

Real-time demonstration of 2.5 GBaud coherent free-space optical communication through an atmospheric channel emulator

Carl Valjus, Raphael Wolf, and Sebastian Schmerberg

German Aerospace Center (DLR), Münchener Strasse 20, 82234 Wessling, Germany

ABSTRACT

Coherent optical detection offers superior receiver sensitivity at the cost of increased digital signal processing (DSP) and photonic complexity compared to on-off keying systems. For free-space optical (FSO) links, this additional sensitivity can be exploited to increase link margin or increase link range, improve availability under adverse atmospheric conditions, or relax terminal constraints such as aperture size and transmit power. At moderate data rates, coherent techniques further enable robust operation in power-limited and turbulence-dominated regimes where simpler intensity-modulation schemes may be insufficient. In FSO links, atmospheric turbulence causes signal fades lasting several milliseconds. This results in substantial data packet loss, especially when using transceivers designed for terrestrial fiber conditions, since they are not designed to handle severe and sudden changes in signal quality. The increasing availability of commercial off-the-shelf (COTS) high-speed analog-to-digital and digital-to-analog converters in the GHz range, driven by software-defined radio technologies in combination with field-programmable gate arrays (FPGAs), enables the application of coherent transceiver architectures at data rates well below those of conventional fiber-optic systems without the need for custom silicon development. This significantly reduces the development cost of custom transceivers with DSP tailored for the FSO channel and custom communication protocols, which would increase the robustness of the system during signal fades and could enable the integration of novel concepts such as time transfer and integrated sensing. In this work, we investigate the practical implementation aspects of coherent transceivers based on COTS FPGA platforms with a focus on the DSP, including carrier recovery, timing recovery, equalization and forward error correction (FEC). We present a real-time 2.5 GBaud coherent transceiver and evaluate its performance using a fading-channel testbench. Implementation challenges and design trade-offs relevant to coherent FSO systems are discussed. Finally, the system performance is assessed in terms of recovery time under representative atmospheric channel conditions.

Keywords: Coherent optical transceiver, optical satellite links, atmospheric turbulence, Digital signal processing, FEC

1. INTRODUCTION

Coherent optical communication systems at moderate data rates (1-10 Gb/s) are attractive for optical satellite links as they offer higher sensitivity than intensity-modulated techniques. The receiver sensitivity advantage of coherent detection allows system designers to increase link margin, improve availability, or reduce terminal aperture size and transmit power.¹ In many applications, these advantages could outweigh the additional implementation complexity associated with coherent communications.

For higher data rates (100+ Gb/s), coherent transceivers designed for terrestrial fiber networks are readily available and already integrated into in-orbit systems.² However, the implementation complexity of coherent optical transceivers has limited their commercial viability for moderate data rates. Hence, applications demanding coherent communications in this regime require tailored transceiver solutions. While the development of a custom transceiver is a complex undertaking, it allows optimized implementation to the unique requirements of satellite communication applications. For example, by implementing a custom digital signal processing (DSP)

Send correspondence to C.V firstname.lastname@dlr.de

stack, it could be possible to design a system that is more robust to scintillation caused by the atmospheric channel. Additionally, the development of custom transceivers also enables the implementation of application-specific communication protocols, opening the door to new applications or highly optimized systems for specialized use-cases.

Traditionally, the development of advanced transceiver architectures has required complex hardware designs combining discrete analog front ends and digital processing components. Modern field-programmable gate arrays (FPGAs) with integrated high-speed analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) now provide a highly integrated mixed-signal platform, enabling the development of flexible transceiver solutions for applications where hardware complexity previously represented a major limitation.³ These advances provide an opportunity to explore mixed-signal system-on-chip platforms as enablers of custom coherent transceiver solutions for optical satellite links at moderate data rates. An experimental demonstrator is therefore required to validate the feasibility of customized transceiver architectures, assess their performance under realistic conditions, and identify unforeseen implementation challenges. Beyond technical validation, such a demonstrator can provide insights into potential new applications and use cases.

This paper introduces a custom transceiver design for optical satellite links along with preliminary testing results. First, the implementation architecture of the custom transceiver is described, followed by initial performance assessment using a fading channel emulator to observe DSP chain behavior under dynamic channel conditions.

2. TRANSCEIVER IMPLEMENTATION

The focus of the paper is on the receiver, which accounts for the majority of the system complexity and therefore represents the most challenging aspect to the system design. An overview of a coherent optical receiver system is shown in Figure 1. In modern high-speed optical systems, DSP within coherent receivers takes responsibility for compensating transmission impairments, shifting both design complexity and performance bottlenecks to the DSP circuit. The receiver can be divided into three sequential functional blocks, optical, analog and finally the digital signal processing stage. The optical front-end first splits the incoming signal into its two orthogonal polarizations and coherently mixes each with a local oscillator (LO). Four balanced photodiodes then convert these optical fields into electrical baseband signals, yielding the in-phase (I) and quadrature (Q) components for both polarization states. These analog outputs are subsequently amplified before being digitized by four high-speed ADCs. Finally, the DSP block performs carrier frequency and phase recovery, adaptive polarization demultiplexing, symbol timing synchronization, and forward-error-correction decoding to reconstruct the original data stream.

While terrestrial systems typically target dual-polarization quadrature phase shift keying (DP-QPSK) modulation or even higher modulation formats, optical satellite links are often constrained by the link budget rather than spectral efficiency. Therefore, this paper investigates the implementation complexity of 2.5 GBaud single-polarization binary phase shift keying (SP-BPSK) transceivers as this reduces the complexity of the DSP. The main focus is on the FPGA receiver architecture which was developed with the objective of balancing communication performance and implementation complexity. While coherent DSP algorithms are well established, their practical realization depends strongly on system-level considerations, including laser and ADC performance as well as available FPGA resources, and the dynamics of the channel. The following section therefore focuses on the trade-off behind each processing stage, exploring how these constraints influence both the algorithmic choices and their hardware realization.

The four ADCs of the system sample the incoming signal at a sample rate of 5 GHz with a 14-bit resolution. A digital gain control stage then scales the signal to align its dynamic range within the word length of the DSP, ensuring optimal bit utilization before the following DSP stages. This is then followed by a polarization recovery stage, which restores the original state-of-polarization (SOP) by combining the two polarization branches into one signal. A timing recovery stage is then used to resample and synchronize the signal to the transmitter symbol rate. An equalization stage is then used to compensate for the channel impulse response. Next, the carrier recovery stage is used to compensate for the offset in carrier frequency and carrier phase between the transmitter and receiver laser. Finally, the receiver detects the frames and applies forward error correction (FEC) to decode




Figure 1. Overview of the coherent receiver system presented in this paper.

and correct errors. Although the transmitted signal is encoded in a single polarization, the received optical field generally does not remain aligned at the receiver. The implemented receiver therefore captures both polarization components using a dual-polarization coherent front-end and then reconstructs the signal. one polarization is transmitted, the objective differs from conventional dual-polarization coherent receivers, where a 2×2 MIMO equalizer is required to separate the two symbol streams. If the transmitted signal is given by $S(t)$, then the received electrical field can be described by,

$$E_X(t) = \sqrt{\alpha}S(t) \quad (1)$$

$$E_Y(t) = \sqrt{1 - \alpha}S(t) \exp(j\phi) \quad (2)$$

$$(3)$$

where α denotes the power splitting ratio and ϕ is the relative phase difference between the two polarization components. The maximum ratio combining solution to recover the original polarization state is given by,⁴

$$E_Z(t) = \sqrt{\alpha}E_X(t) + \sqrt{1 - \alpha}E_Y(t) \exp(-j\phi) \quad (4)$$

where α and ϕ can be estimated by,

$$\alpha = |E_X|^2 / (|E_X|^2 + |E_Y|^2) \quad (5)$$

$$\phi = \angle(E_X E_Y^*) \quad (6)$$

$$(7)$$

Unlike fiber-optic coherent systems, where random birefringence leads to rapid polarization rotations and polarization-mode dispersion, free-space optical propagation does not introduce comparable polarization mixing mechanisms. Consequently, polarization tracking requirements are expected to be reduced compared with conventional fiber coherent receivers enabling slower and more robust polarization tracking. Thus, α and θ is estimated by calculating $|E_X|^2, |E_Y|^2$ and $E_X E_Y^*$ at a fraction of the sample rate and then filtering with a first-order low-pass filter. Lookup tables are then used to determine the gain coefficients $\sqrt{\alpha}$ and $\sqrt{1 - \alpha}$ as well as phase offset ϕ . Additional care is required when one polarization dominates as this leads to an inaccurate estimation of the phase offset and as the weaker polarization mainly contributes noise. To prevent this, the system suspends the updates of ϕ when there is not sufficient power in one of the polarizations. polarization recovery is then followed by the clock recovery stage to align the symbol timing of the signal. Maintaining stable symbol synchronization is critical as the following DSP stages often rely on the symbols being synchronized. The dynamics of the atmospheric channel in combination with Doppler shifts makes this a challenging problem for optical satellite links compared to terrestrial fiber links as deep fades could cause the clock recovery system to diverge and Doppler introduces an additional clock frequency offset. Clock recovery could be achieved by fine tuning the ADC sample clock frequency, however this often requires a control loop from the DSP to the oscillator, resulting in both higher system complexity and higher loop delay. Instead, an all-digital solution allows

free-running oscillators and shorter update paths for the clock recovery loop, which could allow the system to operate with higher clock frequency offsets and higher noise tolerance. The receiver implements the Gardner timing recovery loop,⁵ where the timing error $e[n]$ is estimated by,

$$e[n] = \mathbb{R}((r[2n+1] + r[2n-1])r^*[n]) \quad (8)$$

where $r[n]$ is the resampled signal at sample position n . A loop filter is then used to filter the error signal. By implementing the loop filter with high resolution and adaptable gain parameters, the timing recovery loop can remain stable through varying channel conditions. The signal is finally resampled using a 4-tap Lagrange Cubic Interpolator.

Doppler shifts can introduce large clock frequency offsets which can go up to 50 ppm for inter-satellite links.⁶ However, when designing custom transceivers, the clock recovery requirements can be partially reduced by choosing a high-quality oscillator at the cost of SWaP. As a result, the loop bandwidth can be reduced, resulting in improved noise performance and better tracking through the atmospheric channel.

The clock recovery stage is followed by an equalization stage which compensates for the channel impulse response. The most common equalization technique for modulation formats such as BPSK and QPSK is the constant modulus algorithm (CMA).⁷ CMA has been reported to diverge during deep fades and it has been proposed to adapt the step size according to the channel conditions.⁸ However, for coherent optical satellite links operating at lower data rates, the channel remains relatively static due to minimal impact of polarization mode dispersion and residual chromatic dispersion, thereby diminishing the necessity for dynamic equalization. Additionally, Parallelized adaptive equalization algorithms are costly to implement on an FPGA. Instead, the receiver implements a 6-tap equalizer where the equalizer is only calibrated during initial link acquisition. This results in a low complexity equalizer as no equalizer update function needs to be implemented using FPGA resources.

Next follows the carrier synchronization stage, which is split into a carrier frequency estimation stage followed by a phase estimation stage. Carrier synchronization presents significant challenges at lower data rates in optical communication systems due to the much smaller electrical bandwidth compared to the carrier frequency. For example, Doppler can introduce up to 10 GHz carrier frequency offsets in inter-satellite links.⁶ However, by using tunable lasers the frequency offset can often be coarsely tuned to align the signal within the bandwidth of the receiver. As a result the DSP only needs to compensate for the residual frequency offsets. Additionally, since the satellite orbit is typically known a priori, the transmit side could also perform coarse tuning to reduce the required tunable frequency range of the receiver laser. implements a feedforward frequency offset algorithm, where the estimated frequency offset f_{est} over an integration window of N symbols is given by,⁹

$$f_{\text{est}} = \frac{1}{4\pi T_s} \angle \left\{ \sum_{n=0}^{N-1} r[n]^2 (r[n-1]^2)^* \right\}, \quad (9)$$

where $r[n]$ denotes the received symbol and T_s is the symbol duration. Further a first-order low-pass filter is then used to provide additional filtering of the estimate. The number of symbols that can be used to for the frequency offset estimation is determined by the frequency stability of the lasers used in the system, thus by designing a custom transceiver a trade-off can be made between the laser size and power consumption and the DSP frequency offset tracking. Following the carrier frequency compensation, residual phase noise remains due to laser phase fluctuations and imperfect carrier frequency offset compensation. The receiver implements the phase estimation using the blind Viterbi-Viterbi algorithm¹⁰ where the phase θ is estimated by,

$$\theta = \frac{1}{2} \angle \sum_{n=0}^N r[n]^2 \quad (10)$$

where $r[n]$ is the input samples to the phase estimator. However, this results in a phase ambiguity problem as the receiver cannot tell the two phases apart. One method to circumvent this is to use differential encoding of the symbols, but this comes at the cost of a sensitivity penalty.⁶ Another solution is to periodically insert pilot symbols to provide an absolute phase reference. This is a common technique used in coherent optical

communications and is promising technique to achieve phase synchronization in optical satellite links. However, in this receiver another technique is explored where only the frame alignment preamble is used to align the absolute carrier phase. This method is sensitive to cycle slips but does on the other hand not require the overhead of periodically inserted pilots and avoid the sensitivity penalty from differential encoding. The frame alignment is achieved by using a predetermined preamble sequence of 64 symbols. The cross-correlation between the received signal and the known preamble is calculated and the position of the frame is determined by finding the maximum of the correlation. To determine whether a valid frame boundary exists, a hypothesis test is performed by comparing the correlation output against a predefined detection threshold. It is important to note that optical satellite links typically employ continuous transmission with a fixed frame structure. As a result, the temporal spacing between consecutive frame headers remains constant. This structural regularity implies that only an initial frame acquisition is strictly necessary. Once synchronization is established, subsequent frame boundaries can be reliably tracked by symbol counting as long as the clock recovery remains locked. It is therefore not necessary to perform repeated correlation-based searches making reliable frame tracking possible even in fading channels.

FEC is indispensable for optical satellite communications due to the severely degraded channel conditions inherent to space-based links. Implementing advanced FEC schemes at gigabit data rates introduces significant computational complexity, which results in substantial resource consumption when deployed on FPGA fabrics. To mitigate these issues, modern FPGA architectures increasingly provide FEC encoder/decoder blocks as dedicated hard macros, dramatically reducing area usage and power consumption. Although, this evolution is mainly driven by the development of 5G equipment, optical satellite links could benefit tremendously by utilizing similar codes. In fact, the Space Development Agency’s Optical Communications Terminal Standard (SDA OCT) already takes advantage of 5G New Radio (NR) Low-Density Parity-Check (LDPC). For example, the AMD RFSoc platform has dedicated blocks for 5G NR LPDC encoding and decoding with a supported throughput of up to 3 Gbit/s.¹¹ The AMD RFSoc platform features a total of 6 decoding blocks such that transceivers with 10 Gbit/s throughput could potentially be achieved through parallelization of the decoders.

The complete system architecture is implemented on an AMD RFSoc ZU48DR platform, which provides a high-performance FPGA fabric integrated with advanced digital upconversion and downconversion blocks as well as integrated LDPC encoder/decoders. The data path employs an input parallelization factor of 16 samples at a clock frequency of 312.5 MHz, equivalent to eight complex symbols processed per clock cycle throughout the DSP, where the input data from the ADC has been quantized to 7 bits. Manchester encoding and decoding was applied before framing and deframing, halving the final bit rate. The hardware resource consumption is summarized in Table 1, which details the utilization of lookup tables (LUT), flip-flops (FF), DSP slices, and block random access memories (BRAMs) across the entire design.

Table 1: Resource utilization of both the transmitter and receiver DSP for an AMD RFSoc ZU48DR FPGA.

Resource	Used	Available	Utilization
LUT	91,347	425,280	21.48%
FF	82,614	850,560	9.71%
BRAM	129	1,080	11.94%
DSP	477	4,272	11.17%




Figure 2. Overview of experimental setup. An RFSoc with integrated DACs generates the baseband signal which is modulated onto an optical carrier. Subsequently, a fading testbed emulates dynamic channel attenuation, after which a coherent receiver captures the signal. Finally, the RFSoc’s integrated ADCs sample the received signal and DSP recovers the transmitted data.

3. EXPERIMENTAL SETUP AND RESULTS

The experimental setup utilized is illustrated in Figure 2. The setup is built around an AMD RFSoc with integrated ADCs and DACs dedicated to transmitter and receiver DSP as well as forward error correction. The DAC generated the transmission signal at 5 GSPS, which was amplified to perform BPSK modulation of a nano-integrated tunable laser assembly (Nano-ITLA) with a specified linewidth of 300 kHz. The fading testbed attenuated the signal using multiple variable optical attenuators (VOA) stages,¹² followed by an EDFA pre-amplifier. Finally, electro-optical conversion was done using a coherent receiver and a Nano-ITLA before the ADCs sample the signal and the DSP is performed.

First, the impact of deep signal fades was evaluated by applying rectangular deep fades to the system. The received power before the EDFA was set -57 dBm for 40 ms followed by 10 ms at -67 dBm and the pattern was then repeated. The system keeps track of the arrival time of each frame as well as DSP status parameters. Additionally, frame numbers and transmit timestamps were included into the frame header to validate the results. The frame rate was set to approximately 80 kHz to achieve high resolution of the recovery time. In Figure 3, the measurements of the frame detection cross-correlator at the frame start, the estimated carrier frequency offset, the normalized timing offset deviation as well as the normalized phase offset between the two polarizations is shown for one deep fade measurement. It can be seen that once the fade starts, the magnitude of the correlation drops drastically for the duration of the fade. Note, that once the signal comes back the sign of the correlation inverts, indicating that a carrier phase cycle slip occurred during the fade. For the clock recovery system, the fade results in it slowly starting to deviate from the symbol center. However, the clock recovery loop remains within one sample error during the deep fade. The carrier frequency offset estimation algorithm on the other hand results in noisy estimates during the fade making it impossible for the receiver to track the carrier frequency and phase. Finally, the phase offset between the two polarizations is not affected as much by the fade as the phase can be averaged over a very large number of symbols. By processing multiple fade measurements, it is possible to gain insights into the statistics of the recovery time of the receiver. The outage time for each fade was calculated as the frame which the communication links was lost until frames which could be decoded were received again. The resulting histogram is shown in Figure 4, where the distribution of the total outage duration can be observed. The average outage lasted approximately 10.25 ms, thus resulting in a typical recovery time of 0.25 ms. Further, 1 ms fades were also tested to evaluate the impact the fade duration has on the recovery time. The resulting outage duration histogram is shown in Figure 5. Notably, the outage statistics remain similar between 1 ms and 10 ms fades suggesting that the slow drift from the clock recovery is not the main contributor to the recovery time.




Figure 3. Tracked DSP parameters during a deep signal fade. The following parameters are shown, the cross-correlation value at the start of each frame, the estimated carrier frequency offset, estimated symbol timing offset normalized to samples as well as the estimated phase offset between the two received polarizations




Figure 4. Histogram showing the outage duration distribution of 10 ms fades.




Figure 5. Histogram showing the outage duration distribution of 1 ms fades.

4. CONCLUSION

A 2.5 GBaud SP-BPSK transceiver was implemented with DSP customized for atmospheric links. The recovery time after deep fades was evaluated by testing the system performance with rectangular deep fades. The receiver showed an average recovery time of approximately 0.25 ms after 10 ms deep fades. In this case, the recovery time of the system would only introduce a minor penalty to the overall link budget. However, similar recovery time statistics were observed for 1 ms fades, suggesting that the recovery time is an important system design parameter to consider. For example, the timing recovery block as well as the carrier frequency recovery block could be implemented using more complex estimation algorithms to potentially further reduce the recovery time of the system. The carrier frequency recovery block is especially difficult to estimate during fades due to the rapid changes in carrier frequency. Thus, for systems aiming to achieve robust frequency synchronization through fades, a more stable laser source can substantially improve the overall system performance. The transceiver was implemented on a single FPGA requiring only about 10% DSP utilization and 20% LUT utilization, showing that custom coherent optical communication systems can be implemented without any major issues on modern FPGAs. Further, the low resource utilization shows promising results for expanding the system. For example, higher ADC quantization could be employed or more complex algorithms could be considered to improve the system performance. Higher data rates could also be achieved by either increasing the data rate or by implementing higher modulation formats and utilizing both polarizations. To conclude, transceivers implemented by utilizing FPGAs could enable rapid adoption of coherent optical satellite links at moderate data rates. The RFSoc platforms are especially attractive for this as they offer ADC, DAC, FEC and an FPGA in a single package, reducing the system complexity and creating the possibility to deploy highly customized coherent optical transceivers.

REFERENCES

- [1] Geisler, D. J., “Coherent architectures for free-space optical communications,” in *[2020 Conference on Lasers and Electro-Optics (CLEO)]*, 1(2), 1–2 (2020).
- [2] Schieler, C. M., Riesing, K. M., Bilyeu, B. C., Chang, J. S., Garg, A. S., Gilbert, N. J., Horvath, A. J., Reeve, R. S., Robinson, B. S., Wang, J. P., Piazzolla, S., Roberts, W. T., Kovalik, J. M., and Keer, B., “On-orbit demonstration of 200-Gbps laser communication downlink from the TBIRD CubeSat,” in *[Free-Space Laser Communications XXXV]*, Hemmati, H. and Robinson, B. S., eds., **12413**, 1241302, International Society for Optics and Photonics, SPIE (2023).
- [3] Farley, B., McGrath, J., and Erdmann, C., “An all-programmable 16-nm rfsoc for digital-rf communications,” *IEEE Micro* **38**(2), 61–71 (2018).
- [4] Costa, P. S., Drummond, M. V., Ribeiro, R. S., Oliveira, C., Ribeiro, N., Nogueira, R., and Monteiro, P. P., “Experimental evaluation of a polarization tracking algorithm for single-polarization m-(d)psk signals using coherent detection,” in *[2012 14th International Conference on Transparent Optical Networks (ICTON)]*, 1–4 (2012).

- [5] Gardner, F., “A bpsk/qpsk timing-error detector for sampled receivers,” *IEEE Transactions on Communications* **34**(5), 423–429 (1986).
- [6] Valjus, C., Wolf, R., and Poliak, J., “Review and analysis of digital signal processing algorithms for coherent optical satellite links,” *International Journal of Satellite Communications and Networking* **43**(3), 229–250 (2025).
- [7] Spinnler, B., Hauske, F., and Kuschnerov, M., “Adaptive equalizer complexity in coherent optical receivers,” in *[2008 34th European Conference on Optical Communication]*, 1–2 (2008).
- [8] Matsuda, K., Koshikawa, S., Yoshida, T., Suzuki, N., and Ando, T., “FPGA implementation of scintillation tolerant adaptive DSP for 4 Gbps coherent reception,” in *[Free-Space Laser Communications XXXII]*, Hemmati, H. and Boroson, D. M., eds., **11272**, 149 – 157, International Society for Optics and Photonics, SPIE (2020).
- [9] Leven, A., Kaneda, N., Koc, U.-V., and Chen, Y.-K., “Frequency estimation in intradyne reception,” *IEEE Photonics Technology Letters* **19**(6), 366–368 (2007).
- [10] Viterbi, A. J. and Viterbi, A. M., “Nonlinear estimation of psk-modulated carrier phase with application to burst digital transmission,” *IEEE Transactions on Information Theory* **29**(4), 543–551 (1983).
- [11] Finnerty, A. and Lee, M., “Integrated sd-fec in zynq ultrascale+ rfsocs for higher throughput and power efficiency,” Tech. Rep. WP498, Xilinx, Inc. (May 2018).
- [12] Raffa, S., La Torre, G., Reger, L., Torres, J. S., Valjus, C., and Poliak, J., “Optical feeder link system design via simulation and emulation tools,” in *[2025 IEEE International Conference on Space Optical Systems and Applications (ICSOS)]*, 1–4 (2025).