Optical Networking for Aerospace Systems Provisioned through Plug-and-play Avionics

Keith Avery, Nathaniel (Shane) Francis, James Lyke
Air Force Research Laboratory
3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117

Patrick Collier
SAIC
2109 Air Park Rd SE Albuquerque, NM 87113

ABSTRACT

A high performance extension of the space plug-and-play avionics (SPA) technology is described based on the addition of optical fiber interconnection. This optical SPA (“SPA-O”), due to the bandwidth potential, provides the possibility of moving unprecedented amounts of data and rendering super-computing systems rapidly in orbital platforms. By exploiting the ideas of SPA, the optical properties of SPA-O components can be added to electronic datasheets embedded within the component so that the system-at-large can discover these components and exploit them in dynamically provisionable optical networks formed on demand.

INTRODUCTION

The space plug and play avionics (SPA) architecture [1] consists of a set of technologies focused on shortening the time it takes to construct a complex satellite system. In this architecture, systems are modularized as smart building blocks, analogous to the components of a personal computer, that when assembled, self-organize to form a desired full system. Software and hardware alike are modularized, abstracted individually into an electronic description referred to as extensible Transducer Electronics Datasheet (xTEDS), which is embedded in the hardware component or the source code of the software

connecting the features described in the data sheets to the components that require them. It creates a sort of electronic brokerage that eventually links together the dozens or hundreds of data dependencies within a system resulting in the “magical” effect that is sometimes called “plug-and-play”.

To the detached observer, a system appears as a set of black boxes that emit signals over the single SPA links that emanate from them. The observer may find in examining many complex systems, such as spacecraft, aircraft, and ground systems, a distribution or spectrum of bandwidths associated with their many components. A complex system will typically consist of many very simple components. There will be components that require more signaling bandwidth, but not as many of them. At each identified tier of performance, this relationship repeats, with the finding that above a performance level x , a number of components will require a bandwidth greater than x , but fewer than those requiring less than x . This relationship amounts to a sort of power law distribution (the so-called “scale-free network” described in [2]), and it is roughly captured in Figure 1 in a pyramid diagram, in which the height of the pyramid represents bandwidth demand, and the girth represents the relative quantity of the devices in a system that require that demand.

In the parlance of “SPA”, a SPA- x interface refers to a connector that would be found on a SPA device allowing the device to join a plug-and-play (SPA) network. The suffix letter refers to a class of interface, driven usually by bandwidth. In the history of SPA, the

Figure 1. The use of SPA across the bandwidth spectrum.

component. A software framework, referred to as “satellite data model” (SDM) discovers the xTEDS of components when they are added to the system,

interfaces were actually not developed in the order depicted in Figure 1. SPA-U, based on the USB standard, was the first SPA interface technology developed [3]. This was because USB already embodied the essential features of plug-and-play (in the personal computer), and the other interface approaches largely did not exhibit plug-and-play characteristics. To access higher performance levels, the SPA-S interface technology was created based on the European Spacewire standard [4], building upon the knowledge gained in SPA-U to add the needed plug-and-play protocols. To improve the efficiency of SPA for extremely simple devices, a development program for SPA-1 has been recently commissioned [5]. The final, important tier of performance is addressed in this paper.

MOTIVATION AND VISION FOR SPA-O

It is apparent that, as with terrestrial systems, space systems are experiencing an insatiable demand for bandwidth, and the growth characteristics will drive the interior components of the spacecraft and payload towards multi-Gbps bandwidths. Under the best circumstances, SPA-S is limited by the skew of the clock and data signals (which are transported on separate copper connections) of the Spacewire physical layer. As such, a high-performance camera (say a 4Gbps “device”) will require ten parallel SPA-S links (assuming a 400Mbps transport, close to the upper limits that have ever been demonstrated with Spacewire – practical implementations are typically more modest in performance). Even then, some exotic channel bonding concepts will be necessary, and the burdens to the wiring harness mass and complexity need no further elaboration.

At the same time, the performance level of processors continues to increase as we expect due to the advancement of semiconductor technology. The advent of field programmable gate arrays (FPGAs) makes it possible to create custom processing engines that may be tuned to specific algorithms with a much higher performance than can be achieved with traditional processor architectures alone. Since computation is necessarily distributed in these high-performance systems, the movement of data requires high-speed networks. As more FPGA-based processors move into space system designs, they too drive the need for extremely high-bandwidth transport between components.

Gilder’s *Telecosm* [6] paints a picture of a world of “infinite bandwidth”, in which every person (or object) possesses an individual wavelength (like a phone number) and a connection to a ubiquitous network referred to as the “telecosm”. Individuals need only “dial the wavelength” of another to send (on demand) arbitrary torrents of data with little thought or consequence of the limits of bandwidth. A

“telecosmic” perspective informs us of the potentialities of architecture when we do not worry about interconnect limitations.

While it seems fantastic to think a time might exist when the physical spectrum could ever admit a true telecosm, it is insightful, particularly in a more compact system such as a single spacecraft (where the number of objects is far smaller), to consider the architectural possibilities that “bandwidth for free” might allow. In the case of sensors, we could envision the ability to simulcast raw video streams or entire radar feeds to many points within a system. Any new “object” added to the system could “subscribe” to the feed. These objects could be processor farms (including FPGA-based), mass memory storage systems, or transponders. The memory systems could (assuming they too are seemingly “infinite”) ingest these feeds as is, or subscribe to a “compression on demand” service from a processor elsewhere in the spacecraft. Processing farms could execute one (or many) algorithm(s) on the raw data to produce a variety of useful mission products, and these products can also be simulcast to raw storage, to other processors, or off-platform through a transponder (high-speed rf or laser). The memory storage systems that might evolve under the environment of “bandwidth for free” would be very powerful, not only having vast storage banks (petabytes or more), but would be capable of managing dozens (or more) transaction streams “in flight”. The processing banks would interact with each other through the “fat pipes” and/or use the memory resources with equal facility. The fluidity of powerful memory and storage facilities might promote new styles of algorithm development in which algorithms are designed to be dynamically amortized over the pool of available resources. This might mean that systems can directly execute such algorithms with whatever pool of resources they possess, and system designers can achieve performance gains by “merely” adding more processor blocks, memory blocks, etc. These concepts lend to the notion of freely scalable, plug-and-play supercomputing on orbit, where limits are not measured by the tedious and error prone ability of human teams to painstaking map custom computation to elaborate proprietary networks, but instead to eventually build powerful algorithms once, and let the hardware (and the system) deal with the detailed mapping of computation to resources and the scheduling of transactions over the fat pipes, since we would have “bandwidth for free”.

Of course, we are quickly reminded that not only is bandwidth, memory, and processing not free, but that the medium of an orbiting spacecraft offers a particularly limited canvas. Nevertheless even space systems have received the benefits of Moore’s Law, achieving performance doublings every 18-24 months (theoretically, space systems deal with additional

complications, such as radiation and fault tolerance). We may not have infinite bandwidth but have the ability to access enough of it to responsively meet our needs on demand. This is a proposition not of “bandwidth for free” but “pay as you go”, whether we mean actual currency (expressed in the fiscal constraints in constructing a system in the first place) or possibly the currency of energy in orbit (in which case we budget for the streams we use based on availability of battery and solar cell generation). Engineering a system based around the “fat pipes” in this case doesn’t mean we use them to capacity, but eliminate them as a primary bottleneck constraint. When the freeway systems (such as the US interstate network) were commissioned in the mid-20th Century, no concept probably existed for the types of traffic that would course through them in the 21st Century. They were built for the highest conceivable capacity, and in many cases that has been “good enough”

These considerations underlie the definition of SPA-O, a plug-and-play “superfreeway” in which the highest performance objects of plug-and-play systems may be freely intermingled. Many aspects of complex networks must be considered when planning for plug-and-play capabilities. There are always new capabilities, functions, and hardware just around the corner. Reducing the amount of integration, construction, and testing time must always be considered. It takes up to fifteen years to design, build, and fly most satellites, counting interminable overruns [7]. The goal of SPA, in general, is to get that time down to just over one week.

DIRECTION OF SPA-O IMPLEMENTATION

We first make an assertion that SPA-O *needs* to be optical. While it is possible to achieve fixed targets (e.g., 10 Gbps) through increasingly elaborate engineering exercises with copper high-speed serial transport, the movement to optical provides far more performance and flexibility. With optical comes wavelength diversity, the gateway to terabit/s performance. Optical leads to a “colorless” system, having no pre-ordained wavelengths. By employing a tandem copper network as a control plane, it is possible to detect components, even those that would otherwise be optically incompatible, and seek valid constructions within a given system to satisfy mission requirements through reconfiguration. The dynamic management of optical resources across a network in the face of disparate wavelengths, bandwidths, spatial path assignments, protocols, and modalities is referred to as

provisioning. Provisioning is managed by SPA, as another application. Plug-and-play provisioning leads to ambition to be totally inclusive to all optical and/or electrical components, and to be backwards compatible with the other SPA interfaces. This optical medium is ultimately intended to be completely scalable up to even multi-terabit/sec rates.

These ambitions have also been reflected in terrestrial proposals, such as the Generalized Multi-Protocol Label Switching (GMPLS) Architecture [8], in which disparate components are managed in automatic switched networks. The difference in our approach is that we make the mechanisms for “automatic” more explicit by defining them within the framework of SPA. In the SPA approach to provisioning, SPA-O devices use the xTEDS to describe the optical capabilities and restrictions of the underlying component. When a number of SPA-O components are placed in a SPA system, their xTEDS are ingested and reconciled throughout the system to match component subscribers of optical streams to the components that produce them. Applications that run over SDM, combined with the control cues embedded in SPA-O devices, form the control plane infrastructure for achieving this.

As a simple example (Figure 2), if we have a optical multi-spectral sensor producing raw imagery on fiber number three connected to a high-speed processor on fiber number five (Figure 2a), we will take steps to align the stream by reassigning the sensor output (Figure 2b), processor input (Figure 2c), or flipping the appropriate switches in an optical switch (between the two) using the SPA-managed control plane (Figure 2d). More sophisticated examples include transactions requiring traffic grooming, traffic de-commutation, wavelength translation, protocol conversion, and bandwidth re-alignment. There may be cases when an optical sensor needs to be connected in the SPA network to a non-optical component. In this case, some (possibly constricted) form of the stream may be converted from the optical domain to the electrical domain and routed as traffic over the control plane itself. Dynamic provisioning operations of this kind add flexibility, scalability, and fault tolerance to the system at large. Redundant resources can be added, such that failures can be circumlocuted through remapping streams to the backup components. Even in the case of optical link breakage, some possible mission recovery options might be enacted (even with degraded performance) by dynamically mapping a lower bandwidth version of the same “service” over the electrical control network.

Figure 2. Example of simple optical fiber provisioning. (a) Sensor and processor misalignment. (b) Reconfiguration of sensor. (c) Reconfiguration of processor. (d) Use of configurable switch.

More exciting is the possibility of network upgrades, using sophisticated future optical components, including all-optical routers, processors, and storage components. Using the SPA paradigm, these resources can be quickly assimilated and exploited by existing components and software. From the black box perspective (and the detached observer vantage) such components simply seem to be more capable black boxes (even if “more” is a million-fold more!).

CURRENT RESEARCH

The current work is in setting up a testbed that will utilize optical fiber for higher data rates beyond the 400Mb/s limit with SPA-S (Typical “low end” data rates for the optical layer are 1Gb/s to 2Gb/s.). This testbed is located at the Air Force Research Lab at Kirtland AFB, NM. It is designed to have two networks with fundamental differences to allow as many options, and integration issues to be solved as possible. Optical SPA will still use SPA-S as the underlining command and control protocol. Header information, routing commands, and even most of the low data rates will still be on Spacewire (electrical). A non-blocking electrical router will connect to the ASIMs of each component and will handle transmissions up to 400Mb/s. In addition to this, an optical switch will be controlled from the ASIMs via its electrical interface (RS-232). One network will consist of sing-mode (SM) fiber, and the other will implement

multi-mode (MM) fiber. The two networks will be connected via a free space link and will be gradually modified to encompass a true ‘network in the sky.’ Designs have 10Gb/s (aggregate) as the ceiling, but scalability to 100Gb/s (aggregate) and beyond are feasible.

The testbed housed at Kirtland AFB, NM, is intended to allow any component or application developer to bring in hardware and software for integration and testing. As the system progresses to flight, bottlenecks and shortfalls will be discovered. A majority of the current research dedicated to SPA-O is funded through the Air Force Small Business Innovative Research (SBIR) program. There are significant core funds dedicated to simulation and modeling, processing and mass memory storage, as well as other hardware components. Figure 3 shows the current design for the testbed being constructed for multiple 2010 demonstration. It will entail three to four ‘sensors’ comprised of a 10.7 Gb/s data generator, a 3.2 Gb/s camera, and a DVD ‘live’ video feed. The ‘heart and brains’ will be the optical switch, controlled by an ASIM, integrated with an electrical router. The processor and memory will be co-located and not radiation hardened at this time. An oscilloscope and monitor will also be connected as output interfaces per network. The crosslink to the second network will require an ‘o-e-o’ conversion to allow single mode free space transmission.

SPA-O Layout (June 2010)

Distribution subject to Notice on Cover Page of Briefing

20

Figure 3. SPA-O Testbed layout

A simulation and modeling effort has been initiated with Sandia National Labs. Initial investigations will focus on network architectures using current hardware to create accurate models. The effort will look at quality of service, bottlenecks, and overall performance of selected network topologies. The end goal is to provide the ability to test and gage hardware before initial trials in the Kirtland testbed. The focus is on delay, bit error rate, and system setup. Using this software will enable one to check the efficiency of using one protocol over another, or which hardware provides the best link margin, or even test the potential lockup of the communication network to verify that a thruster will fire upon request. The software is utilizing 'Simi-Link' and has only gone through basic system setup at this time.

Contracts are developing the best design of an optical switch integrated with an electrical router. Other contracts are researching future components involving true optical routing, IP and other routing protocols, high-speed optical transceivers, and high-speed interfacing, processing, and mass storage for space. All components will have to be developed using radiation hardened designs and hardware. Most of these hardware development programs will also use core funds. As these contracts output a prototype or even flight-ready hardware, they will come to the testbed for integration and final testing. The testbed

will then pass them off as ready to fly when the next launch calls for the particular capability.

CONCLUSION

Space Plug-n-Play Architecture allows components to be integrated and tested rapidly so that complicated systems may be integrated and launched in days not years. Legacy systems have proven this mentality and have advanced to a point of allowing many differing components to be recognized and many data rates utilized. Low data rates will remain in the electrical medium and sometimes be transmitted wirelessly, while the high data rates up to and beyond hundreds of Gbps will be sent through the optical domain. The testbed found at Kirtland AFB, NM, will allow integration and testing to continue for future components. This open-source topology will allow most components to be quickly and readily flown, while still allowing the proprietary 'black-box' component, owned by a company, to do its full mission. The future of satellite construction is promising to be faster, cheaper, flexible, and scalable. While components mature and integration becomes easier, space and airborne networks will integrate smoothly and transmit more data at the speed of light.

Acknowledgments

The authors would like to thank AFRL and Dr. Steven Lane along with the many contractors and companies who have contributed to this effort.

References

1. Lyke, J., "Space Plug-and-play Avionics (SPA): A Three-Year Progress Report," *AIAA Infotech@Aerospace*, AIAA, Rohnert Park, CA, 2007.
2. A. Barabasi, *Linked: The New Science of Networks*. Perseus Publishing, 14 May 2002.
3. Lyke, J., Cannon, S., Fronterhouse, D., Lanza, D., and Byers, T. "A Plug-and-play System for Spacecraft Components Based on the USB Standard", proceedings of the 19th Annual AIAA/USU Conference on Small Satellites, Logan, UT, 8-11 August, 2005.
4. SPA-S
5. SPA-1
6. Gilder, George. *TELECOSM: How Infinite Bandwidth will Revolutionize Our World*, Free Press, September 2000.
7. Cáceres , Marco. "Cost overruns plague military satellite programs", *Aerospace America* (publication of AIAA), January 2006, pp 18-20, 23.
8. "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", document of the Internet Engineering Task Force, *Request for Comments (RFC) 3945*, October 2004 (available at <http://www.ietf.org/rfc/rfc3945.txt>, accessed April 2010).