How Can Commercial Fiber Equipment Cope with the Random Fadings of FSO Links?

G. Cossu[(1)], A. Sgambelluri[(1)], F. Paolucci[(2)], E. Ciaramella[(1)]

[(1)] Scuola Superiore Sant'Anna, Via Moruzzi 1, 56124 Pisa, Italy. giulio.cossu@santannapisa.it
[(2)] Consorzio Nazionale Interuniversitario per le Telecomunicazioni, Via Moruzzi 1, 56124 Pisa, Italy.

Abstract *In future satellite links, Free Space Optics (FSO) must cope with random deep fading. Here, we experimentally investigate recovery times of different commercially available solutions, with some impressive difference among the tested equipment.This should stimulate further optimization.*

Introduction

Future satellite communications will be based on high-speed Free Space Optics (FSO) links, largely exploiting the heritage of existing components and subsystems for optical fiber communications. To reach this goal, it is evident that the key difference between fiber and satellite communications lies in the channel. Among the various differences, the fiber impairments (dispersion, nonlinear effects, etc.) are complex, yet stable, features of the specific link (with relevant exception of the PMD). On the other side, in a real FSO link, especially in those passing through the atmosphere, the received signal shows random intensity fluctuations (due to combination of many different effects[1]). This eventually results into random fading effects, typically with duration of tens of ms[2],[3], which can determine a complete loss of signal.

We highlight that fading in today optical fiber communication networks is not present nor considered: the design process of systems and networks assumes that the links can either be always on or forever off, the last due to some unrecoverable issue. This assumption is reflected in the design of the equipment and network protocols. Unfortunately, when trying to move today equipment from terrestrial networks to satellite, we should ask questions about the capability of the fiber-based hardware to manage random fading events. Namely, once a deep fading is present, we expect that the link will show a non-negligible recovery time, with additional loss of packets: this is affected by the different features of the equipment, such as the physical layer specifications (e.g. CDR, FEC), but also the type of signal framing and alarm flags. Namely, non-trivial interaction about the different communication layers can have a major impact.

Here, we provide the results of the first experimental assessment of fading robustness in various types of commercial equipment, emulating a transmission with controlled fading events. To perform these tests, we setup a testbed that emulates the fading loss by means of an intensity modulator on bidirectional links, with different equipment: in that link, from the measured packet loss we derive the recovery time, which we found, in some cases, much longer than the fading itself.

Experimental setup

In order to evaluate the effects of the fading over the commercially available optical transceivers, we setup the testbed shown in Fig. 1. In particular, a Spirent N4U traffic generator and analyzer equipped with 1 Gigabit Ethernet (GbE) interfaces was exploited to generate the probing Ethernet traffic streams to be injected in the transmission system. A programmable P4-native switch[4] was adopted to maintain the Spirent interfaces alive during the experimental validation, while the fading affects the optical channel. The switch was configured with a set of permanent flow entries. Each flow entry allows to forward the packets from an input port to an output port. Two pairs of commercially available muxponders have been considered to test the performance of OTN-based equipment over the end-to-end transmission, i.e., 10G and 100G. In the case of 10G muxponders, up to 6 1GbE tributaries are aggregated in a DWDM 10G OOK optical line. While, for the 100G muxponder, up to 10 10GbE tributaries are aggregated in a DWDM 100G DP-QPSK with coherent detection.

Three different options were considered as transceiver interfaces, each connected to the proper port of the programmable switch. First, we characterized a pair of 100 Gbit/s coherent transceivers, supported by Digital Signal Processing (DSP) and fully compliant with Optical

Transport Network (OTN) (ports 4 and 7), then we tested 10 Gbit/s direct detection transceivers that again implement OTN (ports 3 and 8). Finally, we assessed two 10 Gbit/s direct detection transceivers designed for GbE links (ports 5 and 6). The last ones do not include the OTN domain.

In order to emulate the fading effects, an Acousto-Optic Modulator (AOM) is placed in only one direction of the optical line, considering 10G and 100G muxponders. An AOM can simulate deep fading events since it can be switched instantaneously on and off, with a very high extinction ratio (typically 30 dB or higher). A waveform generator drives the AOM to produce a deterministic loss of signal over a configurable time window. To verify the link interruption, we split 1 % of the AOM output signal and we sent it to a Photo-Diode (PD) to measure the average optical power (this part is not shown in Fig. 1, for the sake of simplicity).

The setup is conceived to enable the continuous transmission of Ethernet frames from/to the Spirent interfaces, avoiding that the AOM trigger affects the interface status of the measurement device. Moreover, the programmable switch was selected to minimize the switching delay, keep the link stable during the AOM-triggered fading events and eliminate the effect of further layer-2 protocols (e.g., MAC learning). As an example, when the 10G OOK system is tested, the switch is programmed in a bidirectional way to forward all the packets from port 1 to port 3 and from port 8 to port 2 (and vice versa).

In each test, a single optical line is employed with a single AOM. Two unidirectional Ethernet streams are generated in the Spirent to test the performance of the transmission system. Each stream generates 10^6 frame per second resulting in a constant bit rate (CBR) traffic at 800 Mbit/s. The Spirent measures the number of packets lost due to the AOM-triggered fading event. Being each packet generated at 1 µs rate, each packet lost corresponds to an additional 1 µs, accounting for the considered transceiver failure and recovery time.

Results

In this section, we report and discuss the experimental measurements of the recovery times obtained with the three different commercial equipment, emulating a transmission through the atmosphere with single event of deep fading (i.e., no signal) within a specific time-period. Clearly, this

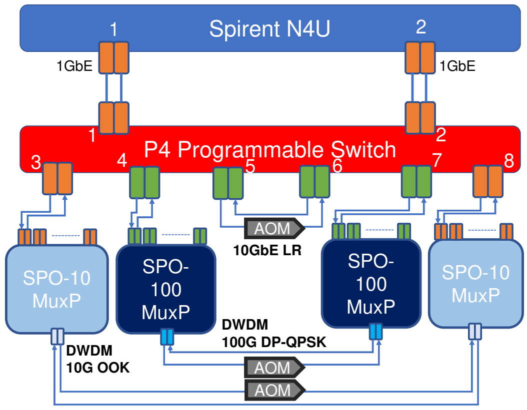

Fig. 1: Scheme of the realized testbed.

is a great simplification of the complex variations of the received signal power after the real atmospheric turbulence; however, this allows us to isolate and estimate the recovery times from a single event of selected duration.

In the measurements, the pulse width was logarithmically varied from 1 µs to 100 ms, covering (also) the expected fading values of a typical Earth-to-Space optical link (i.e., 1–50 ms[2]). The link was periodically interrupted with a period of 1 s: this is quite larger compared to fading and recovery times. The tests were automated by configuring in the Spirent a sequencer routine. The routine was repeated for 100 times running the two traffic streams for 10 s, collecting the metrics of each execution in a csv file. The selected metrics included the number of Ethernet frames lost, the minimum, maximum end-to-end latency. At any fading event, the Spirent gives the number of lost packets, which are due to those transmitted during the fading and those lost after the fading, i.e., during the recovery time. Surprisingly, we found that the second contribution can largely exceed the first one.

In Fig. 2, we report the net recovery times (in µs) subtracting the fading time from the total measured loss time as a function of the fading time (in ms), for the three considered transmission protocols. As expected, for any fixed configuration, the recovery time is not exactly fixed, i.e., is a statistical variable. Yet, very large differences in the statistics are experimentally obtained among the coherent transceivers, supported by DSP, and the direct detection transceivers that use OTN or designed for GbE links.

In Fig. 2a, we report the results at 100 Gbit/s using OTN, which is based on DP-QPSK modulation format and coherent detection. As can be noted, the recovery time is almost independent

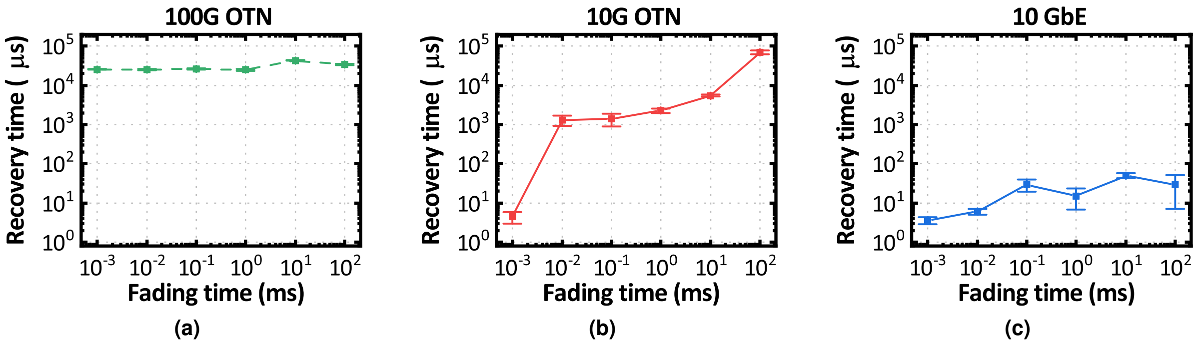

Fig. 2: Measured recovery times vs. fading time for the three different transmission protocols: (a) 100G using OTN; (b) 10G using OTN; (c) 10G Ethernet transceiver.

on the fading time. this can be explained considering that the recovery time is dominated by the DSP re-syncronization time around 30 ms, even for very short fading times of 1 µs. This is a significant problem in the FSO links, since after any fading event several gigabit of data would be lost.

Considering Fig. 2b, the 10 Gbit/s OTN recovery time is shown. The considered muxponders do not perform the end-to-end link status control, maintaining alive the transceivers towards the programmable switch (ports 3 and 8). The curve highlights two main contributions on the recovery time: fading effects with short duration (below 10 µs) are dominated by the OTN failure detection protocols with a recovery in the order of 5 µs, while longer fading effects trigger the Ethernet based recovery schemes with a resulting recovery time directly dependent on the fading duration (i.e., a fading duration of 10 ms is recovered in around 5 ms).

Finally, Fig. 2c reports the recovery time of the 10GbE transceiver. For very short fading times ($<$10 µs), the recovery time is dominated by the analogue circuitry for the signal resynchronization and reaches approximately 6 µs. Slightly higher values are measured with fading times $>$10 µs; here, the recovery times are lower than 100 µs; the fluctuations are due to not-deterministic Ethernet protocol, which includes the loss of optical power alarm.

Conclusion

In order to assess the potential for optical satellite communications, we have presented the first experimental emulation of random fading events on commercial fiber network equipment, including different types of transceivers (100 Gbit/s coherent, with OTN, 10 Gbit/s direct-detection with OTN or 10GbE): this was obtained by emulating a short (µs to ms) deep fading events, measuring the number of packet lost and then calculating the recovery time. It was found that the coherent transceivers with OTN have much larger issues than direct detection transceivers: this is likely due to the impact of the DSP, which gives recovery times on the order of 30 ms, even for very short fading events. It is possible that this might be improved by means of an optimized DSP, able to cope with the long fade events.

On the other side, the two direct-detection solutions typically have quite faster recovery. However, the one based on OTN still has significant recovery times (around 10 ms for a fading of 10 ms). Definitely, the transceivers based on 10GbE proved to be by far much faster (typically lower than 100 µs). The apparent differences in recovery times have a straightforward correspondence to the excess loss of signal packets: even assuming a retransmission solution, which is not simple to realize and implement[5], they involve much different additional complexity.

These results stimulate future extensive investigations, to further assess the potential of fiber equipment for space applications. Whilst in the long-term view, they can also encourage the industry to further optimize existing hardware, in the near future they can provide a useful hints to select the most suitable equipment for immediate space missions.

Acknowledgements

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”).

References

[1] M. A. Khalighi and M. Uysal, “Survey on free space optical communication: A communication theory perspective”, *IEEE Communications Surveys & Tutorials*, vol. 16, no. 4, pp. 2231–2258, 2014. DOI: 10.1109/COMST.2014.2329501.

[2] D. L. Fried, “Greenwood frequency measurements”, *J. Opt. Soc. Am. A*, vol. 7, no. 5, pp. 946–947, May 1990. DOI: 10.1364/JOSAA.7.000946. [Online]. Available: https://opg.optica.org/josaa/abstract.cfm?URI=josaa-7-5-946.

[3] W. Liu, K. Yao, D. Huang, X. Lin, L. Wang, and Y. Lv, “Performance evaluation of coherent free space optical communications with a double-stage fast-steering-mirror adaptive optics system depending on the greenwood frequency”, *Opt. Express*, vol. 24, no. 12, pp. 13 288–13 302, Jun. 2016. DOI: 10.1364/OE.24.013288. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-24-12-13288.

[4] Intel, *Intel tofino*, https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-series.htm [Accessed: (May 5 2023)].

[5] C. M. Schieler and B. S. Robinson, “Data volume analysis of a 100+ Gb/s LEO-to-ground optical link with ARQ”, in *Free-Space Laser Communication and Atmospheric Propagation XXX*, H. Hemmati and D. M. Boroson, Eds., International Society for Optics and Photonics, vol. 10524, SPIE, 2018, p. 105240L. DOI: 10.1117/12.2292910. [Online]. Available: https://doi.org/10.1117/12.2292910.