sensors

MDPI

Article

Phase-Combiner Sensitivity and Paired-Error Branch-Tracking Robustness in Alamouti Coherent-Lite Receivers

Mahmoud Alhalabi [1,*], Temel Sönmezocak [1] and Fady I. El-Nahal [2,3,4]

[1] Department of Electrical and Electronics Engineering, Istanbul Yeni Yuzyil University, Istanbul 34000, Türkiye
[2] Institute for Communications Engineering, Technical University of Munich (TUM), Theresienstrasse 90, 80333 Munich, Germany
[3] Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK
[4] Department of Electrical Engineering, Islamic University of Gaza, Gaza P.O. Box 108, Palestine
* Correspondence: mahmoud.alhalabi@yeniyuzyil.edu.tr

Abstract

Broad-linewidth lasers challenge Alamouti coherent-lite receivers because carrier recovery combines conjugated branch terms from adjacent time slots. This work compares single-error and projected paired-error complex branch-correction rules in a four-branch Alamouti coherent-lite receiver using identical frozen FIR branch streams. To avoid the Nyquist-edge ambiguity of the 100 gigasamples per second (GSa/s) reference model, the key experiments are repeated with a 200 GSa/s front end. The 200 GSa/s results form the principal basis for the robustness conclusions, while the 100 GSa/s results are retained as reference-model diagnostics. Across two independent static, lossless Jones-matrix ensembles at 2 MHz per laser, threshold failures decrease from 21/240 with single-error tracking to 2/240 with paired-error tracking, and the paired rule gives lower BER in 217/240 matched records. FIR-training controls show that training quality materially affects both trackers, so the receiver-level improvement cannot be attributed solely to the downstream paired-error correction. These conclusions are limited to the simulated static, lossless SOP transformations and idealized receiver model considered here.

Keywords: carrier recovery; coherent-lite receiver; heterodyne detection; passive optical network; wide-linewidth laser

1. Introduction

The continuing growth of residential, enterprise, mobile-fronthaul, and cloud-access traffic is pushing passive optical networks (PONs) beyond the capacity range in which conventional intensity-modulation/direct-detection links can simultaneously provide a high line rate, large splitting loss, and long reach. Coherent PON architectures are increasingly being considered for high-capacity access networks, including 100G/λ-and-beyond and emerging 6G systems [1–3]. Earlier work also examined coherent PON transport for 5G and beyond [4,5]. Coherent reception is attractive in this setting because it offers high sensitivity, wavelength selectivity, linear field detection, and digital compensation of channel impairments. Reviews and recent tutorials, however, consistently identify optical-network-unit (ONU) cost, local-oscillator control, analogue bandwidth, data-converter power, and burst acquisition as the main barriers to widespread coherent access deployment [1,2,6].

A practical coherent PON must therefore reduce the optical front end and the associated digital workload, since direct transfer of a long-haul receiver architecture to the access

Check for updates

Academic Editor: Nageswara Lalam

Received: 16 October 2026
Revised: 13 November 2026
Accepted: 25 November 2026
Published: 27 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

network would retain unnecessary complexity. Digital signal processing (DSP)-enabled detection and time/frequency-division multiple-access architectures illustrate how digital processing can trade hardware complexity for sensitivity and flexible resource allocation [7,8]. This trade-off is particularly important at the ONU, where photodiode count, optical-hybrid complexity, analogue-to-digital converter (ADC) bandwidth, power consumption, and calibration burden directly affect the feasibility of high-volume deployment.

The 0.5–2 MHz stress range is intended to represent low-cost wide-linewidth laser operation rather than narrow-linewidth long-haul practice. Experimental short-reach coherent work has evaluated 1 MHz and 3 MHz distributed-feedback (DFB) lasers and demonstrated 1 Tb/s 256-QAM transmission with a 3 MHz DFB source, explicitly motivating MHz-class phase recovery for cost-sensitive coherent links [9]. For a directly comparable Alamouti/PTBC receiver context, Oh et al. reported sensitivity penalties of 0.37 dB and 0.79 dB at 1 MHz and 2 MHz Tx/Rx laser linewidths, respectively, in a 100 Gb/s 16-QAM system [10]; the different symbol rate and receiver architecture preclude treating those values as a prediction for the present receiver. In the Alamouti receiver the difficulty is structural: the two source symbols of a pair occupy adjacent symbol slots and enter the decoder through conjugated relations, so phase accumulated over one symbol period affects the odd and even branch combinations differently. In a 100 Gb/s/λ Alamouti PON demonstration, Fan et al. used a DFB local oscillator with linewidth above 1 MHz at the ONU. After 25 km, the measured optical power budgets were 33.9 dB with that DFB and 34.2 dB with an external-cavity laser [11]. In a separate 200G PON component comparison, Fang et al. allocated approximately 1.2 W to thermoelectric cooling in temperature-controlled laser configurations [12]. These figures are specific to the reported devices and architectures. Laser linewidth, electrical power, and component pricing depend on packaging, temperature control, drive electronics, and production volume.

Alamouti space-time block coding (STBC), originally introduced for wireless transmit diversity [13], provides one route to polarisation-insensitive coherent-lite reception. In optical access systems, Alamouti coding can be combined with heterodyne detection so that two transmitted polarisations are reconstructed using a single-polarisation optical front end and a single balanced photodetector. The resulting architecture has been demonstrated as a candidate for coherent PONs beyond 50 Gb/s/λ [14]. Its reduced optical complexity is obtained at the cost of a distinctive DSP problem: successive symbols are transmitted through conjugated odd- and even-slot branch relations, and the receiver must combine those branches while tracking laser phase noise.

Let s_i, $i = 0, 1, \ldots$, denote the Gray-coded 16-QAM source symbols. Symbol and Alamouti-pair indices are zero-based: pair k contains s_{2k} and s_{2k+1}. The sequences transmitted on the two optical polarisations, denoted x and y, are

$$x = [s_0, -s_1^*, s_2, -s_3^*, \ldots], \tag{1}$$

$$y = [s_1, s_0^*, s_3, s_2^*, \ldots]. \tag{2}$$

where $(\cdot)^*$ denotes complex conjugation. Conventional coherent DSP provides adaptive multiple-input multiple-output (MIMO) equalisation and a range of feed-forward or decision-directed carrier-recovery strategies [15,16]. Those methods are well established for quadrature amplitude modulation (QAM) streams, but the conjugated two-slot Alamouti structure changes how phase-correction states enter the odd and even decisions. A tracker that updates two complex branch states from only one recovered-symbol error can therefore behave differently after state-of-polarisation (SOP) mixing, even when the

underlying finite-impulse-response (FIR) branches are common and their coefficients are held fixed after training.

Previous work has established coherent-lite front ends, Alamouti decoding, wide-linewidth carrier recovery, and high-rate transmission through access fibre. Oh et al. [10] already use both recovered-symbol errors and the conjugate-consistent projection $p = (p_1 + p_2^*)/2$ in a simplified two-filter PTBC decoder. Here, we retain four separate FIR coefficient banks from the Faruk receiver architecture [17] and compare single-error and averaged paired-error downstream updates on identical fixed branch streams. The present work therefore focuses on a controlled four-branch comparison on identical fixed branch streams, with tracker-neutral training controls used to separate the downstream correction effect from FIR-training assistance. Section 3.2 gives an explicit relation to the shared-filter special case of Oh et al. [10]

The present study addresses two receiver-DSP questions: (i) how the single-error correction rule of Faruk et al. [17] and a projected paired-error extension compare on identical trained branch streams that are held fixed during the comparison, and (ii) how the Alamouti pair combiner changes when a common pair correction is relaxed to independent slot corrections in a known-symbol pair-level diagnostic. The latter is used only to probe combiner sensitivity and is not treated as a practical receiver.

This work makes four primary contributions:

1. We compare single-error and projected paired-error correction on identical fixed four-branch FIR streams and use tracker-neutral training controls to isolate the downstream correction effect.
2. We extend the comparison to 20 km linear standard single-mode fibre (SSMF) and a received-optical-power sweep, including transmitter (Tx)-only and local-oscillator (LO)-only phase-noise controls to assess effects compatible with equalisation-enhanced phase noise (EEPN).
3. We use known-symbol one-phase and two-phase pair-level searches as a secondary diagnostic of phase-combiner sensitivity and verify numerical stability by grid refinement.
4. We add a 200 GSa/s sampling control, expanded independent-seed/static-Jones statistics, phase-noise-free duplicate-record and same-record true-phase-trajectory FIR-training controls, a blind-phase-search reference, and power-restored Tx-/LO-only fibre controls; these controls provide the basis for the robustness assessment.

Unless otherwise stated, the 100 GSa/s reference-model linewidth/SOP and fibre analyses use ten independent realisations per condition; the 100 GSa/s reference-model received-optical-power (ROP) robustness study uses 30 independent realisations per received-power point, and the disjoint tracker-tuning set uses a separate ten realisations. The 200 GSa/s 2 MHz principal-SOP control uses 20 independent stochastic seeds per angle; the corrected Jones replication crosses ten seeds with two independently generated 12-matrix ensembles.

Related Work and Scope

System-level coherent-access research spans wavelength-division-multiplexed (WDM) PONs, 200 Gb/s/λ coherent-detection links, and simplified bidirectional ONUs. Bidirectional intensity-modulated/direct-detection optical orthogonal frequency-division multiplexing (OFDM) WDM-PON systems represent a lower-complexity direction for access networks. One study simulated a long-reach link with dispersion-compensating fibre [18], while a later study evaluated a 30 km link without it [19]. Coherent optical OFDM simulations have also examined QPSK and 16-QAM over 500 km of uncompensated fibre [20]. More recently, Taşpınar and Alhalabi investigated a 40 Gb/s CO-OFDM system with

optimization-assisted partial transmission sequence processing for PAPR reduction [21]. The present work focuses on coherent-lite phase tracking.

Real-time coherent ultra-dense WDM (UDWDM)-PON transmission has been demonstrated over 50 km with a large power budget, while 200 Gb/s/λ polarisation-division-multiplexed four-level pulse-amplitude modulation (PDM-PAM-4) coherent detection has been shown over 20 km. More recent simplified-ONU studies report bidirectional 20–40 km operation and 200 Gb/s coherent PON reception with reduced-complexity front ends [22–25]. These works establish system feasibility, reach, and power budget. The present study focuses on the subsequent branch-correction update on common fixed branch streams, with a separate tracker-neutral control used to quantify training dependence.

The most directly related receiver family combines Alamouti coding with heterodyne detection. Single-balanced-photodiode long-reach WDM-PON reception and dedicated DSP for single-carrier Alamouti signals established the optical and algorithmic foundation. Parallel work on single-ended and minimal coherent receivers reduced the front-end component count through different signal and detection assumptions. Experimental 100/200 Gb/s Alamouti-coded downstream links further showed that simplified coherent receivers can operate over access-fibre distances [17,26–29]. A recent 50 GBaud, 20 km short-preamble study further used a four-branch gear-shifted LMS equalizer with a single-balanced-detector/one-ADC Alamouti coherent-lite ONU, focusing on acquisition and preamble-length trade-offs rather than phase-combiner robustness [30]. In contrast, the present paper holds the trained FIR branch extractor fixed and compares only the subsequent complex correction rules.

Carrier-phase estimation (CPE) for coherent QAM includes blind phase search, pilot-aided estimation, blind joint equalisation and phase recovery, linewidth-tolerant reduced-complexity feed-forward estimators, joint- polarisation estimation, and soft-decision-aided maximum-likelihood methods [31–36]. These algorithms provide important benchmarks for linewidth tolerance and complexity. Their primary formulations operate on conventional QAM streams or exploit shared-polarisation/subchannel statistics, while the present receiver uses conjugated two-slot Alamouti branch equations. Recent residual-carrier-aided coherent-lite work confirms that multi-megahertz linewidths can be addressed with low-complexity decision-directed tracking [37]. A direct algorithm-by-algorithm comparison with conventional CPE would require adapting each method to the conjugated two-slot Alamouti branch equations and retuning its acquisition and phase-recovery parameters. The present study therefore fixes the branch extraction and focuses on the single-error versus paired-error correction update. Section 3.3.3 includes a 200 GSa/s Alamouti-adapted blind phase search (BPS) benchmark as a practical CPE reference.

The 20 km extension considers deterministic linear SSMF with verified chromatic-dispersion compensation. Its fibre model includes attenuation and chromatic dispersion but not polarisation-mode dispersion (PMD) or Kerr nonlinearity; optical-amplifier noise, burst-mode acquisition, and experimental hardware are also outside this extension. Relative to system-level coherent-PON and simplified-ONU work [22–25,27,28], the present study isolates the subsequent four-branch correction update; the Alamouti-PON literature [17,26,29] provides the receiver foundation.

2. System Model and Methods

This section describes the transmitter architecture, channel model, receiver signal processing, bit-error-rate (BER) evaluation window, and construction of independent stochastic realizations used throughout the study. All numerical simulations, signal processing algorithms, and system performance evaluations were implemented using custom-developed Python code (Python 3.15.2). No commercial simulation software or experimental hardware

was employed. To ensure consistency and reproducibility, the same branch-training procedure and BER evaluation window were maintained across the phase-combiner diagnostic, tracker comparison, and fibre-extension experiments.

2.1. Transmitter and Channel

The transmitter generates 2^{15} Gray-coded 16-QAM symbols. Alamouti STBC encodes successive pairs according to Equations (1) and (2). Each polarisation is shaped by a root-raised-cosine (RRC) filter with roll-off 0.01 and span 128 symbols at 2 samples/symbol. The symbol rate is 50 GBaud, and transmitter quantisation is represented by a 6-bit effective number of bits (ENOB). Laser phase noise is modelled as a discrete-time Wiener process [15],

$$\phi[n] = \phi[n-1] + \Delta\phi[n], \qquad \Delta\phi[n] \sim \mathcal{N}\left(0, \frac{2\pi\Delta\nu}{f_s}\right). \tag{3}$$

where $\phi[n]$ is the laser phase at sample n, $\Delta\phi[n]$ is its sample-to-sample phase increment, and $\mathcal{N}(\mu, \sigma^2)$ denotes a Gaussian distribution with mean μ and variance σ^2. The quantities $\Delta\nu$ and f_s denote the laser linewidth and sampling frequency, respectively. Independent phase-noise processes are generated for the transmitter and local oscillator. Let $\Delta\nu_{\mathrm{TX}}$ and $\Delta\nu_{\mathrm{LO}}$ denote the transmitter- and local-oscillator-laser linewidths, respectively. Unless otherwise stated, the linewidth values reported throughout this paper are per-laser values with $\Delta\nu_{\mathrm{TX}} = \Delta\nu_{\mathrm{LO}} = \Delta\nu$. The corresponding relative-phase process therefore has a total linewidth contribution $\Delta\nu_{\mathrm{TX}} + \Delta\nu_{\mathrm{LO}} = 2\Delta\nu$, and its sample-to-sample phase increment has variance $2\pi(\Delta\nu_{\mathrm{TX}} + \Delta\nu_{\mathrm{LO}})/f_s$. The principal sweep applies the real SOP rotation

$$\mathbf{H}(\theta) = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}, \tag{4}$$

where θ is the SOP rotation angle; the principal sweep uses $\theta \in \{0, \pi/8, \pi/4\}$. A broader static-SOP validation additionally uses lossless Jones matrices

$$\mathbf{J} = \begin{bmatrix} a & b \\ -b^* & a^* \end{bmatrix}, \qquad |a|^2 + |b|^2 = 1, \tag{5}$$

where $a, b \in \mathbb{C}$ are normalized complex Jones coefficients. These matrices span the three physically relevant degrees of freedom of a unitary 2×2 Jones transformation after removal of common phase. Twelve fixed Haar-distributed special-unitary $\mathrm{SU}(2)$ matrices, generated from normalized complex-Gaussian pairs (a, b), are kept identical across both trackers and the ten benchmark realisations. They are evaluated at 500 kHz, 1 MHz, and 2 MHz, giving 360 matched realisation–linewidth–Jones records. No optical amplification is applied; amplified-spontaneous-emission (ASE) noise is therefore absent. The baseline excludes fibre propagation. The 20 km extension uses linear SSMF at 1546.119 nm with 0.2 dB/km loss, $D = 17$ ps/(nm km), and $S = 0.08$ ps/(nm^2 km), where D is the chromatic-dispersion coefficient and S is the dispersion slope; polarisation-mode dispersion (PMD) and Kerr nonlinearity remain disabled.

2.2. Receiver and BER Window

The single balanced photodetector (BPD) heterodyne receiver uses a local oscillator (LO) offset of $1.01R_s/2$, where R_s = 50 GBaud is the symbol rate, and includes independent LO phase noise. The launch and LO powers are 0 dBm and 15.75 dBm, respectively, and the photodiode responsivity is 0.5 A/W. Shot and thermal detector-current noise are generated on the same 100 GSa/s grid as the received optical field before ADC-rate resampling. We use one-sided current-noise power spectral densities. Let $f_{s,\mathrm{det}}$ = 100 GSa/s denote the

detector-current sampling rate. Because the detector-current grid has Nyquist bandwidth $B_{\text{gen}} = f_{s,\text{det}}/2 = 50$ GHz, the per-arm discrete sample variances are

$$\sigma^2_{\text{th},i} = \frac{4k_B T B_{\text{gen}}}{R_L}, \qquad \sigma^2_{\text{sh},i} = 2q\mathcal{R}P_i B_{\text{gen}}, \tag{6}$$

where $\sigma^2_{\text{th},i}$ and $\sigma^2_{\text{sh},i}$ denote the thermal- and shot-noise sample variances in detector arm i, k_B is Boltzmann's constant, $R_L = 50\ \Omega$ is the load resistance, $T = 298.15$ K (25 °C) is the receiver absolute temperature, q is the electron charge, $\mathcal{R}$ is the photodiode responsivity, $i \in \{1, 2\}$ indexes the two photodiodes, P_i is the optical power incident on photodiode i. Independent noise from the two photodiodes is subtracted with the balanced currents, giving $\sigma^2_{\text{th,bal}} = 8k_B T B_{\text{gen}}/R_L$ and $\sigma^2_{\text{sh,bal}} = 2q\mathcal{R}(P_1 + P_2)B_{\text{gen}}$ before the subsequent multi-rate filter chain. Thus, the 50 GHz quantity in Equation (6) is the noise-generation Nyquist band, not a second post-filter equivalent-noise bandwidth.

The balanced current is then resampled from 100 GSa/s to 200 GSa/s, filtered by a fourth-order 50 GHz Butterworth low-pass filter (LPF), and resampled back to 100 GSa/s. An impulse-response energy calculation for this exact linear cascade gives a discrete white-noise variance gain $G_n = 0.9062$; a two-million-sample Monte Carlo check gives 0.9058. Hence the implemented post-chain effective one-sided noise bandwidth is $G_n B_{\text{gen}} = 45.31$ GHz. This explicit check confirms that the noise variance is not integrated over 50 GHz and then inadvertently integrated a second time by the LPF; the LPF intentionally shapes the already sampled noise. At a representative 0 dBm signal power and 15.75 dBm LO, the pre-filter balanced thermal and shot variances are 3.292×10^{-11} and 3.091×10^{-10} A^2, respectively, and the theoretical post-chain values are 2.983×10^{-11} and 2.801×10^{-10} A^2; the corresponding Monte Carlo values are 2.982×10^{-11} and 2.799×10^{-10} A^2. This agreement provides a direct implementation check of the discrete-time noise-bandwidth mapping. Shot-noise variance is recomputed as the received optical power changes.

The initial detector sampling grid imposes a separate limitation. With the stated RRC roll-off $\alpha = 0.01$, the nominal ideal-RRC baseband support extends to $(1 + \alpha)R_s/2 = 25.25$ GHz. At the chosen heterodyne offset, the real detector spectrum therefore reaches 50.5 GHz, slightly exceeding the 50 GHz Nyquist limit of the initial 100 GSa/s grid. Consequently, a narrow portion of the outer RRC transition region may fold at the first sampling operation, and subsequent interpolation to 200 GSa/s cannot restore those components. This front-end sampling effect is common to both trackers because they are compared on identical sampled branch streams; the reported comparisons therefore apply to the specified discrete-time model.

Because the two trackers can respond differently to that shared distortion, the 200 GSa/s validation in Section 3.3 is used as the principal basis for robustness conclusions. The 100 GSa/s results are retained as reference-model diagnostics and to document sensitivity to the front-end sampling model.

The Tx laser phase process, LO phase process, and detector-noise samples are all generated at 100 GSa/s. Only the real balanced detector current is resampled to 200 GSa/s for the Butterworth LPF, after which it returns to 100 GSa/s; the 6-bit effective-number-of-bits (ENOB)-equivalent noise is then added on that 100 GSa/s grid. ENOB is represented by additive zero-mean Gaussian quantization-equivalent noise. The following variance and SNR expressions are specified for each real-valued component sequence $z[n]$. The peak-to-average power ratio (PAPR) is $\text{PAPR} = \max|z[n]|^2/\mathbb{E}|z[n]|^2$, where $\mathbb{E}$ denotes the sample average over the complete processed record. With full-scale peak A, an ideal uniform N-bit quantizer has noise variance $A^2/(3\,2^{2N})$. In the ENOB-equivalent model, N_{ENOB} denotes the effective bit depth, giving the arbitrary-waveform component signal-to-noise ratio (SNR) $6.02N_{\text{ENOB}} + 10\log_{10}3 - 10\log_{10}(\text{PAPR})$, where $10\log_{10}3 = 4.77$ dB. For an ideal N-

bit quantizer and a sinusoid with 3.01 dB PAPR this reduces to the familiar $6.02N + 1.76$ dB expression. In the 100 GSa/s reference implementation, the transmitter ENOB-equivalent noise is applied to the shaped transmitter waveform before optical propagation, while the receiver ENOB-equivalent noise is applied to the real post-LPF 100 GSa/s detector-current samples before digital downconversion. Its variance is recomputed per record from the corresponding component power/PAPR. No hard quantizer transfer function, clipping threshold, ADC full-scale value, or automatic-gain-control (AGC) loop is modelled. Known digital mixing removes the configured frequency offset, so residual frequency-offset estimation is not evaluated. For the 20 km extension, frequency-domain chromatic-dispersion compensation (CDC) is applied to the complex baseband before matched filtering and synchronisation. The simplified end-to-end system and receiver/DSP chain are shown in Figure 1.

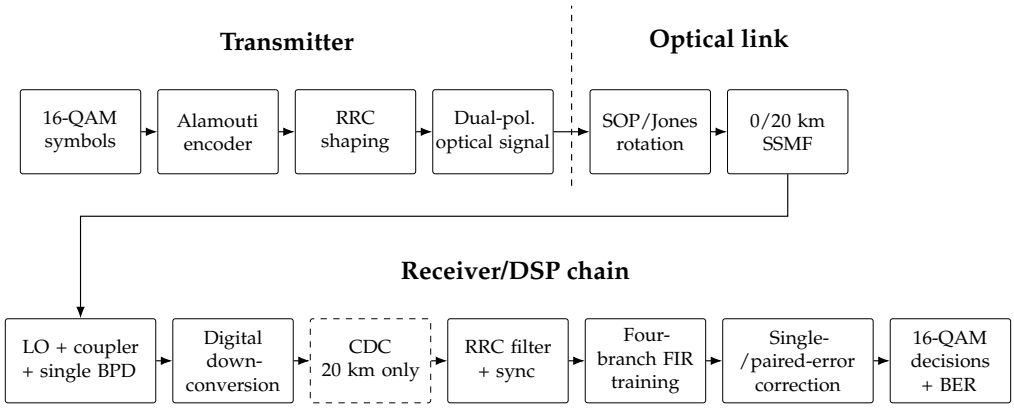

Figure 1. Simplified simulated Alamouti coherent-lite system. The optical link includes static SOP/Jones rotation and 0/20 km SSMF; the receiver uses a local oscillator (LO), optical coupler, and single balanced photodetector (BPD), with chromatic-dispersion compensation (CDC) only for the 20 km case. Both trackers operate on the same fixed four-branch FIR streams.

In the 200 GSa/s control model, the RRC-shaped dual-polarisation optical field, Tx and LO Wiener phase processes, balanced-photodetector current, 50 GHz fourth-order Butterworth filter, receiver ENOB-equivalent impairment, and digital downconversion are all represented natively at 200 GSa/s. Complex baseband is reduced to 100 GSa/s only after matched filtering. With $R_s = 50$ GBaud, $\alpha = 0.01$, and the unchanged 25.25 GHz heterodyne offset, the approximately 50.5 GHz upper signal edge is therefore below the 100 GHz Nyquist frequency. Receiver thermal and shot noise are generated on the 200 GSa/s grid and passed through the same Butterworth filter. The filter has a discrete white-noise variance gain of 0.5. At the numerical check point, the analytical post-filter variance is 3.380×10^{-10} A^2, and a two-million-sample Monte Carlo calculation gives 3.387×10^{-10} A^2, about 0.2% higher. The same data and seed identities are retained when the 100 GSa/s and 200 GSa/s models are compared.

Frame timing correlates both coded polarisation references after RRC matched filtering. The least-mean-square (LMS) branch extractor uses $n_{\text{tap}} = 11$, corresponding to $L_{\text{tap}} = 4n_{\text{tap}} + 2 = 46$ complex coefficients in each of the four FIR branches, and the first 4000 symbols (2000 pairs) form a known preamble. The FIR-weight least-mean-square adaptation uses a step size of $\mu_w = 5 \times 10^{-4}$ for each branch. This preamble drives data-aided phase de-rotation and common FIR training without using the simulated laser-phase trajectory. In the primary branch-training procedure, known preamble symbols form odd/even errors that update an auxiliary paired-error complex correction state using the distinct training-state adaptation step size $\mu_{p,\text{train}} = 0.05$. Once the FIR weights are held fixed after training, that auxiliary state is discarded. For the downstream tracker description, k

denotes the Alamouti-pair update index, p_1 and p_2 are the two complex correction-state accumulators, and $p = (p_1 + p_2^*)/2$ is the projected coefficient applied to the branch outputs; the update equations are given in Section 3.2. Both compared trackers use the same bootstrap initialization: the state accumulators start from $p_1^{(0)} = p_2^{(0)} = 0$, while $p^{(0)} = 1$ is used only to form the first odd/even outputs and their first errors. After that first state update, the applied coefficient is set by the projection $p^{(k+1)} = (p_1^{(k+1)} + p_2^{(k+1)*})/2$ and the recursion proceeds normally. Thus, $p = 1$ is not asserted to equal the projection of the zero accumulators at initialization. As a sensitivity check, all 30 $\theta = \pi/4$ benchmark records were rerun with projection-consistent initialization $p_1^{(0)} = p_2^{(0)} = p^{(0)} = 1$; every evaluated error count and BER was identical, with maximum absolute BER difference zero. Both trackers then traverse the same known preamble and operate decision-directed on identical fixed branch streams. The comparison therefore isolates the downstream correction rule: the single-error baseline is the correction rule of Faruk et al. [17] applied to these common four-branch outputs, rather than a full end-to-end reproduction of the original receiver. To isolate the downstream correction rule from FIR-training assistance, a strict tracker-neutral control disables the auxiliary branch-correction state entirely during the known preamble: the branch correction is held at unity, no auxiliary correction state is created or updated, and the four FIR banks are then held fixed before either downstream tracker is applied. Exact checks confirm that the two downstream trackers receive bit-identical FIR coefficients and branch streams. This neutral-training comparison is reported in Section 3.2 and repeated over the full general-Jones ensemble in Section 3.2.3.

The 16,384 transmitted pairs are numbered 0–16,383. The FIR tap window requires an 11-pair margin at each end, giving valid branch-pair indices 11–16,372 inclusive. The preamble occupies pairs 0–1999, and the post-training guard occupies pairs 2000–2199. At the tail, the implementation sets an exclusive BER cutoff by subtracting 200 from the last valid branch-pair index, giving $16{,}372 - 200 = 16{,}172$. The BER calculation accepts only indices below this cutoff, so the tail exclusion covers all 201 valid pairs from 16,172 through 16,372. The scored window is therefore pairs 2200–16,171 inclusive, containing 13,972 pairs, 27,944 symbols, and 111,776 bits per record.

To summarise unsuccessful realisations consistently, a run is counted as a tracker failure when its BER exceeds 3×10^{-2}. This criterion was predefined before the reported benchmark analysis as an algorithmic loss-of-robustness classification: it separates the low-BER tracking cluster from high-BER loss-of-tracking events in this simulation. It is not intended as a communications-system acceptability or forward-error-correction (FEC) threshold. Trajectory-level behaviour is examined separately ,and threshold sensitivity at the pre-FEC BER reference of 10^{-2} and at 2×10^{-2} is reported later in Section 3.2. Runs at or below the operational tracker-failure criterion of 3×10^{-2} are referred to as non-failed realisations.

2.3. Independent Realisations and Reproducibility

Independent stochastic realisations are generated from deterministic random-number seeds. The benchmark set uses base seeds 200–209, the received-optical-power sweep extends the same base ensemble to 200–229, and the disjoint tuning set uses 300–309; the 12-matrix Jones ensemble uses ensemble seed 6 September 2026. These numerical labels were chosen only to make the simulation ensembles deterministic and reproducible, with the tuning set kept separate from the benchmark data; they have no physical significance. Within each realisation, distinct deterministic sub-seeds are used for the transmitted data, Tx and LO phase noise, receiver noise, and Tx/Rx ENOB noise so that the stochastic components do not reuse the same random stream. Repeatability checks reproduce identical results for repeated runs of the same base seed. For the 200 GSa/s validation, the ex-

panded principal-SOP set uses seeds 200–219. The two independent 12-matrix Haar-SU(2) ensembles use Jones draw seeds 6 September 2026–17 September 2026 (ensemble A) and 3 November 2026–14 November 2026 (ensemble B), each crossed with stochastic seeds 200–209. Crossed seed–Jones intervals use 20,000 bootstrap replicates; stochastic seeds and Jones matrices are resampled independently with replacement and the reported 95% intervals are the 2.5th and 97.5th percentiles. The principal simulation parameters are summarised in Table 1.

Table 1. Simulation parameters.

Parameter	Value
Modulation/symbol rate	16-QAM/50 GBaud
Symbols/training symbols	32,768/4000
RRC roll-off/span	0.01/128 symbols
Branch filter n_{tap}/coeffs.	11/46
FIR-weight LMS step size μ_w	5×10^{-4}
Transmitter/receiver ENOB	6/6
Sampling grids	Tx/Tx phase/LO phase/detector noise: 100 GSa/s; LPF: 200 GSa/s; post-LPF/ENOB: 100 GSa/s
Optical amplification	None; ASE disabled
Principal SOP angles	0, $\pi/8$, $\pi/4$
General-SOP validation	12 fixed Haar SU(2) Jones matrices
Per-laser linewidths ($\Delta\nu_{\mathrm{TX}} = \Delta\nu_{\mathrm{LO}}$)	100, 500, 1000, 2000 kHz
Benchmark/received-power/tuning seeds	200–209/200–229/300–309
Operational tracker-failure criterion	3×10^{-2}

Table 1 describes the 100 GSa/s reference implementation. The 200 GSa/s control changes only the front-end sampling realisation described above; the symbol rate, modulation, SOP/Jones definitions, FIR length, training length, and BER window are retained.

3. Results

For comparison and reproducibility, Figures 2–4, Tables 2 and 3, and the associated paragraphs report the 100 GSa/s reference implementation. Section 3.3 provides the 200 GSa/s validation and is the principal basis for the robustness conclusions in Sections 4 and 5; the 100 GSa/s results are retained as reference-model diagnostics.

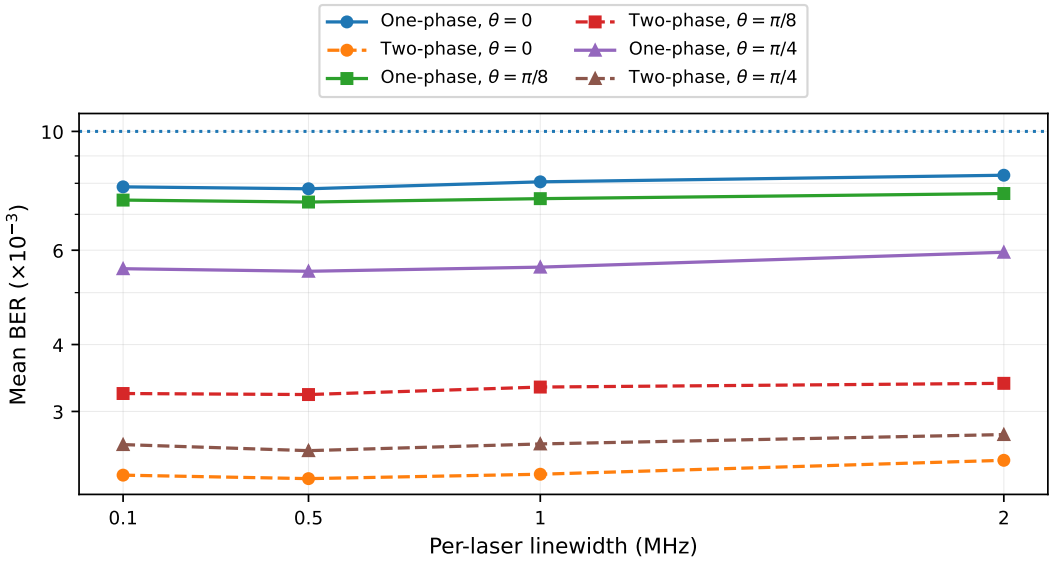

Figure 2. Mean BER for the one-phase and two-phase pair-level diagnostics over ten independent realisations and three SOP angles. The dotted line denotes the pre-FEC BER reference of 10^{-2}. This is a 100 GSa/s reference-model diagnostic; 200 GSa/s tracker evidence is reported in Section 3.3.

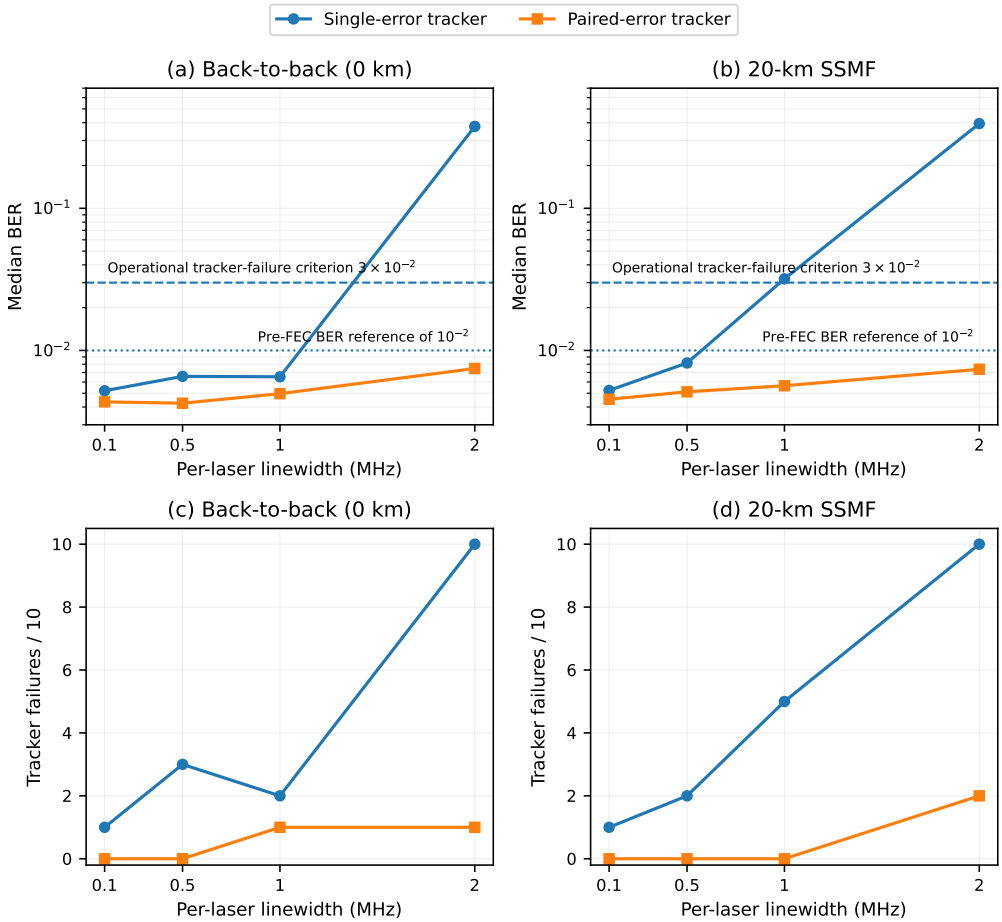

Figure 3. Performance versus per-laser linewidth at $\theta = \pi/4$ for the single-error and projected paired-error trackers: (**a**) median BER, back-to-back; (**b**) median BER, 20 km; (**c**) tracker failures out of ten, back-to-back; and (**d**) tracker failures out of ten, 20 km. This is a 100 GSa/s reference-model result; 200 GSa/s fibre controls are reported in Section 3.3.3.

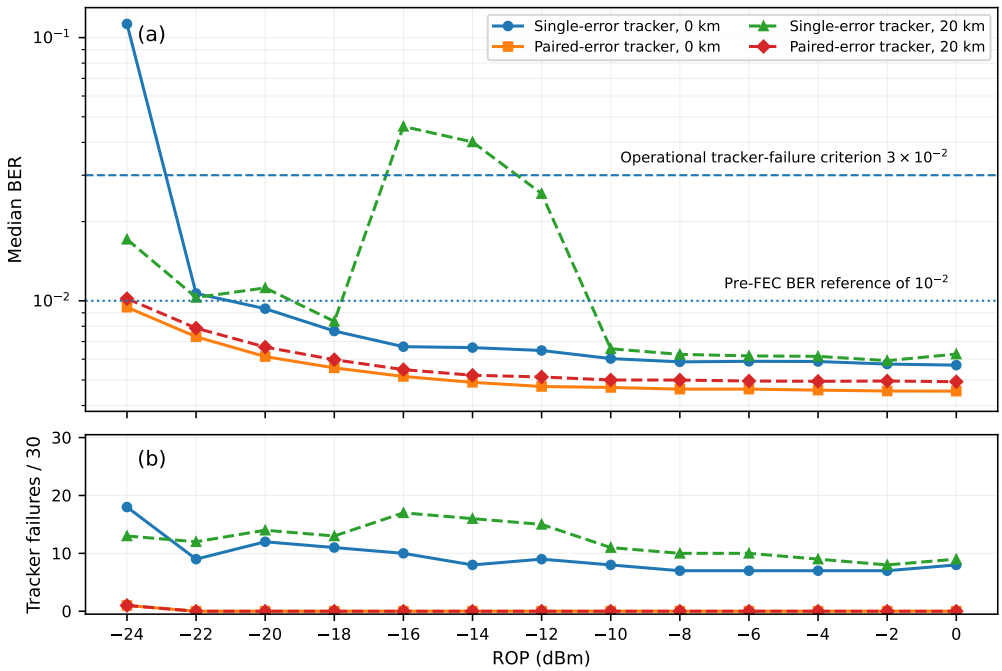

Figure 4. Median BER and threshold-failure counts versus received optical power for the single-error and projected paired-error trackers in the 100 GSa/s reference model at $\Delta\nu_{\mathrm{Tx}} = 100$ kHz, $\Delta\nu_{\mathrm{LO}} = 1$ MHz, and $\theta = \pi/4$: (**a**) median BER; (**b**) failures out of 30 independent realisations.

Table 2. Principal SOP sweep for both trackers. Each cell gives median BER with failures out of ten in parentheses using the predefined operational tracker-failure criterion of 3×10^{-2}. This is a 100 GSa/s reference-model result.

Per-Laser Linewidth	$\theta = 0$		$\theta = \pi/8$		$\theta = \pi/4$	
	Single	**Paired**	**Single**	**Paired**	**Single**	**Paired**
100 kHz	1.25×10^{-1} (5/10)	2.01×10^{-3} (1/10)	1.91×10^{-2} (5/10)	3.35×10^{-3} (0/10)	5.21×10^{-3} (1/10)	4.35×10^{-3} (0/10)
500 kHz	3.87×10^{-2} (5/10)	2.35×10^{-3} (1/10)	5.37×10^{-3} (3/10)	3.70×10^{-3} (0/10)	6.57×10^{-3} (3/10)	4.25×10^{-3} (0/10)
1 MHz	1.73×10^{-1} (9/10)	3.19×10^{-3} (1/10)	1.61×10^{-1} (7/10)	3.98×10^{-3} (1/10)	6.53×10^{-3} (2/10)	4.97×10^{-3} (1/10)
2 MHz	3.74×10^{-1} (9/10)	9.95×10^{-2} (7/10)	4.28×10^{-1} (10/10)	4.23×10^{-2} (5/10)	3.77×10^{-1} (10/10)	7.47×10^{-3} (1/10)

Table 3. Summary of common fixed-branch correction comparisons. All rows use per-laser linewidths of 500 kHz, 1 MHz, and 2 MHz, the ten benchmark realisations, and the operational tracker-failure criterion BER $> 3 \times 10^{-2}$. S and P denote the single-error and projected paired-error rules.

SOP Ensemble	FIR Training/Step Choice	μ_p^{S}	μ_p^{P}	S Failures	P Failures	P BER Lower
$\theta = \pi/4$	Assisted/common	0.02	0.02	15/30	2/30	—
$\theta = \pi/4$	Assisted/tuned	0.01	0.02	10/30	2/30	26/30
$\theta = \pi/4$	Neutral/tuned	0.02	0.02	24/30	21/30	23/30
12 Jones matrices	Assisted/common	0.02	0.02	234/360	62/360	331/360
12 Jones matrices	Neutral/reused	0.02	0.02	287/360	254/360	269/360

Tuned steps are selected independently using only the disjoint tuning realisations at $\theta = \pi/4$; the selected neutral-training steps are reused unchanged in the Jones ensemble. "P BER lower" counts matched records with lower paired-error BER. The dash denotes an aggregate matched-BER count not reported for that configuration. Per-linewidth assisted-training medians and failure counts remain in Table 2.

3.1. Phase-Combiner Sensitivity

This section examines the effect of constraining the two symbols in each Alamouti pair to share one complex correction. Known-symbol one-phase and two-phase searches are used to isolate pair-combiner sensitivity before the causal tracker implementations are evaluated. For pair k, let a_o, a_e and b_o, b_e denote the odd/even outputs of the two fixed Alamouti branch filters. The quantities v_o, v_e are the combined odd/even soft symbols, e_o, e_e are reference or decision errors, and p_1, p_2 are the implemented complex branch-correction states. Here and below, $j = \sqrt{-1}$. A one-phase combiner applies the complex correction $p = e^{-j\hat{\phi}}$ and its conjugate, where $\hat{\phi}$ is the common phase estimate for the Alamouti pair,

$$v_o = a_o p + b_o p^*, \qquad v_e = a_e p + b_e p^*. \tag{7}$$

For an illustrative noiseless derivation, let h_x, h_y be unit-norm polarisation-channel coefficients, $|h_x|^2 + |h_y|^2 = 1$. The received pair is $r_o = (h_x s_{2k} + h_y s_{2k+1}) e^{j\phi_{2k}}$ and $r_e = (-h_x s_{2k+1}^* + h_y s_{2k}^*) e^{j\phi_{2k+1}}$, where the phases denote values at the two symbol slots. Choosing the common correction $p = e^{-j\phi_{2k}}$ and forming $v_o = \bar{h}_x r_o p + h_y r_e^* p^*$ gives, for $\delta = \phi_{2k} - \phi_{2k+1}$,

$$v_o = s_{2k}\left(|h_x|^2 + |h_y|^2 e^{j\delta}\right) + \bar{h}_x h_y s_{2k+1}\left(1 - e^{j\delta}\right). \tag{8}$$

Here, $\bar{h}_x$ denotes complex conjugation. For small δ, the cross-symbol term is approximately $-j\delta\bar{h}_x h_y s_{2k+1}$. If the channel coefficients include an additional common gain, the gain-normalised combiner output requires division by $|h_x|^2 + |h_y|^2$. The known phase used in this illustrative derivation is not supplied to the causal trackers. Equation (8) isolates an

intra-pair drift contribution; the numerical search also captures decision geometry, receiver noise, quantisation, fixed branch estimates, and finite-grid effects.

For each pair the known-symbol diagnostic minimises the known-symbol squared-error objective

$$\mathcal{J} = |v_o - s_{2k}|^2 + |v_e - s_{2k+1}|^2. \tag{9}$$

The one-phase search evaluates $p_m = \exp(j2\pi m/N_1)$, $m = 0, \ldots, N_1 - 1$, with $N_1 = 256$ and substitutes the same p_m and p_m^* into both outputs of Equation (7). The two-phase search uses two independent uniform full-range grids, $p_r = \exp(j2\pi r/N_2)$ and $p_s = \exp(j2\pi s/N_2)$, with $r, s = 0, \ldots, N_2 - 1$ and $N_2 = 128$, and evaluates $v_o = a_o p_r + b_o p_s^*$ and $v_e = a_e p_r + b_e p_s^*$. The minimum-cost hypothesis is selected independently for each pair; in an exact numerical tie, the first minimum in increasing grid order is retained. Both searches therefore cover $[0, 2\pi)$ uniformly, use known transmitted symbols, fixed branch estimates, and no temporal continuity. The two-phase search has an additional degree of freedom, so the result is interpreted as a degrees-of-freedom/combiner sensitivity diagnostic rather than a laser-phase-noise penalty. Accordingly, the comparison is used only as a constrained-versus-less-constrained numerical diagnostic with known symbols and pair-local optimization, not as a comparison between practical competing receivers. The resulting linewidth dependence is shown in Figure 2.

Figure 2 shows that the one-phase search produces a consistently higher mean BER than the two-phase search for every evaluated SOP angle and linewidth. Across 100 kHz–2 MHz, the mean BER within each search family changes only modestly compared with the persistent separation between the one-phase and two-phase results. The one-phase-to-two-phase mean-BER ratio ranges from 2.13 to 3.52 across the 12 principal linewidth–SOP conditions. In a separate zero-linewidth check, the ratios are 3.46, 2.24, and 2.17 for $\theta = 0$, $\pi/8$, and $\pi/4$, respectively. The persistence of this separation without stochastic laser-phase drift indicates that the numerical gap also includes decision geometry, receiver noise, quantisation, fixed branch estimates, finite-grid resolution, and the additional degree of freedom provided by the two-phase search.

Increasing the one-phase grid from 256 to 512 changes aggregate BER by at most 2.02% and its diagnostic ratio by 2.02%. For the two-phase search, a grid-refinement check increased the two-dimensional phase-hypothesis grid from 128×128 to 192×192, i.e., from 128 to 192 candidate phases for each of the two independently searched corrections. The check used three benchmark realisations for three condition groups: 100 kHz at $\theta = 0$, 1 MHz at $\theta = \pi/4$, and 2 MHz at $\theta = \pi/4$; their condition-mean BER and diagnostic-ratio changes were at most 3.15% and 3.05%, respectively. A subsequent 2 MHz stress test selected, before refined-grid calculation, the realisation with the largest archived 128×128 two-phase BER at each SOP. Two of three selected cases remained below the predefined 5% criterion; the $\theta = \pi/4$, selected stress realisation changed by 6.27% in BER and 5.90% in ratio. Refining that same case from 192×192 to 256×256 changed BER and ratio by only 0.51%. Moreover, for the full ten-realisation ensemble at 2 MHz and $\theta = \pi/4$ plotted in Figure 2, the mean two-phase BER changes by 2.08% and the mean-BER diagnostic ratio by 2.04% from 128×128 to 192×192. Thus, the plotted ensemble result is numerically stable under refinement, although an individual selected stress realisation can exceed the coarse-to-intermediate 5% criterion. The 5% value is only a numerical convergence criterion for the finite-grid search, not a communications-system tolerance or accuracy requirement.

3.2. Paired-Error Branch Tracker

This section compares the implemented single-error and paired-error correction updates on identically trained branch streams held fixed during comparison. Hereafter, "single-error tracker" denotes the baseline rule that updates both branch-correction states

from the odd-symbol error only, whereas "projected paired-error tracker" denotes the proposed rule that averages the odd- and even-symbol error contributions before projection to the common conjugate-consistent coefficient. For these causal trackers, p_1 and p_2 denote the complex correction-state accumulators and $p = (p_1 + p_2^*)/2$ denotes the applied projected coefficient; unlike the one-phase diagnostic of Equation (7), this tracker coefficient is not constrained to unit magnitude. The update rules are introduced first, followed by the independent-realisation benchmark, the 20 km linear-SSMF extension, and the tracker comparison.

3.2.1. The SOP-Dependent Tracker-Failure Problem

On the common fixed branch streams, the single-error correction rule of Faruk et al. [17], implemented here after the common fixed branch extractor, updates both branch states from the odd-symbol error $e_o = d_o - v_o$, where d_o is the known preamble symbol or the current hard decision; d_e is defined similarly for the even output. With μ_p denoting the branch-correction adaptation step size, the rule is

$$\begin{aligned} p_1 &\leftarrow p_1 + \mu_p e_o a_o^*, \\ p_2 &\leftarrow p_2 + \mu_p e_o b_o^*, \\ p &\leftarrow (p_1 + p_2^*)/2. \end{aligned} \tag{10}$$

For the paired-error rule, define the instantaneous constrained cost

$$\mathcal{J}_k(p) = \frac{1}{2} \sum_{t \in \{o,e\}} |d_t - a_t p - b_t p^*|^2, \tag{11}$$

where d_o, d_e are held fixed during the update. With $e_t = d_t - a_t p - b_t p^*$, Wirtinger differentiation gives

$$\frac{\partial \mathcal{J}_k}{\partial p^*} = -\frac{1}{2} \sum_{t \in \{o,e\}} (e_t a_t^* + e_t^* b_t). \tag{12}$$

The branch correction requires a conjugate-consistent pair (p, p^*). For provisional complex states (p_1, p_2), its Euclidean projection is

$$p = \arg\min_{c \in \mathbb{C}} \{|c - p_1|^2 + |c^* - p_2|^2\} = \frac{p_1 + p_2^*}{2}. \tag{13}$$

The implemented averaged paired-error recursion is

$$\begin{aligned} p_1 &\leftarrow p_1 + \frac{\mu_p}{2}(e_o a_o^* + e_e a_e^*), \\ p_2 &\leftarrow p_2 + \frac{\mu_p}{2}(e_o b_o^* + e_e b_e^*), \\ p &\leftarrow (p_1 + p_2^*)/2. \end{aligned} \tag{14}$$

The factor $1/2$ averages the two error contributions. After the exceptional first bootstrap update defined in Section 2, $p = (p_1 + p_2^*)/2$ holds at each update. Combining the state increments and projection gives

$$\begin{aligned} p^{(k+1)} &= p^{(k)} + \frac{\mu_p}{4} \sum_{t \in \{o,e\}} (e_t a_t^* + e_t^* b_t) \\ &= p^{(k)} - \frac{\mu_p}{2} \frac{\partial \mathcal{J}_k}{\partial p^*}. \end{aligned} \tag{15}$$

Thus, the applied coefficient follows an exact instantaneous constrained LMS update with effective Wirtinger step $\mu_p/2$. This local relation does not establish global convergence or decision-directed tracking stability.

The relation to Oh et al. [10] can be made explicit. After conjugating their even recovered output to match the source-symbol convention here, identify $a_o = y_o w_1$, $b_o = y_e^* w_2$, $a_e = y_o w_2^*$, and $b_e = -y_e^* w_1^*$, where y_o, y_e and w_1, w_2 are their received inputs and shared FIR coefficients, using their filter-product notation. The even reference and error are conjugated correspondingly, $e_e = (e_{e,\mathrm{Oh}})^*$. The applied-coefficient increment from their Equation (10) then equals that of Equation (14) with their step $\mu_{\mathrm{Oh}} = \mu_p/2$. Their shared-filter relations need not hold for the four separate FIR coefficient banks used here.

During the preamble, both errors use known symbols; hard 16-QAM decisions are used thereafter. For the assisted-training downstream comparison, both trackers use $\mu_p = 0.02$; independently tuned values are reported later in this section. The coefficients p_1, p_2, p are unconstrained complex values and can correct residual gain as well as phase. Relative to Equation (10), Equation (14) adds two complex multiplications and two complex additions per pair. The even-symbol error is already available, no additional correction-state memory is required, and the added products can be evaluated in parallel.

The observed trajectories are consistent with loss of tracking associated with updating both branch states from only the odd-symbol error in some SOP/linewidth conditions. Fixed branch weights isolate the downstream correction stage conditional on the chosen FIR-training procedure; the neutral-training control separately quantifies this training dependence. Joint FIR-weight and branch-correction adaptation is reserved for future work.

3.2.2. Independent-Realisation Benchmark

At the balanced-mixing SOP condition $\theta = \pi/4$, the paired-error median BERs at 100 kHz, 500 kHz, 1 MHz, and 2 MHz are 4.35×10^{-3}, 4.25×10^{-3}, 4.97×10^{-3}, and 7.47×10^{-3}, respectively. No paired-error realisation exceeds the predefined operational tracker-failure criterion of 3×10^{-2} at 100 or 500 kHz, whereas one of ten realisations exceeds it at each of 1 and 2 MHz, as reported in Table 2. The known-symbol one-phase and two-phase diagnostics are already presented in Figure 2; they are not repeated here. Table 2 extends the causal-tracker comparison across all three SOP angles.

Across all 12 SOP–linewidth conditions in Table 2, the paired-error tracker gives both a lower median BER and fewer threshold failures than the single-error tracker. At 2 MHz, the paired-error failure counts are 7/10, 5/10, and 1/10 for $\theta = 0$, $\pi/8$, and $\pi/4$, respectively, showing that the balanced-mixing condition is not the universal worst case.

The causal tracker exploits temporal continuity. The pair-level diagnostic optimises each pair independently using known symbols and a finite grid, so the two analyses provide complementary information. The non-monotonic SOP dependence at the largest linewidth is therefore reported empirically; the present simulations do not isolate a single physical cause.

3.2.3. General-Jones SOP Validation

To extend the SOP evaluation beyond the three real-rotation angles of Table 2, a separate validation was performed using 12 fixed Haar-distributed lossless Jones matrices from Equation (5), the three per-laser linewidths 500 kHz, 1 MHz, and 2 MHz, and ten independent realisations. This gives 360 matched records for each training procedure. With the primary paired-error-assisted FIR training, threshold-defined failures fall from 234/360 to 62/360 and the paired-error BER is lower in 331/360 records. Of all matched records, 178 are rescued and six worsen. Aggregating the 36 repeated conditions within each independent base realisation, all ten realisations improve and none worsen (exact

two-sided sign test $p = 0.002$). The crossed realisation–Jones bootstrap gives an absolute failure-fraction reduction of 47.8 percentage points with a 95% confidence interval of 37.5–58.1 percentage points. The 360 crossed records are not interpreted as 360 independent samples. In this 100 GSa/s reference-model bootstrap, a base stochastic realisation and a Jones matrix are the two resampling units; both dimensions are resampled independently and their crossed cells are retained within each bootstrap replicate.

To test the FIR-training dependence directly, the complete 360-record Jones study was then repeated with the strict tracker-neutral training defined in Section 2, using the same stored Jones matrices, realisations, linewidths, and independently selected downstream steps $\mu_p = 0.02$ for both trackers. Failures are 287/360 for the single-error tracker and 254/360 for the projected paired-error tracker, a reduction of 33 rather than 172 records. Paired-error BER is lower in 269/360 matched records, 33 single-error failures are rescued, and no non-failed single-error record becomes a paired-error failure. On the 73 records non-failed for both trackers, the matched-subset medians are 1.21×10^{-2} and 1.01×10^{-2}, respectively. Seven of ten independent base realisations improve in failure count, three tie, and none worsen (exact two-sided sign test $p = 0.016$). The absolute failure-fraction reduction is 9.2 percentage points; the realisation-cluster and crossed realisation–Jones bootstrap 95% confidence intervals are 4.7–13.6 and 3.6–15.8 percentage points, respectively. At 2 MHz, both neutral-trained trackers fail all 120 records. The neutral result therefore retains a smaller, statistically consistent downstream advantage, while showing that most of the larger assisted-training gain arises from the interaction between FIR training and downstream paired-error correction.

3.2.4. 20 km Linear-SSMF Transmission

The fibre study compares back-to-back and 20 km transmission across four per-laser linewidths, three state-of-polarisation (SOP) rotation angles, and ten independent realisations at each of two distances, giving 240 simulation records ($4 \times 3 \times 10 \times 2$) or 120 matched back-to-back/20 km pairs. The 20 km standard single-mode fibre (SSMF) model includes 4 dB attenuation and second- and third-order chromatic dispersion. Frequency-domain chromatic-dispersion compensation (CDC) is applied before matched filtering and frame synchronisation; polarisation-mode dispersion, Kerr nonlinearity, optical amplification, and amplifier noise are excluded.

Compensation was verified by a noise-free propagation/CDC round trip, with normalised mean-square error of order 10^{-30} after removing attenuation. For transmitter and LO linewidths of 500 kHz each, an SOP rotation angle of $\theta = \pi/4$, and a representative benchmark realisation, omitting CDC gives BER 5.01×10^{-1} for both trackers, whereas applying CDC gives paired-error BER 4.71×10^{-3}.

Figure 3 compares the two distances at $\theta = \pi/4$ in the 100 GSa/s reference model. At 20 km, across the tested per-laser linewidths of 500 kHz, 1 MHz, and 2 MHz, failures are 17/30 and 2/30 for the single-error and projected paired-error trackers, respectively. Paired-error BER is lower in 39/40 matched records across all four linewidths. Attenuation-only, compensation-disabled, and power-restored controls show that the nonmonotonic means over all realisations primarily reflect changes in the high-BER failure set.

To assess equalisation-enhanced phase noise (EEPN) [38], transmitter-only and LO-only phase noise were evaluated separately at 1 and 2 MHz, using the ten benchmark realisations and $\theta = \pi/4$. The controls covered back-to-back transmission, 20 km with CDC and 4 dB loss, and 20 km with received power restored. The projected paired-error tracker had no LO-only threshold failures. For LO-only phase noise, median BER increased from 4.55×10^{-3} to 5.28×10^{-3} at 1 MHz and from 5.47×10^{-3} to 5.70×10^{-3} at 2 MHz between the back-to-back reception and 20 km with CDC and 4 dB loss. The only additional paired-

error failure occurred in the 2 MHz transmitter-only case and disappeared on restoring received power.

For consistency with the standard analytical estimate, the EEPN phase-noise variance is approximated by

$$\sigma^2_{\mathrm{EEPN}} \approx \frac{\pi\lambda^2 DL\Delta\nu_{\mathrm{LO}}}{2cT_s}, \tag{16}$$

where σ^2_{EEPN} is the EEPN phase-noise variance, λ is the optical carrier wavelength, D is the chromatic-dispersion coefficient, L is the fibre length, $\Delta\nu_{\mathrm{LO}}$ is the LO linewidth, c is the speed of light in vacuum, and $T_s = 1/R_s$ is the symbol period. Equation (16) gives root-mean-square (RMS) phase values of approximately 0.836° and 1.182° for 1- and 2-MHz LO linewidths, respectively, with the present 50 GBaud, 20 km parameters. The small LO-associated BER increase is compatible with an EEPN contribution, but the reported back-to-back/20 km medians also include the 4 dB received-power reduction and therefore do not isolate an EEPN-induced BER penalty. No LO-only threshold failures were observed in the evaluated paired-error controls. The EEPN interpretation therefore draws on the native 200 GSa/s Tx-only and LO-only controls with received power restored (Section 3.3.3).

3.2.5. Received-Optical-Power Robustness

To complement the fixed-operating-point tracker comparison, an additional received optical power (ROP) sweep was performed under an asymmetric laser-linewidth condition representative of a narrower transmitter source and a broader receiver local oscillator. Here, ROP denotes the total signal power across both polarisations immediately before the receiver/LO coupler. Figure 4 shows the simulated tracker response to received power under ideal record-based digital normalisation. Hardware receiver sensitivity is outside the scope of this diagnostic because the model omits a practical AGC loop, fixed ADC full-scale selection, clipping, and saturation.

The transmitter and LO linewidths were set to 100 kHz and 1 MHz, respectively, with $\theta = \pi/4$. For both 0 and 20 km, ROP was swept from -24 dBm to 0 dBm in 2 dB steps using 30 independent realisations. Before applying the ROP-sweep power scaling, the received powers are 0 dBm for the back-to-back configuration and -4 dBm for the 20 km fibre link. At each ROP point, the signal field is multiplied by the square root of the required power ratio. LO power, thermal-noise model, filters, ENOB, FIR/tracker settings, and realisation assignments remain fixed; shot-noise variance is recomputed from the power incident on the detected coupler arms. The same 30 independent realisations and their underlying data/phase/noise Gaussian sequences are reused at every ROP to form paired power-point comparisons. Receiver/ENOB noise is inserted before the subsequent digital normalization. After downconversion, optional CDC, matched filtering, and synchronization, a single scalar gain $g = \sqrt{\mathbb{E}|s_{\mathrm{ref}}|^2/\mathbb{E}|r_{\mathrm{sync}}|^2}$ is calculated, where s_{ref} denotes the known reference coded-signal samples, r_{sync} denotes the noisy synchronized receiver samples, and $\mathbb{E}$ denotes the record average. The gain g is applied to that same noisy waveform before common FIR training, scaling signal and noise together. Automatic-gain-control (AGC) dynamics, analog-to-digital-converter (ADC) clipping, and full-scale adaptation are not modelled. Figure 4 therefore reports the simulated BER dependence on ROP under ideal digital gain normalization; each point gives the median over 30 independent realisations. The dotted horizontal line denotes the pre-FEC BER reference of 10^{-2} and is distinct from the predefined operational tracker-failure criterion of 3×10^{-2}.

For the projected paired-error tracker, median BER remains below the pre-FEC BER reference of 10^{-2} at all sampled ROP values in back-to-back operation and from -22 dBm upwards for the 20 km link. At -24 dBm, 20/30 and 12/30 realisations achieve BER at or below the pre-FEC BER reference of 10^{-2} for the 0 and 20 km configurations, respectively.

At −22 dBm, the counts rise to 29/30 in both configurations, with no paired-error records exceeding the predefined operational tracker-failure criterion of 3×10^{-2}. The 20 km single-error curve exhibits a pronounced nonmonotonic excursion between −18 and −10 dBm, peaking near −16 dBm. This behaviour does not represent a conventional power-dependent BER penalty; rather, it results from changes in which realisations enter the high-BER loss-of-tracking state at each ROP, as also reflected by the corresponding failure counts in Figure 4b. Figure 4b reports failure counts across ROP: 7–18/30 at 0 km and 8–17/30 at 20 km for the single-error tracker, compared with 1/30 at −24 dBm and 0/30 at higher sampled powers for the paired-error tracker at either distance.

3.2.6. Common-Branch Comparison and Tracker-Failure Diagnostic

The assisted-training common-branch comparison evaluates the downstream correction rules at the same step size on identical fixed branch streams. Table 3 summarizes the assisted, independently tuned, and tracker-neutral comparisons. As reported in Table 2, the paired-error tracker reduces failures from 3/10 to 0/10 at 500 kHz, from 2/10 to 1/10 at 1 MHz, and from 10/10 to 1/10 at 2 MHz. Because the failure mixture is strongly non-Gaussian, the primary descriptive quantities are the failure counts and all-record medians; medians conditioned on separately defined non-failed subsets are not compared as if they were matched samples.

For an independently step-size-tuned tracker comparison under the primary assisted FIR training, the two downstream trackers are tuned independently on exactly the same disjoint tuning ensemble at the balanced-mixing SOP condition $\theta = \pi/4$ for the tested per-laser linewidths of 500 kHz, 1 MHz, and 2 MHz. Candidate steps are 0.005, 0.01, 0.02, and 0.05; selection minimizes failures, with median and mean BER as tie breakers. Independent step-size tuning selects a downstream branch-correction adaptation step size of $\mu_p = 0.01$ for the single-error tracker and $\mu_p = 0.02$ for the projected paired-error tracker. Applying these values unchanged to the held-out benchmark realisations gives 10/30 single-error versus 2/30 paired-error failures; nine failed single-error records are rescued, one previously non-failed record worsens, and paired-error BER is lower in 26/30 matched records. The original common-$\mu_p = 0.02$ result is used as the assisted-training downstream comparison, in which both trackers use the same step size; it is interpreted separately from the independently tuned and tracker-neutral comparisons.

The stronger isolation control uses the strict tracker-neutral FIR training defined in Section 2 and independently selects the downstream step for each tracker using only the disjoint tuning realisations. The existing failure/median/mean selection objective chooses $\mu_p = 0.02$ for both trackers; these step sizes are then kept unchanged for held-out benchmark realisations. On the resulting 30 matched records, failures are 24/30 for the single-error tracker and 21/30 for the projected paired-error tracker. The all-record median BERs are 0.350 and 0.254, respectively, and the paired-error BER is lower in 23/30 matched records. Only six records are non-failed for both trackers; on this identical common subset the medians are 1.290×10^{-2} and 1.158×10^{-2}. At 2 MHz both trackers fail all ten held-out records. Exact identity and immutability checks give zero numerical difference between the fixed FIR coefficients and branch streams supplied to the two trackers. The common non-failed subset and marginal failure counts imply three rescued records and no record-level worsening. This three-record failure reduction is descriptive at the balanced SOP: at most three independent realisation clusters can improve, so an exact two-sided realisation-level sign test cannot give $p < 0.25$. The statistically supported neutral-training advantage over the finite general-Jones ensemble in Section 3.2.3 is a separate result.

Because linewidths reuse the same ten benchmark realisations, statistical inference is performed at realisation level. A realisation is classified as improved when the paired-error

tracker has fewer threshold failures than the single-error tracker across that realisation's three 500 kHz, 1 MHz, and 2 MHz records; equal failure counts are ties, and BER magnitude does not break a tie. Under the assisted comparison with $\mu_p = 0.02$ for both trackers using the operational tracker-failure criterion of 3×10^{-2}, failures are 15/30 versus 2/30: 13 records are rescued and none worsen; nine realisations improve, one ties, and none worsen (exact sign test $p = 0.004$; realisation-cluster bootstrap reduction 43.3 percentage points, 95% confidence interval 26.7–60.0 percentage points). As a sensitivity check, using the pre-FEC BER reference of 10^{-2} as the classification threshold gives 17/30 versus 3/30 failures ($p = 0.008$), while using 2×10^{-2} gives 16/30 versus 2/30 ($p = 0.004$), with no realisation-level worsening.

Figure 5 presents a selected illustrative tracker-failure case at a per-laser linewidth $\Delta\nu = 2$ MHz and $\theta = \pi/4$ for a selected benchmark realisation. The upper panel shows the unwrapped phase of the applied complex branch-correction coefficient $p = (p_1 + p_2^*)/2$, where p_1 and p_2 are the adaptive complex correction states defined in Section 3.2. The two tracker trajectories initially evolve similarly but separate during the displayed interval, after which the single-error tracker follows a substantially different phase trajectory. The lower panel shows the 200-pair running mean of the total tracker error power, $|e_o|^2 + |e_e|^2$. The single-error tracker exhibits pronounced error-power excursions, whereas the paired-error trajectory remains lower over most of the displayed interval. The plotted trajectories show the first 7000 pairs, while BER is evaluated over pairs 2200–16,171. Over this 111,776-bit evaluation window, the single-error and paired-error trackers produce 70,587 and 756 bit errors, corresponding to BERs of 6.315×10^{-1} and 6.764×10^{-3}, respectively.

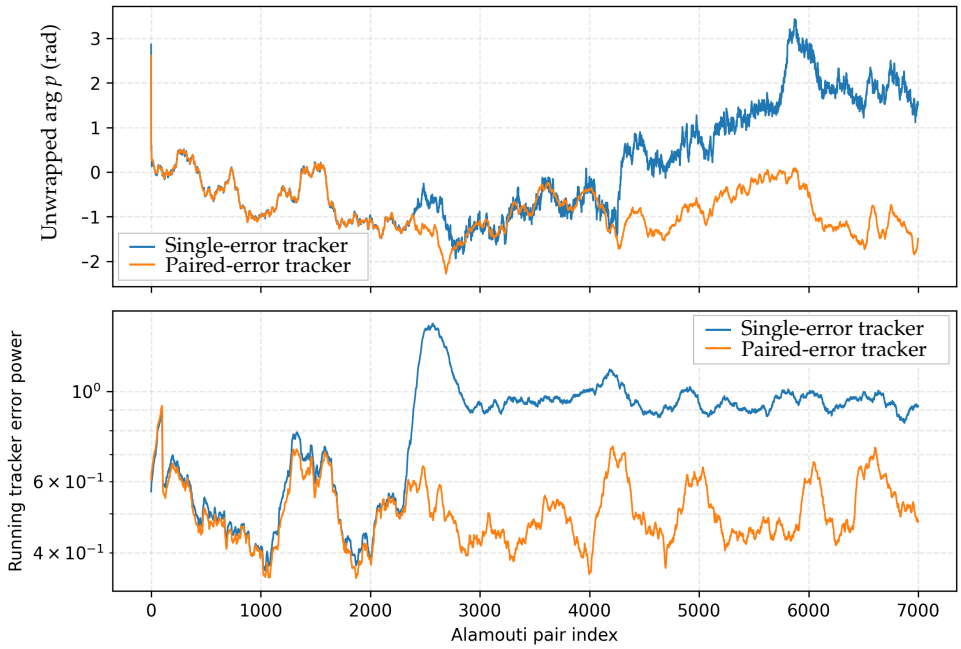

Figure 5. Illustrative 2 MHz, $\theta = \pi/4$, selected benchmark realisation fixed-branch tracker diagnostic. Upper panel: unwrapped phase of the applied coefficient $p = (p_1 + p_2^*)/2$. Lower panel: 200-pair running mean of $|e_o|^2 + |e_e|^2$. 100 GSa/s reference-model example; multi-run 200 GSa/s trajectories are shown in Figure 6.

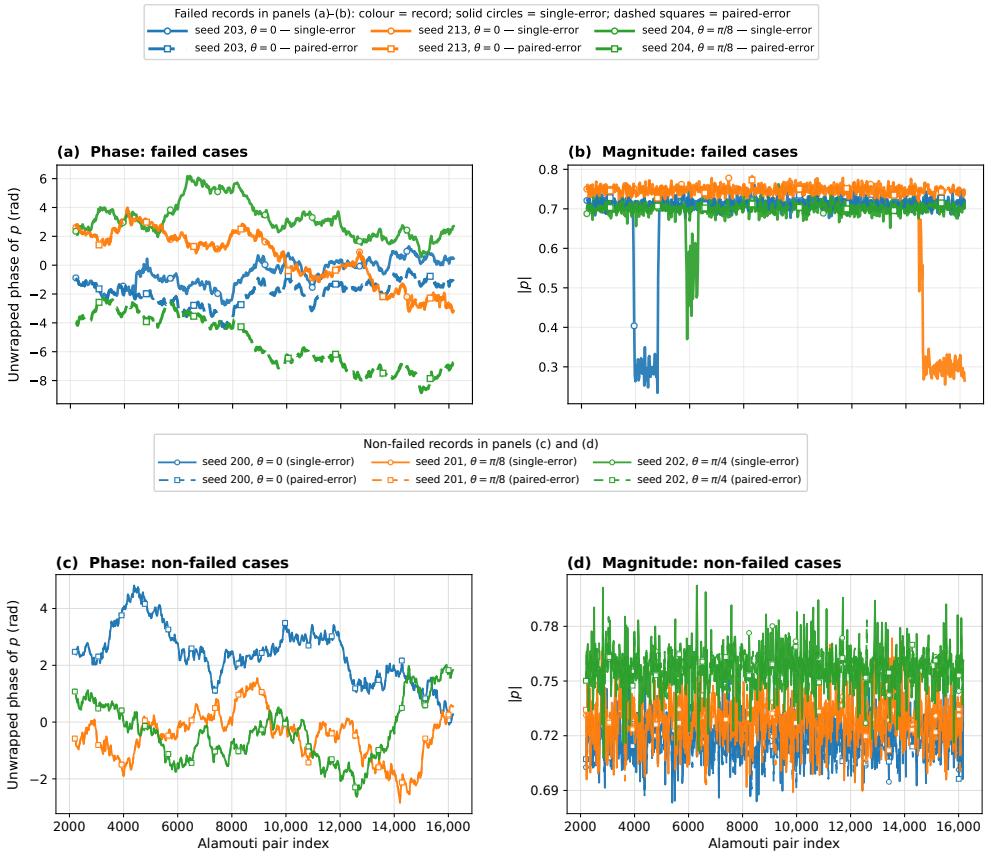

Figure 6. The 200 GSa/s 2 MHz coefficient trajectories for the six representative trajectory-diagnostic records: three observed single-error failures and three non-failed records, one per SOP angle. Panels (**a**,**b**) show single-error failures and (**c**,**d**) show non-failed records. Panels (**a**,**c**) plot unwrapped phase, and (**b**,**d**) plot $|p|$. Colour identifies the record. Solid lines with circles show the single-error tracker, and dashed lines with squares show the paired-error tracker. Phase offsets from unwrapping can differ across records, so phase comparisons are made within each record.

3.3. 200 GSa/s Validation and Robustness Controls

The native 200 GSa/s receiver described in Section 2.2 places the heterodyne edge inside the initial Nyquist band. The results in this section therefore provide the principal robustness evidence; the 100 GSa/s results above are retained as reference-model diagnostics. The controls below separate static-SOP robustness, FIR-training dependence, phase recovery, and fibre effects.

3.3.1. Static-SOP Robustness

With the common downstream step $\mu_p = 0.02$, the 200 GSa/s 10-seed sweep over three SOP angles and four per-laser linewidths contains two single-error failures and no paired-error failures in 120 records. At $\theta = \pi/4$, neither tracker fails in the 30 records spanning 500 kHz–2 MHz, compared with 15/30 single-error and 2/30 paired-error failures in the 100 GSa/s reference model. The 2 MHz comparison was then extended to 20 seeds; Table 4 reports BER distributions and Wilson 95% intervals. The paired rule has a lower median BER at all three static SOP angles, while the failure-count separation is much smaller than in the 100 GSa/s reference model.

Table 4. A 200 GSa/s 2 MHz static-SOP comparison with $\mu_p = 0.02$ and 20 independent stochastic realisations per angle. BER entries are median [IQR]; intervals are Wilson 95% confidence intervals for the failure probability.

SOP	Single BER	Paired BER	S Failures	P Failures	Failure-Probability 95% CI (S/P)
$\theta = 0$	2.68×10^{-5} [8.95×10^{-6}, 4.92×10^{-5}]	8.95×10^{-6} [6.71×10^{-6}, 2.68×10^{-5}]	2/20	0/20	0.028–0.301/0–0.161
$\theta = \pi/8$	1.52×10^{-4} [1.07×10^{-4}, 2.42×10^{-4}]	4.47×10^{-5} [3.58×10^{-5}, 7.16×10^{-5}]	1/20	1/20	0.0089–0.236/0.0089–0.236
$\theta = \pi/4$	2.68×10^{-5} [8.95×10^{-6}, 4.92×10^{-5}]	1.34×10^{-5} [6.71×10^{-6}, 2.01×10^{-5}]	0/20	0/20	0–0.161/0–0.161

For each fixed SOP angle, we resampled the 20 stochastic seeds with replacement, keeping the single-error and paired-error BERs together for each seed. The median matched BER differences (single minus paired) and pointwise 95% percentile intervals from 20,000 replicates were 8.95×10^{-6} [0, 2.68×10^{-5}] at $\theta = 0$, 8.50×10^{-5} [4.03×10^{-5}, 1.61×10^{-4}] at $\theta = \pi/8$, and 8.95×10^{-6} [4.47×10^{-6}, 2.24×10^{-5}] at $\theta = \pi/4$. Paired-error BER was lower in 13/20, 18/20, and 14/20 realisations, respectively. Table A1 reports the individual bit-error counts from which these BERs are calculated.

Independent step-size screening on disjoint seeds 300–309 evaluates four values: $\mu_p = 0.005$, 0.01, 0.02, and 0.04. Both trackers select the same value, $\mu_p = 0.04$, when ranked first by failure count and then by median BER. Re-evaluating the same 20-seed 2 MHz SOP set with this common independently selected value gives single/paired failure counts of 3/20 versus 0/20 at $\theta = 0$, 2/20 versus 1/20 at $\theta = \pi/8$, and 2/20 versus 0/20 at $\theta = \pi/4$. Thus, the residual advantage is not produced by assigning a more favorable adaptation step to the paired tracker.

For the static-Jones replication, both trackers used the same assisted FIR branch streams and common downstream step size $\mu_p = 0.02$. The design treats the Jones matrices and stochastic seeds as separate sampling units. Two independently generated 12-matrix Haar-SU(2) ensembles were each crossed with ten stochastic realisations at 2 MHz. Ensemble A gives 9/120 versus 1/120 single-/paired-error failures, with paired-error BER lower in 110/120 records; ensemble B gives 12/120 versus 1/120 failures, with paired-error BER lower in 107/120 records. Jointly, across the 240 matched records, failures are 21/240 versus 2/240 and paired-error BER is lower in 217/240 records. The combined BER medians are 9.84×10^{-5} and 3.58×10^{-5}; the IQRs are $[5.37 \times 10^{-5}, 2.51 \times 10^{-4}]$ and $[1.79 \times 10^{-5}, 5.37 \times 10^{-5}]$, respectively. The median matched BER difference (single minus paired) is 5.37×10^{-5}. A crossed bootstrap with 20,000 replicates independently resamples the ten stochastic seeds and 24 Jones matrices with replacement and retains the full crossed product within each replicate. Percentile 95% intervals (2.5th–97.5th percentiles) are 2.08–15.83 percentage points for the absolute failure-fraction reduction and 3.58×10^{-5}–8.95×10^{-5} for the median paired BER difference.

3.3.2. Diagnostics and Training Dependence

An identity-channel, zero-linewidth sanity check gives zero BER for both trackers with receiver noise/ENOB disabled and negligible Tx quantization noise. The odd/even FIR-input condition numbers are 3.23 and 3.22, and the applied-coefficient magnitudes remain close to unity. Figure 6 and Table 5 then compare all three observed single-error failures with three non-failed records, one per SOP angle. Across these records, the FIR-input condition numbers remain approximately 3.0–3.4, providing no evidence of pathological input conditioning.

Table 5 gives BER, FIR-input condition numbers, branch powers, and median $|p|$ for the six records. Condition numbers range from approximately 3.0 to 3.4 across both groups.

Table 5. Representative 200 GSa/s 2 MHz trajectory diagnostics. Branch powers are mean squared magnitudes in simulation units over evaluated pairs 2200–16,171. Median $|p|$ uses the 600 displayed samples per trace. S and P denote the single-error and paired-error trackers, respectively.

Record	S BER	P BER	$\kappa(X_o)$	$\kappa(X_e)$	P_{a_o}	P_{b_o}	P_{a_e}	P_{b_e}	Med. $\|p\|$ S	Med. $\|p\|$ P
$\theta = 0$, 203 (S fail)	0.4296	8.95×10^{-6}	3.048	3.309	19.518	0.0198	0.0263	19.498	0.713	0.713
$\theta = 0$, 213 (S fail)	0.03016	2.68×10^{-5}	3.403	3.415	17.800	0.0196	0.0242	17.638	0.744	0.747
$\theta = \pi/8$, 204 (S fail)	0.3701	3.58×10^{-5}	3.251	3.132	17.084	2.952	3.018	16.833	0.702	0.702
$\theta = 0$, 200	4.47×10^{-5}	0	3.248	3.240	18.985	0.0236	0.0290	19.575	0.720	0.717
$\theta = \pi/8$, 201	1.25×10^{-4}	1.52×10^{-4}	3.197	3.204	15.944	2.721	2.793	15.945	0.729	0.729
$\theta = \pi/4$, 202	3.58×10^{-5}	2.68×10^{-5}	3.155	3.172	8.751	8.597	8.664	8.634	0.756	0.757

In the failure records, the single-error coefficient shows pronounced transient drops in $|p|$ and, for two records, persistent phase separation after the excursion. The non-failed records show no comparable magnitude collapse. Similar median $|p|$ values in Table 5 therefore mask transient behaviour; the underlying failure mechanism remains unresolved.

FIR-training dependence is evaluated at $\theta = \pi/4$ with the same downstream step $\mu_p = 0.02$. Three training conditions are compared: strict unity correction on the phase-noisy record; a phase-noise-free duplicate with the same data/channel/noise seed identities; and same-record true-phase-aided training, in which the exact simulated relative laser phase is removed only during FIR training while detector noise and ENOB remain present. In every case the FIR weights are frozen before downstream tracking, and the true phase is never supplied to either tracker. Table 6 summarizes the matched results.

Table 6. FIR training controls at 200 GSa/s and balanced SOP ($\theta = \pi/4$). The 2 MHz-only rows use the same seeds (200–209). Medians include all evaluated records in each row. S and P denote the single-error and paired-error trackers, respectively.

FIR Training	Linewidth per Laser	Failures		Median BER	
		S	P	S	P
Strict unity correction	500 kHz, 1 and 2 MHz	16/30	15/30	0.0425	0.0258
Standard assisted	2 MHz	0/10	0/10	1.34×10^{-5}	4.47×10^{-6}
Strict unity correction	2 MHz	9/10	9/10	0.3547	0.3179
Phase-noise-free duplicate	2 MHz	1/10	0/10	1.61×10^{-4}	4.03×10^{-5}
Same-record true-phase aid	2 MHz	1/10	0/10	1.79×10^{-4}	5.37×10^{-5}

On the same ten 2 MHz records, both trackers have 9/10 failures after strict unity training and 0/10 after standard assisted training. The duplicate-record and same-record true-phase-aided controls reduce failures to 1/10 versus 0/10, with paired-error BER lower in 10/10 and 8/10 matched records, respectively. Thus, FIR-training quality materially affects both trackers, and the receiver-level improvement cannot be attributed solely to the downstream paired-error update.

3.3.3. Benchmark and Fibre Controls

The Alamouti-adapted blind phase search (BPS) benchmark uses the same assisted FIR branch streams at balanced SOP and 2 MHz. The causal decision-distance window is fixed at 21 pairs throughout tuning; only the phase-hypothesis count $N_\phi \in \{32, 64\}$ and continuity weight $\lambda_c \in \{0.05, 20\}$ are screened on disjoint tuning seeds 300–302, giving four combinations. Ranking first by failures and then median BER selects $N_\phi = 64$ and $\lambda_c = 20$. For each record, one positive real gain $g \in [0.2, 1.5]$ is fitted from known preamble pairs 200–1999 by minimizing, over g, the mean over preamble pairs of the minimum known-symbol squared error across the same phase hypotheses; the first 100 preamble pairs resolve the initial phase ambiguity. Subsequent BPS choices minimize the 21-pair

causal decision-distance metric plus λ_c times squared wrapped phase change. On held-out seeds 200–209, the fitted gains span 0.7140–0.7579 and the median BERs are 1.34×10^{-5} (single-error), 4.47×10^{-6} (paired-error), and 2.77×10^{-4} (gain-calibrated BPS), with 0/10 failures for all three methods. BPS without preamble gain calibration gives 3/10 failures, so the calibrated result is used for comparison.

We compared back-to-back reception with 20 km fibre at restored received optical power. For each Tx-only and LO-only case at 1 and 2 MHz, the two distances used the same 20 seeds (200–219), balanced SOP ($\theta = \pi/4$), assisted FIR training, and a common downstream step size $\mu_p = 0.02$. Table 7 reports the BER spread and within-seed change for each tracker.

Table 7. Native 200 GSa/s BER results with received power restored at 20 km. BER and ΔBER entries are medians [minimum, maximum] in units of 10^{-6}. Failures count BER above 3×10^{-2} out of 20 seeds. ΔBER is the 20 km BER minus the back-to-back BER for the same seed, and up/down gives the numbers of seeds with positive and negative changes, respectively.

Phase Noise	Tracker	Back-to-Back		20 km, Power Restored		Within-Seed Change	
		BER	Failures	BER	Failures	ΔBER	Up/Down
Tx-only, 1 MHz	Single	0 [0, 0]	0/20	0 [0, 0]	0/20	0 [0, 0]	0/0
	Paired	0 [0, 0]	0/20	0 [0, 0]	0/20	0 [0, 0]	0/0
LO-only, 1 MHz	Single	0 [0, 8.95]	0/20	0 [0, 8.95]	0/20	0 [0, 8.95]	1/0
	Paired	0 [0, 0]	0/20	0 [0, 0]	0/20	0 [0, 0]	0/0
Tx-only, 2 MHz	Single	0 [0, 8.95]	0/20	0 [0, 8.95]	0/20	0 [−8.95, 8.95]	4/1
	Paired	0 [0, 0]	0/20	0 [0, 8.95]	0/20	0 [0, 8.95]	1/0
LO-only, 2 MHz	Single	0 [0, 26.84]	0/20	0 [0, 17.89]	0/20	0 [−8.95, 17.89]	2/3
	Paired	0 [0, 8.95]	0/20	0 [0, 8.95]	0/20	0 [−8.95, 8.95]	1/1

Across the power-restored fibre controls, the median BER and median within-seed BER change are zero in every condition, and neither tracker crosses the failure threshold. Equation (16) predicts phase variance rather than BER, so no quantitative BER penalty is inferred from these controls.

A separate three-seed (200–202) phase-variance diagnostic suppresses additive receiver noise by disabling detector thermal/shot noise and receiver ENOB, while setting Tx ENOB to 100 bits. A 512-hypothesis known-symbol one-phase estimate is obtained from the frozen branch streams for the 20 km and back-to-back records; common evaluation-pair indices are aligned, the two phase series are unwrapped, their difference is mean-removed, and its unbiased sample variance is calculated. For 1 and 2 MHz LO-only linewidths, the median variances are 1.5407×10^{-4} and 2.6667×10^{-4} rad^2, respectively, 27.6% and 37.4% below the corresponding Equation (16) predictions. The Tx-only medians are smaller. Because this fitted-phase-difference statistic is not a direct measurement of physical EEPN, quantitative agreement with Equation (16) is not claimed; the result supports only an LO-dominated trend compatible with EEPN.

4. Discussion

This section interprets the pair-level diagnostic and causal tracker results separately. It then considers the remaining tracker failures, relevant pre-FEC reference levels, and the limitations of the simulated evaluation.

4.1. Phase-Combiner Analysis

The pair-level searches are a secondary known-symbol combiner diagnostic rather than a practical carrier-recovery algorithm or a direct estimator of laser-induced intra-pair phase penalty. The two-phase search has one additional per-pair degree of freedom, and the nonzero zero-linewidth separation demonstrates that its advantage also contains

decision geometry, receiver noise, quantisation, fixed branch estimates, finite-grid effects, and fitting freedom. Equation (8) identifies one drift-induced term, but the reported one-phase/two-phase BER ratio is therefore not interpreted as a physical laser-phase penalty. Additional grid-refinement checks were performed for selected 2 MHz stress cases and for the complete ten-realisation 2 MHz, $\theta = \pi/4$ ensemble shown in Figure 2. One selected realisation exceeds the 5% 128×128–192×192 criterion, but changes by only 0.51% on further 192×192–256×256 refinement; the plotted ten-realisation ensemble mean changes by 2.08% in BER and 2.04% in diagnostic ratio.

4.2. Tracker Robustness and Training Dependence

The 100 GSa/s results are retained as reference-model diagnostics because the Nyquist-edge folding can influence the two adaptive trackers differently. They nevertheless show an important training interaction: assisted FIR training produces much larger single-versus-paired failure-count separations than strict neutral training.

The 200 GSa/s validation is the principal basis for the robustness conclusions. Neither tracker fails in the 20-seed balanced-SOP test at 2 MHz with the common step size $\mu_p = 0.02$, whereas across the two independent static Jones ensembles failures decrease from 21/240 with single-error tracking to 2/240 with paired-error tracking. The phase-noise-free duplicate-record and same-record true-phase-aided controls show that FIR-training quality materially affects both trackers. Thus, the paired-error downstream rule retains a reproducible advantage on these common frozen branch streams, but the receiver-level improvement cannot be attributed solely to the downstream paired-error update.

The correction is not phase-only: p_1, p_2, and their projected combination p are unconstrained complex coefficients and can therefore absorb residual complex gain as well as phase. The benefit is consequently described as complex branch correction rather than a purely phase mechanism.

The Alamouti-adapted BPS benchmark in Section 3.3.3 provides one practical carrier-phase-recovery reference at 2 MHz; broader comparison with additional feed-forward and pilot-aided methods remains future work. The 100 GSa/s ROP sweep is retained as an idealized digitally normalized diagnostic and is not used as principal robustness evidence.

The operational tracker-failure criterion of 3×10^{-2} used here is kept separate from FEC engineering limits: it is an algorithmic loss-of-robustness label for the tracker study. Throughout this paper, the pre-FEC BER reference of 10^{-2} is used as a comparison point rather than as the threshold of a particular FEC code. For context, a pre-FEC BER of 3.8×10^{-3} is commonly associated with a 7% overhead hard-decision FEC reference, whereas BER values near 2×10^{-2} generally require stronger soft-decision FEC schemes [39,40]. Post-FEC decoding is not evaluated, and no claim of system acceptability is attached to the operational tracker-failure criterion of 3×10^{-2}. In the 100 GSa/s reference model, the paired tracker also had fewer failures at thresholds of 10^{-2} and 2×10^{-2} in the assisted-training comparison (Section 3.2.6).

4.3. Limitations

The results are limited to the simulated system and evaluated parameter range. Accordingly, the practical interpretation is restricted to the static, lossless SOP transformations and idealized receiver model evaluated here. The 20 km extension uses an idealized linear-SSMF model with 4 dB attenuation and deterministic chromatic dispersion followed by frequency-domain compensation. Polarisation-mode dispersion, Kerr nonlinearity, optical-amplifier noise, residual frequency-offset estimation, component imperfections, and hardware-calibration drift are not included. The power-restored Tx-/LO-separated controls remove the 4 dB attenuation confound and show an LO-dominated fitted-phase-variance trend. The

measured LO-only statistic lies 27.6–37.4% below Equation (16), so quantitative agreement is not claimed; the comparison is limited to a trend compatible with EEPN. No measurable median-BER penalty is observed at those corrected control points. Experimental validation is still needed to assess the tracker advantage under practical transceiver impairments.

The SOP study uses three principal real rotations and a finite ensemble of 12 static Haar-distributed lossless Jones matrices. The corrected validation expands this to two independent 12-matrix static ensembles, but no time-varying SOP trajectory is simulated. The robustness claim is therefore explicitly restricted to static lossless SOP transformations and does not cover SOP tracking, polarisation-dependent loss, or PMD. In the 100 GSa/s reference model, under strict neutral training, both trackers fail all 120 2 MHz Jones records, so the isolated downstream paired-error tracker does not establish robust 2 MHz operation over that ensemble. The 100 GSa/s reference-model linewidth, SOP, Jones-matrix, and fibre analyses use ten independent realisations per condition, while the 100 GSa/s reference-model ROP sweep uses 30. The 200 GSa/s 2 MHz principal-SOP control uses 20 seeds per angle, and the corrected Jones replication crosses ten stochastic seeds with 24 independently generated static Jones matrices. Larger ensembles would further refine low-probability failure estimates.

At the 100 GSa/s detector sampling rate, the 50 GHz Nyquist frequency lies slightly below the approximately 50.5 GHz upper edge of the RRC-shaped heterodyne spectrum, so a small transition-band portion may fold at the first sampling operation. The 200 GSa/s control demonstrates that the large failure-rate advantage observed with the 100 GSa/s model is sampling-model dependent. The ROP study likewise omits practical AGC dynamics, ADC full-scale selection, clipping, and saturation. The ENOB model is an effective simulation impairment rather than a validated converter model. Accordingly, the study demonstrates empirical branch-tracking robustness under the stated conditions rather than formal control-theoretic stability.

5. Conclusions

This work evaluated projected paired-error complex branch correction in a four-branch Alamouti coherent-lite receiver under wide-linewidth laser phase noise. The 200 GSa/s validation is the principal basis for the robustness conclusion. Across two independent static, lossless Jones-matrix ensembles at 2 MHz per laser, projected paired-error correction reduces threshold failures from 21/240 to 2/240. The benefit varies with SOP and is smaller than the separation observed in the 100 GSa/s reference model, confirming sensitivity to front-end sampling assumptions.

FIR-training controls show that training quality materially affects both trackers; the receiver-level improvement therefore cannot be attributed solely to the downstream paired-error update. Within the simulated static, lossless SOP conditions and idealized receiver considered here, the paired update nevertheless provides a reproducible robustness advantage on common frozen branch streams. Time-varying SOP/PMD, realistic AGC/ADC behaviour, residual-frequency acquisition, and experimental hardware remain for future work.

Author Contributions: Conceptualization, M.A., T.S. and F.I.E.-N.; methodology, M.A., T.S. and F.I.E.-N.; validation, M.A., T.S. and F.I.E.-N.; formal analysis, M.A., T.S. and F.I.E.-N.; investigation, M.A., T.S. and F.I.E.-N.; writing—original draft preparation, M.A., T.S. and F.I.E.-N.; writing—review and editing, M.A., T.S. and F.I.E.-N.; visualization, M.A., T.S. and F.I.E.-N.; supervision, F.I.E.-N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: This publication was supported by the Scientific Research Projects Coordination Unit of Istanbul Yeni Yuzyil University.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Individual Results for the Principal Static SOPs

Table A1. Individual 2 MHz native-rate SOP results at $\mu_p = 0.02$. Entries show single-error/paired-error bit-error counts for each seed and SOP angle under standard assisted FIR training. Each count is out of 111,776 evaluated bits.

Seed	$\theta = 0$	$\pi/8$	$\pi/4$	Seed	$\theta = 0$	$\pi/8$	$\pi/4$
200	5/0	9/3	2/4	210	3/3	20/5	20/19
201	1/1	14/17	1/0	211	1/0	84/7148	109/1
202	3/4	6/4	4/3	212	1/0	27/6	4/2
203	48021/1	12/3	4/0	213	3371/3	184/35	7/4
204	1/1	41363/4	0/0	214	2/1	12/11	2/1
205	7/3	13/5	48/2	215	3/2	19/7	5/2
206	4/1	8/5	13/0	216	2/1	8/6	3/1
207	1/0	13/4	1/1	217	3/3	563/18	3/2
208	0/0	19/6	0/0	218	1/1	26/3	1/1
209	7/3	15/5	1/3	219	7/4	27/5	3/2

References

1. Zhang, J.; Jia, Z. Coherent passive optical networks for 100G/λ-and-beyond fiber access: Recent progress and outlook. *IEEE Netw.* **2022**, *36*, 116–123. https://doi.org/10.1109/MNET.005.2100604.
2. Jia, Z.; Zhang, H.; Choutagunta, K.; Campos, L.A. Coherent passive optical network: Applications, technologies, and specification development [Invited Tutorial]. *J. Opt. Commun. Netw.* **2025**, *17*, A71–A86. https://doi.org/10.1364/JOCN.535200.
3. Guo, D.; Zhou, J.; Zhang, J.; Gao, R.; Yu, C.; Xin, X. Flexible coherent PON for 6G fixed networks: Technologies and solutions. *IEEE Commun. Mag.* **2025**, *63*, 94–100. https://doi.org/10.1109/MCOM.002.2400622.
4. Shbair, W.; El Nahal, F. Coherent passive optical network technology for 5G. In Proceedings of the 2019 IEEE 7th Palestinian International Conference on Electrical and Computer Engineering (PICECE), Gaza, Palestine, 26–27 March 2019; pp. 1–4. https://doi.org/10.1109/PICECE.2019.8747183.
5. Shbair, W.W.; El-Nahal, F.I. Coherent passive optical network for 5G and beyond transport. *Optoelectron. Lett.* **2021**, *17*, 546–551. https://doi.org/10.1007/s11801-021-0178-3.
6. Shahpari, A.; Ferreira, R.M.; Luis, R.S.; Vujicic, Z.; Guiomar, F.P.; Reis, J.D.; Teixeira, A.L. Coherent access: A review. *J. Lightw. Technol.* **2017**, *35*, 1050–1058. https://doi.org/10.1109/JLT.2016.2623793.
7. Teixeira, A.; Lavery, D.; Ciaramella, E.; Schmalen, L.; Iiyama, N.; Ferreira, R.M.; Randel, S. DSP enabled optical detection techniques for PON. *J. Lightw. Technol.* **2020**, *38*, 684–695. https://doi.org/10.1109/JLT.2020.2964877.
8. Zhou, J.; Xing, Z.; Wang, H.; Zhang, K.; Chen, X.; Feng, Q.; Zheng, K.; Zhao, Y.; Dong, Z.; Gui, T.; et al. Flexible coherent optical access: Architectures, algorithms, and demonstrations. *J. Lightw. Technol.* **2024**, *42*, 1193–1202. https://doi.org/10.1109/JLT.2024.3355443.
9. Fang, X.; Zhu, Y.; Cai, X.; Hu, W.; He, Z.; Yu, S.; Zhang, F. Overcoming laser phase noise for low-cost coherent optical communication. *Nat. Commun.* **2024**, *15*, 6339. https://doi.org/10.1038/s41467-024-50439-1.
10. Oh, J.-Y.; Moon, S.-R.; Chang, S.-H.; Kang, H.-S. Alamouti-coded DSP algorithm with a simplified PTBC decoder for next-generation optical access networks. *Opt. Express* **2024**, *32*, 18727–18741. https://doi.org/10.1364/OE.525724.
11. Fan, R.; Li, J.; Gong, Z.; Shi, F.; Luo, M.; Li, X. Investigation of laser phase noise tolerance in simplified coherent PON with subcarrier multiplexing. *Opt. Express* **2024**, *32*, 39384–39393. https://doi.org/10.1364/OE.535117.
12. Fang, Y.; Zhu, H.; Cheng, X.; Yuan, P.; Lu, W.; Shi, D.; Ji, H.; Shieh, W. Integrated pre-amplified SC-PON receiver for 200G very high speed PON. *J. Lightw. Technol.* **2026**, *44*, 1765–1775.
13. Alamouti, S.M. A simple transmit diversity technique for wireless communications. *IEEE J. Sel. Areas Commun.* **1998**, *16*, 1451–1458.

14. Erkılınç, M.S.; Emmerich, R.; Habel, K.; Jungnickel, V.; Schmidt-Langhorst, C.; Schubert, C.; Freund, R. PON transceiver technologies for $\geq$ 50 Gbits/s per λ: Alamouti coding and heterodyne detection [Invited]. *J. Opt. Commun. Netw.* **2020**, *12*, A162–A170. https://doi.org/10.1364/JOCN.12.00A162.
15. Savory, S.J. Digital coherent optical receivers: Algorithms and subsystems. *IEEE J. Sel. Topics Quantum Electron.* **2010**, *16*, 1164–1179. https://doi.org/10.1109/JSTQE.2010.2044751.
16. Ip, E.; Kahn, J.M. Feedforward carrier recovery for coherent optical communications. *J. Lightw. Technol.* **2007**, *25*, 2675–2692. https://doi.org/10.1109/JLT.2007.902118.
17. Faruk, M.S.; Louchet, H.; Erkılınç, M.S.; Savory, S.J. DSP algorithms for recovering single-carrier Alamouti coded signals for PON applications. *Opt. Express* **2016**, *24*, 24083–24091. https://doi.org/10.1364/OE.24.024083.
18. Alhalabi, M.; Taşpınar, N. Performance investigation of bidirectional hybrid long-haul optical IM/DD OFDM WDM-PON using OOK-RSOA remodulation. *Müh. Bil. Araş. Derg.* **2022**, *4*, 54–61. https://doi.org/10.46387/bjesr.1082750.
19. Alhalabi, M.; Taşpınar, N.; El-Nahal, F.I. Bidirectional intensity-modulated/direct-detection optical OFDM WDM-PON system. *Optoelectron. Lett.* **2023**, *19*, 83–87. https://doi.org/10.1007/s11801-023-2126-x.
20. El-Nahal, F.I.; Abualoff, N.M. Coherent optical communication systems based on orthogonal frequency division multiplexing. *Optoelectron. Lett.* **2020**, *16*, 303–305. https://doi.org/10.1007/s11801-020-9166-2.
21. Taşpınar, N.; Alhalabi, M. Peak-to-average power ratio (PAPR) reduction using discrete invasive weed optimization (DIWO) in coherent detection optical OFDM (CO-OFDM) communication systems. *Wirel. Pers. Commun.* **2024**, *139*, 715–732. https://doi.org/10.1007/s11277-024-11645-8.
22. Luo, M.; Wu, D.; Li, W.; Zeng, T.; Zhou, L.; Meng, L.; Li, X.; Yu, S. 100 Gb/s (4×25 Gb/s) real-time coherent UDWDM-PON with a large power budget. *J. Opt. Commun. Netw.* **2020**, *12*, A204–A213. https://doi.org/10.1364/JOCN.12.00A204.
23. Zhang, J.; Yu, J.; Li, X.; Wang, K.; Zhou, W.; Xiao, J.; Zhao, L.; Pan, X.; Liu, B.; Xin, X. 200 Gbit/s/λ PDM-PAM-4 PON system based on intensity modulation and coherent detection. *J. Opt. Commun. Netw.* **2020**, *12*, A1–A8. https://doi.org/10.1364/JOCN.12.0000A1.
24. Kovacs, I.B.; Faruk, M.S.; Torres-Ferrera, P.; Savory, S.J. Simplified coherent optical network units for very-high-speed passive optical networks. *J. Opt. Commun. Netw.* **2024**, *16*, C1–C10. https://doi.org/10.1364/JOCN.514867.
25. Li, J.; Fan, R.; He, Q.; Luo, M.; Li, X. Bidirectional TFDM 200 G coherent PON with a simplified receiver at the ONU side. *Opt. Lett.* **2024**, *49*, 5127–5130. https://doi.org/10.1364/OL.536923.
26. Erkılınç, M.S.; Lavery, D.; Shi, K.; Thomsen, B.C.; Bayvel, P.; Killey, R.I.; Savory, S.J. Polarization-insensitive single-balanced photodiode coherent receiver for long-reach WDM-PONs. *J. Lightw. Technol.* **2016**, *34*, 2034–2041. https://doi.org/10.1109/JLT.2015.2507869.
27. Le, S.T.; Aref, V.; Cho, J. Single-ended coherent receiver. *J. Lightw. Technol.* **2022**, *40*, 1382–1399.
28. Kovacs, I.B.; Faruk, M.S.; Savory, S.J. A minimal coherent receiver for 200 Gb/s/λ PON downstream with measured 29 dB power budget. *IEEE Photon. Technol. Lett.* **2023**, *35*, 257–260. https://doi.org/10.1109/LPT.2023.3238666.
29. Faruk, M.S.; Li, X.; Savory, S.J. Experimental demonstration of 100/200-Gb/s/λ PON downstream transmission using simplified coherent receivers. In Proceedings of the Optical Fiber Communication Conference (OFC), San Diego, CA, USA, 6–10 March 2022; Paper Th3E.5.
30. AlQahtani, D.; El-Nahal, F.I. Short-preamble DSP for downstream Alamouti coherent-lite PON using gear-shifted LMS equalization. *Photonics* **2026**, *13*, 695. https://doi.org/10.3390/photonics13080695.
31. Pfau, T.; Hoffmann, S.; Noé, R. Hardware-efficient coherent digital receiver concept with feedforward carrier recovery for *M*-QAM constellations. *J. Lightw. Technol.* **2009**, *27*, 989–999.
32. Morsy-Osman, M.; Zhuge, Q.; Chen, L.R.; Plant, D.V. Feedforward carrier recovery via pilot-aided transmission for single-carrier systems with arbitrary *M*-QAM constellations. *Opt. Express* **2011**, *19*, 24331–24343. https://doi.org/10.1364/OE.19.024331.
33. Fatadin, I.; Ives, D.; Savory, S.J. Blind equalization and carrier phase recovery in a 16-QAM optical coherent system. *J. Lightw. Technol.* **2009**, *27*, 3042–3049.
34. Li, J.; Li, L.; Tao, Z.; Hoshida, T.; Rasmussen, J.C. Laser-linewidth-tolerant feed-forward carrier phase estimator with reduced complexity for QAM. *J. Lightw. Technol.* **2011**, *29*, 2358–2364.
35. Souto, D.V.; Olsson, B.-E.; Larsson, C.; Mello, D.A.A. Joint-polarization and joint-subchannel carrier phase estimation for 16-QAM optical systems. *J. Lightw. Technol.* **2012**, *30*, 3185–3191.
36. Wu, M.-W.; Kam, P.-Y. Soft-decision-aided, maximum-likelihood carrier phase estimation for coherent optical QAM. *J. Lightw. Technol.* **2013**, *31*, 3443–3452.
37. Zhang, T.; Pan, Y.; Fang, X.; Zhu, Y.; Zhou, X.; Zhang, F. Residual-carrier-aided coherent-lite receiver for wide-linewidth links using low-complexity DD-ML phase tracking. *Opt. Lett.* **2026**, *51*, 2112–2115. https://doi.org/10.1364/OL.589664.
38. Shieh, W.; Ho, K.-P. Equalization-enhanced phase noise for coherent-detection systems using electronic digital signal processing. *Opt. Express* **2008**, *16*, 15718–15727. https://doi.org/10.1364/OE.16.015718.

39. Mizuochi, T. Recent progress in forward error correction and its interplay with transmission impairments. *IEEE J. Sel. Top. Quantum Electron.* **2006**, *12*, 544–554.
40. International Telecommunication Union. *Forward Error Correction for High Bit-Rate DWDM Submarine Systems*; ITU-T Recommendation G.975.1; ITU: Geneva, Switzerland, 2004.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.