach was chosen so that SpaceVPX cards could each receive the same Utility and Management plane signals that an OpenVPX card receives with minor adjustments for any changes in topology. This became known as the Space Utility Management (SpaceUM) module and is a major section of the standard. The SpaceUM module contains up to eight sets of power and signal switches to support SpaceVPX modules. It receives one power bus from each of two Power supplies and one set of management signals from each of two System Controller functions required in the SpaceVPX backplane. The various parts of the SpaceUM module are considered extensions of the Power Supply, System Controller or other SpaceVPX modules for reliability calculation and thus do not require their own redundancy. Two Management protocol options are provided for control of the system; one is a subset of VITA 46.11; the other is a simpler protocol developed specifically for SpaceVPX. Both utilize the management plane for access to the managed modules.

A slot profile for a controller for Switch topology 5 is shown in Figure 3. P0/J0 through P6/J6 represents the segments on the SpaceVPX connector. Each segment has either 8 or 16 wafers that contain up to 9 backplane pins. The figure shows how the various connections are mapped to the slot. The standard defines each pin to insure interoperability between modules.

D. SpaceVPX Profiles

A total of 17 backplane profiles were defined to cover the spectrum of potential payload topologies expected in SpaceVPX usage. The first nine cover switched Backplane profiles. The first three use a single fat pipe (4 lanes of data plane fabric) as the connection to the switch, the next three use a double fat pipe (8 lanes of data plane fabric) and the last three use a quad fat pipe (16 lanes of data plane fabric). These are followed by 7 mesh topologies, two each with 1, 2 or 4 fat pipes between each peer and the last one with a special integrated power supply. Each of the groups of three or two are then separated by the location and grouping of the controller and data plane switches. The first of each switch topology group has an integrated data plane switch on the same card as the controller. The second has a separate controller but with connections to the data plane and the third of each group has a separate control plane without a data plane connection. The mesh topologies have a mesh profile with an integrated controller on a mesh peer and a separate profile with a controller without a mesh connection. The profiles are defined for the maximum number of slots possible for the data switch or controller implementation within the topology. Lower numbers of slots are possible by not implementing the additional nodes.

A payload slot profile for a controller for Switch topology 5 is shown in Figure 2. P0/J0 through P6/J6 represents the segments on the SpaceVPX connector. Each segment has either 8 or 16 wafers that contain up to 9 backplane pins. The figure shows how the various connections are mapped to the slot. The standard defines each pin to insure interoperability between modules.

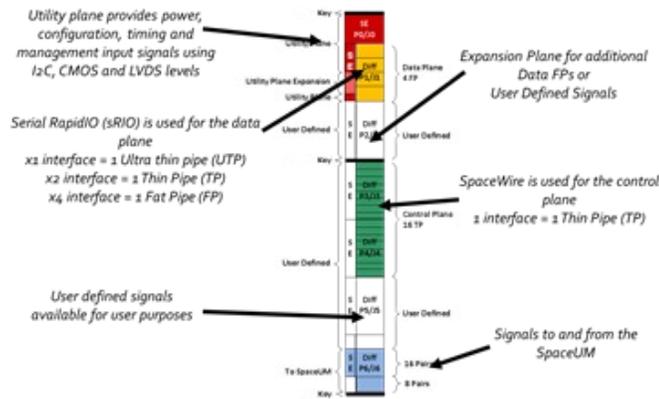


Fig.2. Example Slot Profile with Connection Notations

E. SpaceVPX Specification Progress

After six months of studies and trades and four months of focused standard writing and development, an initial draft of the SpaceVPX standard was created in April, 2013. At that point the VITA 78 study group became a working group and since then, the VITA 78 working group has been reviewing, updating and working the document toward an expected first vote in late 2013. As of July of 2014, the SpaceVPX specification was approved for release at the working group level. The working group will now focus on American National Standards Institute (ANSI) approval. SpaceVPX modules are expected to emerge from development starting in late 2014 to early 2015.

IV. AFRL PHOTONICS RESEARCH AND DEVELOPMENT

Current and future satellite payloads will generate amounts of data that will require copper data communications infrastructure to increase in size, weight, power, and cost. Requirements for future spacecraft subsystem interconnect are growing by orders of magnitude. In the past 5 years, we have seen the industry go from 1553 at 1Mb/s to SpaceWire at 250Mb/s to Time Triggered Gigabit Ethernet at 1Gb/s for Control Plane applications. Current high data rate needs require data rates in excess of 5Gb/s with multiple lanes of traffic operating simultaneously. New systems paradigms, including plug and play architectures and vehicle undoc/re-doc, will require extending interconnect capabilities and reliability levels well beyond those available in current technologies. Figure 3, below, provides an estimate of sensor data generation growth.

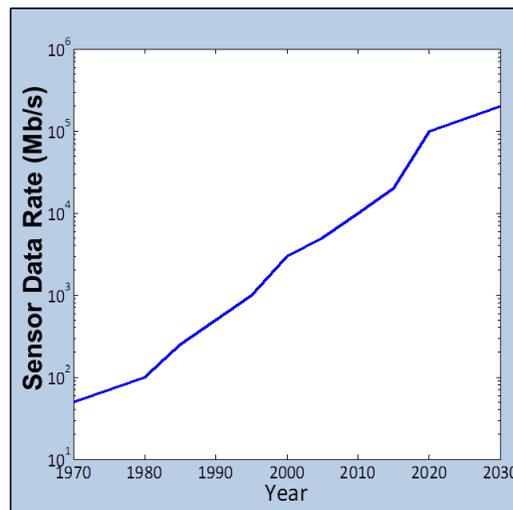


Fig.3. Sensor Data Rate Growth

The Satellite Optical Network (SON), as part of the Space Communication Program's overall intra-satellite communication strategy, is focused on verifying the viability of optical interconnects (cabled or free-space) over their copper counterparts. Optical interconnects offer an increase in bandwidth, scalability through the use of the same infrastructure with a variety of bandwidth needs, a decrease in cost, size, weight, and power (CSWaP)

in comparison to the copper analog, and non-existent Electromagnetic Interference (EMI) (leakage that affects other adjacent transmission signals).

The physical topology (layout) of the network is a Star, i.e. there is a central hub with connections points to all end-points. Each spoke of the central hub has one connection to one end point. End-point to end-point communication must go through the central hub.

The SON test platform, shown in figure 4 and 5, consists of 4 commercial-off-the-shelf (COTS) computers, 4 COTS Field Programmable Gate Array (FPGA) boards, 2 COTS solid state drives (SSDs), 2 government furnished cameras, 1 government furnished optical switch and 1 government furnished control node. One computer and FPGA board represent one satellite processor node. Each of these processor nodes is connected to the optical switch and the optical switch can be configured for a specific network topology through its Ethernet control interface.

The FPGA board is reconfigurable and is used to implement a specific network protocol to transport data over the optical network. The cameras provide test platform stimulus where their image data that will be encapsulated into payload packets by the FPGA board's implemented network protocol. The SSD provide storage for camera image data from the network or directly from the camera. Images can be retrieved from the SSD and encapsulated into payload packets and transported over the network by the FPGA board.

The SON test platform emulates a satellite with 2 camera processing nodes and 2 SSD processing node. Each processor node is connected in a star topology and each FPGA board is configured using the Xilinx Aurora network protocol. The control node configures the optical switch for point-to-point communication between processor nodes. The Payload represents the camera image data that is encapsulated by the Aurora protocol and transported over the network by the FPGA board.

The Polatis Optical Switch is used to route Network Protocol video image payloads to processor nodes over the high-speed data plane. The Optical Switch is an 8-input to 8-output passive optical cross connect, only 4 inputs/outputs are used to connect processor nodes. The Optical Switch only accepts single-mode fiber with FC connectors on its input/outputs. The Optical Switch uses 10Mb/s Ethernet control interface to switch input to output. The 10Mb/s Ethernet connects to the SON test bed Control Node, this link serves as the low-speed control plane (See Figure 4).

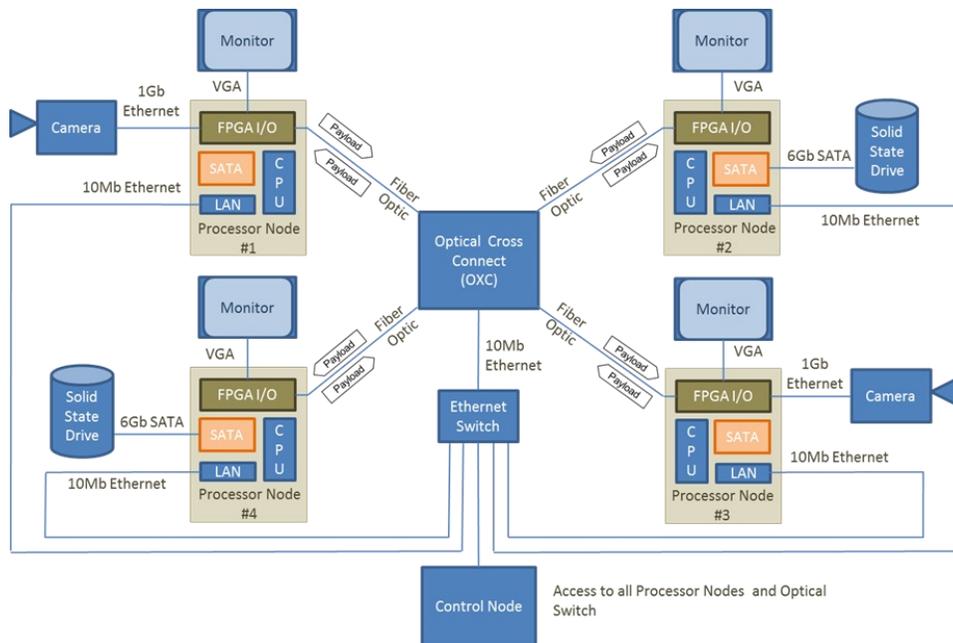


Fig.4. Satellite Optical Network (SON) switched architecture

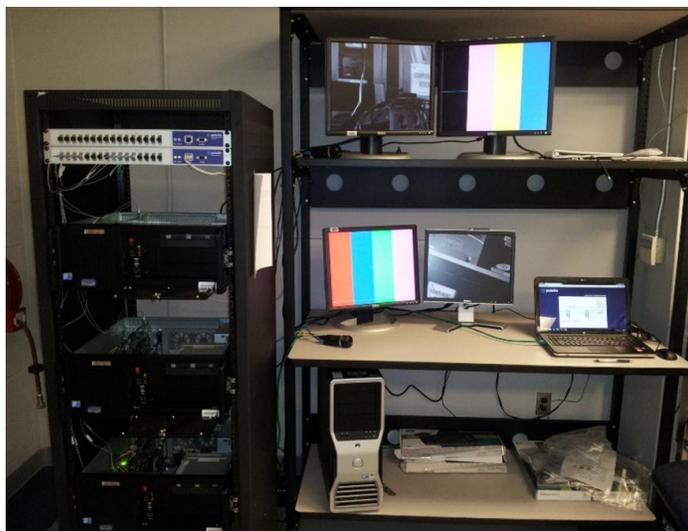


Fig.5. Satellite Optical Network (SON) switched architecture

V. SATELLITE OPTICAL NETWORK (SON) ACCOMPLISHMENTS

- **Developed typical “communication satellite” test platform consisting of centralized optically switched bus and communication payload**
- **Initial protocol testing leverage Xilinx’s Aurora protocol**
 - Encapsulated payload into packets using aurora communication protocol and transferred packets over optical bus (Star Topology)
 - Connections must be actively modified via a control plane interface to the optical cross-connect
- **Transferred payload packets through STAR point-to-point optical bus topology**
- **Used control plane to switch point-to-point processor node connections to demonstrate packet routing**
- **Remotely controlled payload packet routing through optical bus using the control plane to switch point-to-point processor node connections**
- **Remotely controlled processor node solid state drive recording and playback of camera video using distributive data service software**
- **Software driver developed to capture / transmit data to / from the network at high speeds**
- Demonstrated raw throughput of 10 Gb/s
- Demonstrated “real-world” throughput* of ~8 Gb/s
- Hardware interface modifications necessary to achieve faster throughput speeds
- Network speeds of 2.67 Gbps in hardware with upgrades to COTS optical transceiver utilizing SFP form factor, or 10Gb/s with SFP+ form factor or LM hardware
- Leverages COTS protocols on the Ethernet control plane
- **Software development tool-chain for the VPX was confirmed; simple “Hello World” test executable binary built and transferred to system and tested successfully**
- **Communication between the SBCs via the message passing interface (Part 2) is partially confirmed; test information was received on the remote computer and bounced back to initiator, but failed validation check.**

REFERENCES

- [1] IEEE Standard 1471-2000, “IEEE Recommended Practice for Architectural Description of Software-Intensive Systems”, Software Engineering Standards Committee, Approved 21, September 2000.
- [2] VITA 65 (OpenVPX) Standard, v3.00, September 2013
- [3] VITA 48.2, REDI Conduction Cooling Applied to VITA 46, v0.17, April 2010.
- [4] Serial RapidIO Specification Revision 2.2, May 2011.