The Space Ethernet PHYsical layer transceiver (SEPHY) Project: A step towards reliable Ethernet in Space

Authors:

P. Reviriego, Universidad Antonio de Nebrija

J. López, Arquimea

M. Sánchez-Renedo, Thales Alenia Space

V. Petrovic: Innovations for High Performance Microelectronics (IHP)

J. F. Dufour, TTTech

J.S. Weil, Atmel

Corresponding author:

P. Reviriego (email: previrie@nebrija.es, phone: int+34-626068722)

Conflicts of interest and financial disclosure:

We are not aware of any conflict of interests regarding this submission.

This work presented in this paper is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No.64024.

Abstract:

Since its development, Ethernet has experienced an impressive growth and has become the dominant technology for wired local area networks. It has also more recently expanded beyond computer networks to cover also industrial and automotive networks. This adoption is driven by the lower costs enabled by reusing existing technology. For critical applications, Ethernet has to be extended to ensure timely and reliable delivery of frames. A number of technologies that can solve the reliability and real time issues have been proposed, for example Time Triggered Ethernet (TTE). Space systems are an example of critical applications and Ethernet has been used in some missions like NASA´s Orion and in launchers. However, there is an additional problem in space applications that has so far prevented a wider adoption of Ethernet. Electronic circuits that operate in space are exposed to radiation that causes errors and make most commercial devices not suitable for space missions. This means that special components have to be designed for space use. This can be done for purely digital components using radiation tolerant FPGAs qualified for space use or radiation hardened libraries. However mixed-signal devices remain a challenge and specific ASICs are needed in most cases. For Ethernet, most components like switches or Medium Access Controllers (MACs) are purely digital. There is however one exception, the physical layer transceivers (PHYs) that are by nature mixed-signal devices. The availability of rad-hard Ethernet PHYs qualified for space use is crucial to enable the widespread adoption of Ethernet in space. This paper discusses the challenges in designing such transceivers and presents the SEPHY project that is currently developing a 10/100 Mb/s European Ethernet transceiver for space.

Introduction

Space applications are no strangers to the need for speed of modern networks. Scientific missions use increasingly powerful sensors that produce images and data that requires high-speed communication while each new generation of commercial satellites needs to support higher data rates. This trend is expected to continue in the future and therefore there is a need to define a roadmap for space networks that supports this vision. Additionally, networks are needed for a variety of functions in a space vehicle that range from non-critical sensor data collection and processing to tele-commands that are vital for the system operation. Ideally, future space networks can integrate all those functions on a single network thus simplifying the system design and operation. Today, most space networks are based on standards that support only low bit rates like the very deterministic MIL-STD-1553 or that are specific to lower-criticality payload data like Spacewire [1]. In the first case, the standards cannot support the speeds needed for future systems while in the second, the development cost are high and may not justify future evolution to higher speeds.

The two issues faced by space networks: the integration of different functions on a single network and the evolution to higher speeds, have already been solved for terrestrial networks. In the first case, industrial and safety critical applications like transportation have fostered the development of enhancements to Ethernet so that critical applications can be supported [2]. In particular, Time-triggered Ethernet (TTE) provides deterministic performance for critical functions while integrating all the functions on the same networks [3]. Additionally, TTE implements fault tolerance mechanisms so that the network is resilient to failures [4]. This has led to the adoption of TTE for some space missions.

The use of Ethernet also solves the need for higher speeds. The IEEE 802.3 standards already support high data rates (10 Gb/s for twisted pairs) and several task forces are working on future standards that will provide even higher rates [5]. The reuse of Ethernet technology not only provides the standards for future evolution, it also reduces the costs. The development of high-speed communication standards requires a large cost that may not be justified for space applications. Reusing Ethernet means that most of those costs are avoided. In addition technology risks are mitigated as the standards have been widely used in millions of terrestrial systems. The use of Ethernet also facilitates testing the components for standard compliance, interoperability with ground testing equipment and prototyping using commercial components. All these advantages of Ethernet have spurred a wide interest in the space community to use Ethernet in space systems. In particular Ethernet is used in many state-of-the-art radiation-tolerant system-on-chip (SoC) often as a debug or output interface. Deterministic versions of Ethernet, for instance Time-Triggered Ethernet, are also being baselined and used for command and control functions onboard spacecraft. The most advanced multi-purpose crew vehicle (MPCV) of NASA (supported by the European Space Agency) using Time-Triggered Ethernet for its main critical data bus, is a good example of how Ethernet can be used for man-rated space missions. In Europe, TTE has recently been baselined for the avionics central communication system of the new European launcher family, Ariane6.

However, as mentioned in the abstract, there is a specificity of space applications that is preventing a fast adoption of Ethernet. The problem is that space systems are not protected by the earth atmosphere and are subject to much larger radiation levels than terrestrial systems. Radiation causes a number of effects on electronic circuits including transient and permanent failures [6]. This means that most commercial components cannot be used in space systems. In many cases, special radiation hardened devices are used in space, for example for transceivers [7]. This hardening is achieved by using specific manufacturing processes and design libraries. In addition, the designs can also incorporate fault tolerant techniques such as the use of Error Correction Codes to protect memories or registers [8]. For purely digital circuits, radiation hardened FPGAs or design libraries are available to facilitate the implementation of systems. For mixed-signal circuits, the situation is worse as there is no platform on which to implement the system.

This creates a problem as in Ethernet the physical layer transceiver (PHY) is by nature a mixed-signal device that process analog signals from the cable and transforms them into digital signals at the Medium Access Control (MAC) interface. Therefore, a space grade Ethernet PHY is needed to enable widespread adoption of Ethernet in space. So far, this has limited the use of Ethernet to some systems like launchers in which commercial components previously tested can be used as they are exposed to radiation for a very small period of time (at most a few hours). In other cases such as the NASA Orion mission, Ethernet is used with no physical layer transceiver. This however limits the network span and increases the cost.

In the rest of the paper, the issues and challenges in developing a space grade Ethernet PHY are discussed. The Space Ethernet PHY (SEPHY) project that is currently developing a European transceiver is also presented.

Radiation effects and mitigation

As mentioned in the introduction, radiation is a major issue for space electronics [6]. Electronic circuits that operate in space are exposed to a number of particles that interact with the circuit elements creating errors. For example, heavy ions are present in low and medium earth orbits and can cause Single Event Effects, which can include upsets, transients and induced latch-ups among others.

The effects on the circuits can be broadly divided in cumulative and Single Event Effects (SEEs). Total Ionizing Dose (TID) is an example of a cumulative effect. As the circuit is exposed to radiation, increments in leakage currents or threshold voltage variations can be seen. This eventually leads to a failure. On the other hand, single event effects are caused by a particle hit and can have a transient effect changing the value of a circuit node or cause a permanent failure in a circuit element. Single Event Upsets (SEUs) that change the logical value stored in a register or memory cell are a common type of transient SEEs. Single Event Transients (SETs) are voltage transients observed in analog or digital signals due to a particle hit. Some of these effects can be destructive, as Single Event Latchups (SELs), Single Event Gate Rupture (SEGRs) or Single Event Burnout (SEBs). Displacement Damage effects (DD) should also be considered.

These effects pose a specific challenge for space electronics that have to be designed to mitigate them and ensure a reliable operation. The amount of radiation varies greatly with the mission, for example a launcher will be in operation for only a few hours and therefore cumulative effects should not be an issue. On the other hand a satellite may be in operation for many years making cumulative effects an important problem. Since it is not cost effective to develop different circuits for each operating environment, there are a set of commonly accepted requirements for space electronics, such that a component can operate in a range of missions. For example, for communication satellites in a geostationary orbit and 15 years mission, a TID of around 100 krads is typically accepted. In terms of SEU, and SEL tolerance a Linear Energy Transfer Threshold (LET) of 70 MeV*mg/cm^2 is considered as standard.

In order to achieve those requirements, a number of techniques can be used [8]. For example, it is common to use Silicon on Insulator (SOI) technology for the devices as this prevents the occurrence of SELs. The design library can also support mitigation features [9] and redundancy can be added in the form of triplication or error correction codes, to mitigate SEUs impact on circuit functionality. The use of all those mitigation techniques increases the area, power and delay of the circuit and is not attractive for commercial devices. Therefore, the market for space components is in most cases small, although they can be of interest for avionics, military applications and systems that operate in high radiation environments like nuclear plants, some medical systems or high energy physics facilities. In addition, the technologies qualified for space use are commonly several technology generations behind the latest commercial technology.

As discussed before, space grade components in most cases have to be ad-hoc designed for a small market. In the case of digital circuits, an attractive option is to use radiation hardened FPGAs that can withstand radiation and can later serve as the platform on which to implement the system. This reduces the development cost significantly. For mixed-signal circuits, although there have been some efforts to implement programmable devices, there is currently no radiation hardened platform that can be used for implementation. In addition, mixed-signal circuits pose several challenges related to radiation effects, which implies that specific radiation tests need to be performed for every device type.

The need for a rad-hard PHY and the options

From the previous section, it becomes clear that radiation hardened Ethernet components are needed to enable widespread use of Ethernet in space systems. The Ethernet standards are divided in layers that include for example the Physical (PHY) and Medium Access Control (MAC) layers [10]. The MAC layer and the Ethernet switching are purely digital and therefore it is easier to migrate them to a rad-hard version. However, the physical layer transceiver (PHY) interfaces directly with the transmission media and therefore is a mixed-signal device. This makes the development of a rad-hard PHY more complex and creates a potential bottleneck for the adoption of Ethernet in space.

Once the need for a Space PHY is clear, the question is which PHY should be developed. The IEEE 802.3 standard defines many PHYs covering different transmission media and speeds. The most commonly used media in Ethernet are Unshielded Twisted Pairs (UTP). Assuming that the PHY will use UTP, the IEEE 802.3 standard provides several alternatives. The most relevant ones are: 10BASE-T defined in IEEE 802.3i , 100BASE-TX defined in IEEE 802.3u, 1000BASE-T defined in IEEE 802.3ab and 10GBASE-T defined in IEEE 802.3an. Each of those standards provides a 10x speed increase over the previous one, starting with the 10 Mb/s of 10BASE-T. This shows how Ethernet enables the increase in network speed with each new standard. From a performance point of view, the best would be to select the highest speed PHY for rad-hard implementation. However, there are other factors that should be considered when making a decision. As the speed increases, so does the complexity of the PHY. For example, 10GBASE-T PHYs are currently manufactured in 40 or 28 nm technologies and consume several watts. Implementing that PHY on the older technologies qualified for space use will most likely not be feasible. The development cost also increases with speed. This is not a show-stopping issue for commercial applications where the cost is spread among millions of devices. However, this is not the case for the space market where volumes are orders of magnitude lower. Therefore, the selection of the PHY standards to implement for the space market needs to consider both the speed and the cost/complexity.

To discuss the trade-offs, a basic understanding of the UTP medium and the different standards is needed. The UTP cable typically has four pairs as shown in Figure 1. Transmission can take place in some or in all of them in both half duplex or full duplex modes. Transmission on a pair causes echo on the same pair and both near end and far end crosstalk on the other three pairs as shown on Figure 1. Therefore, as more pairs are used, speed increases but also crosstalk that has to be cancelled is created in other pairs. The same applies to the use of the full duplex mode that introduces echo and further complications at the receiver side.

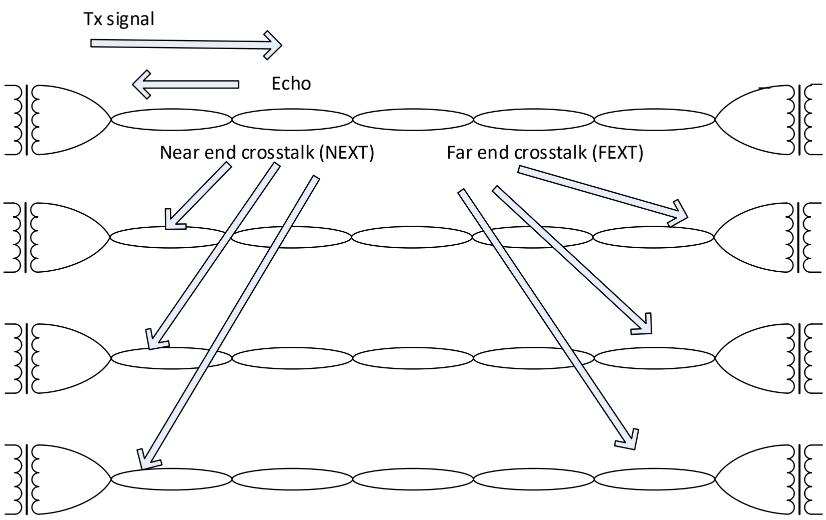

Figure 1 UTP transmission environment

The first standards (10BASE-T and 100BASE-TX) use only one pair in half duplex mode for each direction as shown in Figure 2 (top). Therefore, there is no echo and no far-end crosstalk. This greatly simplifies the transceiver design. For 100BASE-TX the speed increase is achieved by using a larger transmission frequency and number of levels. In the cable, higher frequencies are attenuated and need equalization, also crosstalk and echo grow with frequency. In any case both standards can be implemented with a moderate cost on an old technology. On the other hand, the 1000BASE-T and 10GBASE-T standards use the four pairs in full duplex mode as illustrated in Figure 2 (bottom). This means that the receiver on each pair needs to cancel the echo and the crosstalk from the other three pairs. Additionally, these two standards incorporate a more sophisticated coding scheme (Trellis Code Modulation in the first case and Multilevel Coset Coding in the second) that need complex decoders. This makes the implementation of the transceiver a challenging task. In the case of 1000BASE-T, commercial devices were implemented in old technologies like 180nm, suggesting that a rad-hard implementation may be feasible. On the other hand, a 10GBASE-T transceiver needs ADCs and complex DSP functions running at more than

800Mhz, a speed that it is very difficult to achieve with the available space grade technologies.

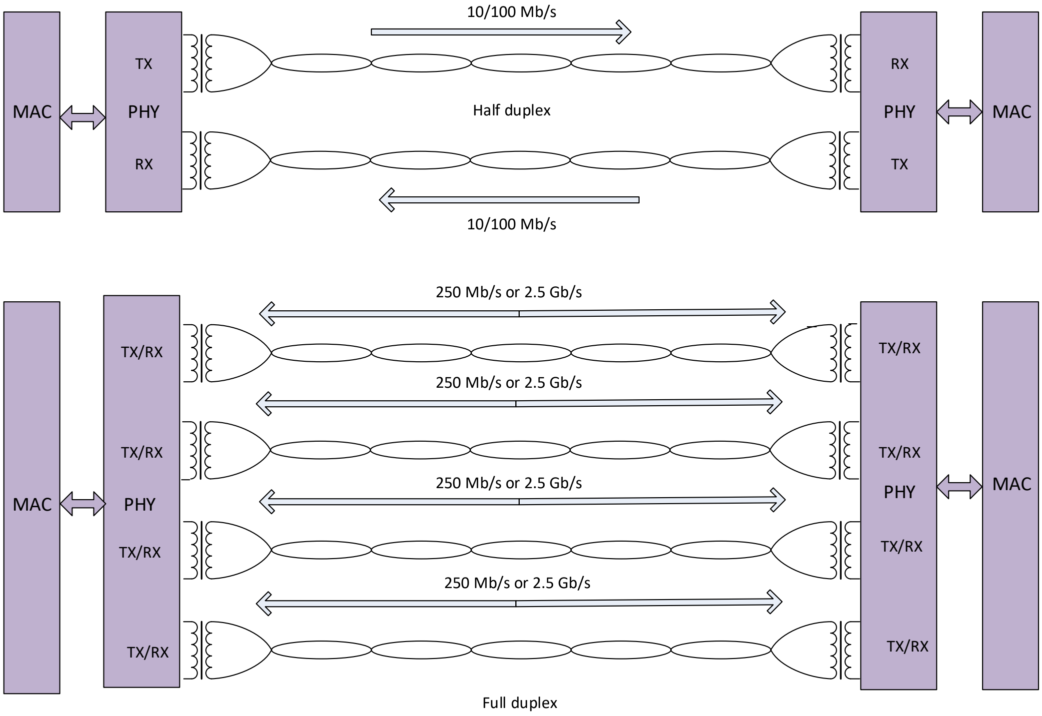

Figure 2 Transmission in 10BASE-T/100BASE-TX (top) and 1000BASE-T/10GBASE-T (bottom)

In addition to the existing IEEE 802.3 standards, the standards currently under development in several IEEE 802.3 task forces can be an option for a rad-hard Ethernet PHY. In particular, the 100BASE-T1 (P802.3bw) and the 1000BASE-T1 PHY (P802.3bp) Task Forces are defining standards for 100 Mb/s and 1 Gb/s communication using a single twisted pair. These standards are targeted to automotive applications on which weight and cabling are important cost factors. The use of a single pair is also attractive for space applications, but the use of an unfinished standard implies significant implementation risks. Therefore these standards are better suited for future generations of rad-hard PHYs once the technology has been validated on terrestrial applications. There are also some additional task forces working on higher speed (2.5, 5, 25 and 40 Gb/s) UTP PHYs, [5], [11] but again those seem more appropriate for future generations of rad-hard PHYs.

A European Ethernet transceiver for space: SEPHY

The Space Ethernet PHYsical layer transceiver (SEPHY) project funded by the European Union Horizon 2020 research program, is currently developing a radiation hardened PHY. The availability of a European PHY is key to ensure that access to the PHY is not restricted by the United States International Traffic in Arms Regulation (ITAR) and the Export Administration Regulation (EAR), and as a consequence, the non-dependence for the European space industry is guaranteed. The goal of the project is to deliver a production-worth PHY in 2017.

The technology selected to implement the SEPHY device is Atmel´s 150 nm Silicon On Insulator (SOI), as it provides a sufficient level of radiation tolerance that qualifies it for space applications[1]. It is also the same technology for which other European Ethernet components are being developed in the FLPP3 Time-trigered Ethernet Space ASIC project.

The project consortium is formed by different European companies and research centers led by Arquimea that will develop the analog components. IHP will focus on the digital design and Universidad Antonio de Nebrija on the verification. Atmel will be in charge of the fabrication of the integrated circuits. Finally, TTTech and Thales Alenia Space Spain will integrate and test the silicon prototypes on a network and perform also radiation testing.

The project targets the implementation of the 10BASE-T and 100BASE-TX standards. This will provide 10Mb/s and 100Mb/s connectivity in space systems. This compares to the solutions currently used in the space domain like the Mil-Std-1553B (low data rate, large cable length) or SpaceWire (high data rate, short cable length). The proposed transceiver

[1] Note that although Atmel is an US company, its Aerospace division is fully based in Europe, only uses non US technologies, and is consequently not subject to US export control laws and regulations (ITAR and EAR).

will meet the cable length (100m) and data transfer (100Mbps) requirements not only for launchers applications where cable length is the main constraint but also for the onboard communication requirements where high data rate is required. As discussed before, those standards can be likewise implemented with a reasonable cost and provide a solution to the industry needs in the short term. The development of a 1000BASE-T PHY would imply much larger cost, time and risk and could jeopardize the adoption of Ethernet. Additionally, starting with lower speeds gives the opportunity to the new standards being developed to mature potentially providing more choices for the second generation of SEPHY. In fact, the project also includes a roadmap activity to identify the best alternative for a second generation of SEPHYs. This will target at least 1Gb/s and its feasibility will also be studied and linked to the future technology planned for space devices.

The SEPHY project uses a number of techniques at different levels to mitigate effectively the radiation effects on both the analog and the digital components of the transceiver. Firstly, a Silicon On Insulator (SOI) technology is used to manufacture the device. This eliminates the Single Event Latchup events that could cause permanent or destructive failures and also provides a good immunity against Total Ionizing Dose (TID). Additional rules are used in the design to minimize the effect of radiation. For example, to protect NMOS transistors against TID, the Width/Length (W/L) ratio and the length (L) have to be larger than the minimum allowed by the technology and the use of Enclosed Layout Transistor (ELT) topologies is preferred. The pads are also designed following special rules to mitigate radiation effects. To reduce the impact of Single Event Transients, nodes with high impedance are avoided and differential circuits are used if possible. On the digital side, Finite State Machines are designed so that any unexpected combination of state and inputs lead to a known state. For the flip-flops that are critical, a cell that incorporates TMR protection is used while registers and memories will use Error Correction Codes to ensure that errors will not corrupt their contents. All these rules will reduce the probability of radiation affecting the operation of the transceiver. Additionally, both the analog and digital blocks include checking mechanisms to detect an abnormal behavior. Those are then used to restart the block or the entire transceiver and bring it back to a normal state. These mechanisms are designed to cope with the few errors that will not be mitigated by the technology and design rules. For example, in the 100BASE-TX mode, the output of the equalizer can be monitored to detect abnormal conditions [12]. In 10BASE-T mode, the reception of Normal Link Pulses (NLPs) can also be used to check that the transceiver is operating correctly. That means that existing features of the transceiver can be monitored to decide if it is operating correctly.

To ensure that all those mitigation techniques are effective it is important to perform accelerated radiation testing experiments. In these experiments, the devices are exposed to the radiation that they will suffer in space but in a short period of time using a powerful radiation source. The SEPHY prototypes will be tested for both Total Ionizing Dose (TID) and Single Event Effects in independent experiments.

Conclusions

In recent years there is a growing interest in using Ethernet in space systems. There are already some examples of space missions and systems that have used Ethernet. However, the further development of radiation hardened Ethernet physical layer transceivers (PHYs) is critical to enable widespread adoption of Ethernet in space systems. This paper discusses the options and challenges to implement space grade Ethernet PHYs and introduces the SEPHY project, currently developing a radiation hardened PHY to tackle some of these challenges. To this aim, particular trade-offs have been made in the goals to facilitate a rapid adoption of Ethernet in space while supporting a long term vision for its evolution in the future.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No.64024.

References

1. R. Klar, et al. “Design Considerations for Adapting Legacy System Architectures to SpaceWire,” International SpaceWire Conference, 2008.
2. J.-D. Decotignie, “Ethernet-based real-time and industrial communications,” Proceedings of the IEEE, vol. 93, no. 6, pp. 1102–1117, Jun. 2005.
3. W. Steiner, G. Bauer, B. Hall and M. Paulitsch, “Time-Triggered Ethernet: TTEthernet, ” In Time-Triggered Communication, R. Obermaisser, editor, CRC Press, 2011.
4. A. Loveless, “TTEthernet for Integrated Spacecraft Networks” AIAA Space and Astronautics Forum and Exposition (SPACE 2015), 2 Sep. 2015.
5. IEEE P802.3bq 25G/40GBASE-T Task Force (available online: http://www.ieee802.org/3/bq/public/index.html).
6. R. D. Schrimpf and D. M. Fleetwood, “Radiation effects and soft errors in integrated circuits and electronic devices”, World Scientific, 2004.
7. T. England, C. Chatterjee, N. Lourenco, S. Finn, L. Najafizadeh, S. Phillips, E. Kenyon, R. Diestelhorst and J. Cressler, "Cold-capable, radiation-hardened SiGe BiCMOS wireline transceivers," IEEE Aerospace and Electronic Systems Magazine , vol.29, no.3, pp.32-41, March 2014.
8. M. Nicolaidis, “Design for soft error mitigation,” IEEE Transaction on Device and Materials Reliability, vol. 5, no. 3, pp. 405–418, Sep. 2005.

9. S. Redant et al., “Radiation Test Results on First Silicon in the Design Against Radiation Effects (DARE) Library”, IEEE Transactions on Nuclear Science, vol. 52, no. 5, October 2005.
10. C. E. Spurgeon and J. Zimmerman “Ethernet: The Definitive Guide”, 2nd Edition, O'Reilly Media, 2014.
11. IEEE P802.3bz 2.5G/5GBASE-T Task Force (available online: http://www.ieee802.org/3/bz/index.html).
12. P. Reviriego, J.A. Maestro, S. Liu, "Efficient Soft Error-Tolerant Adaptive Equalizers", IEEE Transactions on Circuits and Systems I, vol. 57, no 8, August 2010, pp. 2032-2040.