information

MDPI

Article

Phase-Aware Multidimensional Reconciliation for Continuous-Variable Quantum Key Distribution

Yanling Liang, Jisheng Dai *, Han Hai and Xue-Qin Jiang

College of Information and Intelligent Science, Donghua University, Shanghai 201620, China; 2252232@mail.dhu.edu.cn (Y.L.); hhai@dhu.edu.cn (H.H.); xqjiang@dhu.edu.cn (X.-Q.J.)
* Correspondence: jsdai@dhu.edu.cn

Abstract

Continuous-variable quantum key distribution (CV-QKD) relies on efficient information reconciliation to establish identical secret key sequences from correlated modulation data and measurement outcomes. Accurate quadrature measurements are essential for reliable reconciliation, but residual phase mismatch between the signal and local oscillator distorts these measurements in coherent receivers and degrades decoding performance. To improve robustness against residual phase mismatch in homodyne-based CV-QKD, we propose a phase-aware multidimensional reconciliation method that jointly performs syndrome-based low-density parity-check (LDPC) decoding and expectation-maximization-based residual-phase refinement. We derive a discrete-time detection model that explicitly incorporates the receiver's pulse-dependent quadrature-selection phase and construct phase-dependent virtual observations through multidimensional orthogonal mapping. The LDPC decoder provides posterior information for calculating the posterior spherical-symbol moments used in residual-phase refinement, while the refined phase is fed back to update the virtual observations and decoding likelihoods for the next iteration. Numerical results show that the proposed method achieves a lower residual-phase-estimation RMSE than disclosure-based cross-correlation estimation, BER and FER close to those under perfect phase compensation, and a higher finite-size secret key rate.

Keywords: continuous-variable quantum key distribution (CV-QKD); balanced homodyne detection; phase noise; multidimensional reconciliation; LDPC decoding; expectation maximization (EM)

Check for updates

Academic Editors: Wenbin Yu, Yadang Chen and Chengjun Zhang

Received: 31 July 2026
Revised: 1 September 2026
Accepted: 3 September 2026
Published: 5 September 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.

1. Introduction

Reliable cryptographic key establishment is fundamental to secure modern information infrastructures. Quantum key distribution (QKD) offers a compelling solution by enabling distant users to establish shared secret keys whose security is rooted in the principles of quantum mechanics rather than assumptions about computational hardness [1]. In continuous-variable QKD (CV-QKD), a major implementation of QKD, Alice encodes information in the continuous quadratures of optical fields, while Bob measures the received states using coherent detection. Among its implementations, the widely studied Gaussian-modulated coherent-state protocol modulates the two conjugate quadratures of coherent states according to Gaussian distributions, with the receiver obtaining continuous-valued outcomes through homodyne or heterodyne detection [2–4]. These operations can be realized using lasers, optical modulators, balanced receivers, and digital signal-processing techniques widely employed in coherent optical communication systems, making CV-QKD

promising for high-rate operation and integration with existing optical communication infrastructures [5]. After measurement, information reconciliation is required to transform Alice's modulation data and Bob's measurement outcomes into identical binary key sequences with limited public communication, and its efficiency and frame error rate directly affect the achievable secret key rate, particularly in low-signal-to-noise-ratio (SNR) regimes [6,7]. Recent studies have pursued complementary improvements through rate-adaptive reconciliation based on raptor-like LDPC codes and high-efficiency reconciliation based on optimized MGC-LDPC constructions [8,9].

Despite these advantages, practical coherent receivers are subject to various implementation impairments, among which signal–local-oscillator (LO) phase mismatch is particularly critical because it directly affects quadrature measurement. In a homodyne receiver, the received optical field is projected onto a quadrature by interference with a strong LO, with the projection direction determined by their relative phase [10]. Random-basis homodyne detection switches between the two conjugate quadratures through a pulse-dependent LO phase shift [11]. Although the basis-selection sequence is available during reconciliation, an unknown residual signal–LO phase mismatch may persist because of laser phase noise, transmission-induced phase fluctuations, and imperfect synchronization [12,13]. This mismatch rotates the effective measurement axis, distorts the correlation between the modulation data and measurement outcomes, and consequently degrades decoding performance. Accurate residual-phase estimation is therefore essential for reconstructing the measured quadrature information and achieving reliable information reconciliation.

Existing phase-recovery and compensation methods can be grouped according to how the phase information is obtained and processed. Pilot- or reference-assisted schemes recover the relative signal–LO phase from weak reference pulses or multiplexed optical tones [12,14]. Beyond direct pilot recovery, Bayesian inference has been used both to avoid transmitting a strong reference pulse and to improve pilot-assisted phase estimation [15,16]. State-space and smoothing methods, including scalar, vector, and unscented Kalman filters as well as Savitzky–Golay filtering, further suppress noise in the available phase observations [17–19]. In addition, multi-pilot phase recovery combined with Wiener filtering has been experimentally investigated to improve phase-estimation accuracy in local-local-oscillator (LLO) CV-QKD systems [20]. Machine-learning-based approaches have also been investigated for phase-error estimation and other CV-QKD post-processing tasks [21,22]. Alternatively, correlation-based methods reserve and disclose a subset of correlated samples for phase estimation and infer the residual phase from the correlation statistics between the local modulation data and the corresponding measurement outcomes [23,24]. For random-basis homodyne detection, the correlation calculation must also account for the receiver's pulse-dependent quadrature-selection phase. The estimation accuracy is governed by both the SNR and the number of disclosed samples. Using more samples improves the phase estimate but reduces data utilization, or, equivalently, increases the transmission overhead when the retained data-block length is fixed. These trade-offs make robust phase estimation particularly challenging over long-distance, low-SNR channels.

A recent joint parameter-estimation and multidimensional-reconciliation framework has shown that decoder-generated posterior information can be used to refine channel gain and noise variance [25]. However, that framework does not address residual signal–LO phase mismatch or the pulse-dependent quadrature selection of random-basis homodyne detection. In this setting, the residual phase determines the virtual observations and their decoding likelihoods, while the decoder posterior contains information about the transmit-

ted symbols that can, in turn, support phase refinement. This reciprocal dependence calls for phase estimation and information reconciliation to be treated jointly.

To address this gap, we propose a phase-aware multidimensional reconciliation method that couples syndrome-based low-density parity-check (LDPC) decoding with expectation-maximization (EM)-based residual-phase refinement. The proposed method deliberately retains the established multidimensional orthogonal mapping and syndrome-based LDPC structure to ensure compatibility with practical CV-QKD reconciliation. Its key methodological novelty is to embed the residual signal–LO phase explicitly in the phase-dependent virtual observations and likelihoods and to refine it through a locally linearized EM-based update driven by decoder-derived posterior spherical-symbol moments. The updated phase estimate is then used to reconstruct the virtual observations and LLRs for the next LDPC iteration, thereby forming a closed LDPC–EM loop rather than treating phase recovery and reconciliation as two separate processing stages. The main contributions are threefold. First, we derive a discrete-time balanced homodyne detection model that explicitly distinguishes the known pulse-dependent quadrature-selection phase from the unknown residual signal–LO phase. Second, we construct phase-dependent virtual observations through multidimensional orthogonal mapping and formulate their likelihoods within a syndrome-constrained probabilistic model, from which the LDPC decoder obtains both bit decisions and posterior spherical-symbol moments. Third, these posterior moments are incorporated into an EM update of the residual phase, and the refined estimate is fed back to reconstruct the virtual observations and decoding likelihoods for the next iteration.

The remainder of this paper is organized as follows. Section 2 establishes the discrete-time balanced homodyne detection model and reviews the conventional cross-correlation-based residual-phase estimation method. Section 3 presents the proposed phase-aware multidimensional reconciliation method, including syndrome-based LDPC decoding and EM-based residual-phase refinement. Section 4 reports and discusses the numerical results in terms of phase-estimation accuracy, decoding performance, robustness to intra-block phase drift, and finite-size secret key rate. Finally, Section 5 concludes the paper.

2. Data Model and Existing Method Review

This section first derives the discrete-time balanced homodyne detection observation model with pulse-dependent quadrature selection and then reviews an existing cross-correlation method for residual-phase estimation.

2.1. Data Model

In a Gaussian-modulated CV-QKD system, Alice encodes key information into the two conjugate quadratures of coherent states, whereas Bob uses homodyne detection to measure one quadrature of each received state and obtain real-valued electrical samples for reconciliation and decoding [3]. A balanced homodyne detector (BHD) performs this measurement by interfering the weak quantum signal with a strong LO, whose phase determines the projection direction of the measurement [26]. The standard homodyne-based Gaussian-modulated coherent-state CV-QKD protocol implements Bob's random basis selection by applying a pulse-dependent phase shift to the LO [3,12].

Figure 1 illustrates the core physical architecture of the BHD employed at Bob's receiver. Consider a data block consisting of L pulse time slots, where $i \in \{1, 2, \ldots, L\}$ indexes the pulse time slot within the block. For the i-th pulse time slot, the real-valued

optical electric field of the signal received by Bob after propagation through the channel is expressed as a function of continuous time t:

$$\mathcal{E}_{s,i}(t) = \sqrt{\eta_{\mathrm{ch}}}\, A_{s,i} \cos(\omega_s t + \theta_i + \theta_s) \tag{1}$$

where η_{ch} denotes the channel power transmittance, the pair $(A_{s,i}, \theta_i)$ gives the polar-coordinate representation of Alice's two Gaussian-modulated quadrature components in the i-th pulse, and ω_s and θ_s denote the signal carrier frequency and phase offset, respectively. The remaining channel and detection noise contributions are incorporated later into the discrete-time observation model.

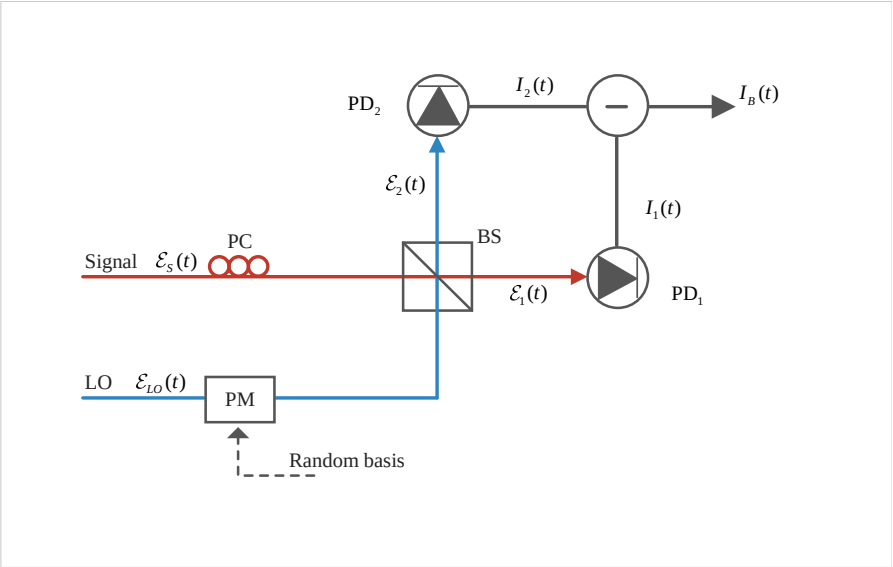


Figure 1. The core physical architecture of the balanced homodyne detector employed at Bob's receiver. LO—local oscillator; PC—polarization controller; PM—phase modulator; BS—beam splitter; PD—photodiode.

Random quadrature-basis selection is implemented by controlling the LO phase. Specifically, the LO passes through the phase modulator, which introduces a pulse-dependent random phase shift $\theta_{\mathrm{PM},i} \in \{0, \pi/2\}$ to select one of the two orthogonal quadratures for each pulse. The resulting LO field in the i-th time slot is

$$\mathcal{E}_{\mathrm{LO},i}(t) = A_{\mathrm{LO}} \cos(\omega_{\mathrm{LO}} t + \theta_{\mathrm{LO}} + \theta_{\mathrm{PM},i}), \tag{2}$$

where A_{LO}, ω_{LO}, and θ_{LO} denote the constant amplitude envelope, angular carrier frequency, and carrier phase offset of the LO, respectively. Bob will announce the basis-setting sequence after detection, making $\theta_{\mathrm{PM},i}$ known to Alice during reconciliation. The signal and LO fields are then combined at the 50:50 beam splitter. Following the standard beam-splitter transformation used in balanced homodyne detection [27], the optical fields at the two output ports of the BS are given by

$$\mathcal{E}_{1,i}(t) = \frac{1}{\sqrt{2}}[\mathcal{E}_{s,i}(t) + \mathcal{E}_{\mathrm{LO},i}(t)], \tag{3}$$

$$\mathcal{E}_{2,i}(t) = \frac{1}{\sqrt{2}}[\mathcal{E}_{s,i}(t) - \mathcal{E}_{\mathrm{LO},i}(t)]. \tag{4}$$

After photodetection by the two photodiodes, the corresponding photocurrents are derived directly from these output fields, without introducing separate optical-power variables [28]:

$$I_{1,i}(t) = R\left\langle (\mathcal{E}_{1,i}(t))^2 \right\rangle = \frac{R}{2}\left\langle \mathcal{E}_{s,i}^2(t) + \mathcal{E}_{\mathrm{LO},i}^2(t) + 2\mathcal{E}_{s,i}(t)\mathcal{E}_{\mathrm{LO},i}(t) \right\rangle, \tag{5}$$

and

$$I_{2,i}(t) = R\left\langle (\mathcal{E}_{2,i}(t))^2 \right\rangle = \frac{R}{2}\left\langle \mathcal{E}_{s,i}^2(t) + \mathcal{E}_{\mathrm{LO},i}^2(t) - 2\mathcal{E}_{s,i}(t)\mathcal{E}_{\mathrm{LO},i}(t) \right\rangle, \tag{6}$$

where R is the photodiode responsivity, and $\langle f(t)\rangle \triangleq T_{\mathrm{d}}^{-1}\int_{t-T_{\mathrm{d}}/2}^{t+T_{\mathrm{d}}/2} f(\tau)\mathrm{d}\tau$ denotes the moving average over a photodiode detection interval of duration T_{d} centered at t. This operation models the finite photodiode response time and averages out rapidly oscillating optical-frequency terms. Differential subtraction of the two photocurrents yields

$$I_{B,i}(t) = I_{1,i}(t) - I_{2,i}(t) = 2R\left\langle \mathcal{E}_{s,i}(t)\cdot\mathcal{E}_{\mathrm{LO},i}(t)\right\rangle. \tag{7}$$

Equation (7) follows directly from Equations (3) and (4) through the square-law photodetection in Equations (5) and (6): the signal and LO self-intensity terms cancel under balanced subtraction, while the signal–LO cross term is retained and carries the relative-phase information. Assuming matched signal and LO carrier frequencies for ideal homodyne detection, i.e., $\omega_s = \omega_{\mathrm{LO}}$, substitution of (1) and (2) into (7) gives

$$I_{B,i}(t) = RA_{\mathrm{LO}}\sqrt{\eta_{\mathrm{ch}}}\,A_{s,i}\cos\left(\theta_i + \theta_s - \theta_{\mathrm{LO}} - \theta_{\mathrm{PM},i}\right). \tag{8}$$

Pulse sampling and normalization by the estimated deterministic gain $RA_{\mathrm{LO}}\sqrt{\eta_{\mathrm{ch}}}$ yield the i-th observation

$$I_{B,i} = A_{s,i}\cos(\theta_i + \phi - \theta_{\mathrm{PM},i}), \tag{9}$$

where $\phi = \theta_s - \theta_{\mathrm{LO}}$ denotes the residual signal–LO phase mismatch caused by transmission-induced phase fluctuations and laser phase noise. The residual phase is assumed approximately constant within one reconciliation block while varying slowly between blocks. Thus, the observation depends on both Bob's known quadrature-selection phase $\theta_{\mathrm{PM},i}$ and the unknown residual phase ϕ. By stacking the observations from all L pulse slots and taking additive Gaussian detection noise into account, the discrete-time observation model is written as

$$\mathbf{y} = \mathbf{x}(\phi) + \mathbf{z}, \tag{10}$$

where $\mathbf{y} = [I_{B,1}, I_{B,2}, \ldots, I_{B,L}]^{\mathrm{T}}$, $\mathbf{x}(\phi) = [A_{s,1}\cos(\theta_1 + \phi - \theta_{\mathrm{PM},1}), \ldots, A_{s,L}\cos(\theta_L + \phi - \theta_{\mathrm{PM},L})]^{\mathrm{T}}$, the superscript $(\cdot)^{\mathrm{T}}$ denotes the transpose operation, and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I}_L)$ denotes the normalized additive Gaussian noise vector. Here, $\mathbf{I}_L$ is the $L \times L$ identity matrix, and σ^2 is the normalized noise variance per sampled pulse, including shot noise, electronic noise, and other Gaussian channel and detection noise contributions.

The resulting discrete-time model captures the BHD output while distinguishing Bob's announced quadrature-selection phase $\theta_{\mathrm{PM},i}$ from the unknown residual phase ϕ. During reconciliation, Alice can directly incorporate the known basis setting $\theta_{\mathrm{PM},i}$, whereas constructing the phase-dependent signal samples $\mathbf{x}(\phi)$ requires an estimate of ϕ. An inaccurate phase estimate produces mismatched virtual observations and consequently distorts the log-likelihood ratios (LLRs) supplied to the LDPC decoder, thereby degrading reconciliation performance. As a baseline, we first adapt the cross-correlation phase estimator (CCPE) in [24] to the observation model in (9) and examine its limitations.

2.2. Review of Cross-Correlation Phase Estimation and Its Limitations

The CCPE in [24] estimates an unknown phase offset by correlating the received observations of the disclosed samples with locally generated references under different candidate phase rotations. We extend this method to estimate the residual phase ϕ in (9), where each disclosed-sample reference must also account for Bob's known random basis-selection phase $\theta_{\mathrm{PM},i}$.

In addition to the L data slots in each block, N_{p} data samples are publicly disclosed for phase estimation. For a candidate phase $\Delta\varphi$, Alice constructs the normalized reference for the publicly disclosed data sample indexed by $p = 1, 2, \ldots, N_{\mathrm{p}}$ as

$$\widetilde{I}_{A,p}(\Delta\varphi) = A_{s,p} \cos\big(\theta_p + \Delta\varphi - \theta_{\mathrm{PM},p}\big). \tag{11}$$

The candidate phase is evaluated using the cross-correlation

$$C(\Delta\varphi) = \frac{1}{N_{\mathrm{p}}} \sum_{p=1}^{N_{\mathrm{p}}} \widetilde{I}_{B,p} \widetilde{I}_{A,p}(\Delta\varphi), \tag{12}$$

where $\widetilde{I}_{B,p}$ denotes Bob's normalized observation in the p-th disclosed sample. Compared with the conventional correlation expression in [24], the pulse-dependent term $-\theta_{\mathrm{PM},p}$ explicitly incorporates the random quadrature selected for each homodyne measurement.

Under zero-mean isotropic Gaussian modulation and independent zero-mean detection noise, its expectation is proportional to $\cos(\Delta\varphi - \phi)$. The residual-phase estimate is therefore obtained as

$$\widehat{\phi}_{\mathrm{cc}} = \underset{\Delta\varphi \in [-\pi, \pi)}{\arg\max}\ C(\Delta\varphi), \tag{13}$$

where the subscript cc denotes cross-correlation estimation. The maximization can be implemented by scanning a set of candidate phase rotations, as in [24]. After estimating $\widehat{\phi}_{\mathrm{cc}}$, Alice uses it to construct the phase-dependent signal samples $A_{s,i} \cos(\theta_i + \widehat{\phi}_{\mathrm{cc}} - \theta_{\mathrm{PM},i})$ for $i = 1, 2, \ldots, L$, while Bob's measured observations $I_{B,i}$ remain unchanged. All L data-slot samples are then supplied to multidimensional reconciliation, whereas the additional disclosed samples are used only for phase estimation.

Although straightforward, this estimator has several limitations. The publicly disclosed samples consume transmission resources without contributing to secret key extraction, and the estimation accuracy depends strongly on the SNR and the number of available disclosed samples. Moreover, the phase is estimated before reconciliation and then treated as fixed, preventing the posterior information generated during LDPC decoding from correcting an inaccurate estimate. These limitations motivate the phase-aware reconciliation method developed in the next section, which iteratively refines the residual phase using decoder-generated posterior information and feeds the updated estimate back into the virtual observations and decoding likelihoods.

3. Proposed Method

This section presents the proposed phase-aware reconciliation method. We first construct the phase-dependent multidimensional virtual channel, then describe syndrome-based LDPC decoding and EM-driven residual-phase refinement, and finally discuss information leakage and computational complexity.

3.1. Phase-Dependent Multidimensional Reverse Reconciliation

Following the block-partitioning procedure used in conventional multidimensional reverse reconciliation [29], Bob divides the observation sequence $\mathbf{y} \in \mathbb{R}^L$ into $G = L/D$ sub-blocks, each containing D consecutive observations, assuming that L is divisible by D. Conventional multidimensional reconciliation typically restricts D to $\{1, 2, 4, 8\}$, and this work adopts the largest conventional dimension, $D = 8$. The cross-rotation method extends reconciliation beyond eight dimensions [30]. Although the proposed phase-aware framework can in principle be generalized to such higher-dimensional constructions, we restrict the present analysis to $D = 8$ in order to focus on the phase-aware reconciliation

mechanism. The observation sequence can then be written as $\mathbf{y} = [\mathbf{y}_1^{\mathrm{T}}, \mathbf{y}_2^{\mathrm{T}}, \ldots, \mathbf{y}_G^{\mathrm{T}}]^{\mathrm{T}}$, with each sub-block satisfying $\mathbf{y}_g \in \mathbb{R}^D$.

Bob then generates a random binary sequence $\mathbf{c} \in \{0,1\}^L$ and partitions it according to the same block structure: $\mathbf{c} = [\mathbf{c}_1^{\mathrm{T}}, \mathbf{c}_2^{\mathrm{T}}, \ldots, \mathbf{c}_G^{\mathrm{T}}]^{\mathrm{T}}$, where each binary sub-block is given by $\mathbf{c}_g = [c_{1,g}, c_{2,g}, \ldots, c_{D,g}]^{\mathrm{T}}$, with $c_{k,g} \in \{0,1\}$ denoting its k-th component. Given an (L, K) LDPC code specified by its parity-check matrix $\mathbf{H} \in \{0,1\}^{(L-K)\times L}$, where K denotes the dimension of the LDPC code, the syndrome associated with Bob's binary vector $\mathbf{c}$ is computed as $\mathbf{s} = \mathbf{Hc} \bmod 2$. Each binary sub-block is then mapped onto a D-dimensional unit spherical vector as

$$\mathbf{u}_g = \frac{1}{\sqrt{D}}[(-1)^{c_{1,g}}, \cdots, (-1)^{c_{D,g}}]^{\mathrm{T}}. \tag{14}$$

This binary spherical mapping is a standard construction in multidimensional reconciliation [29]. It embeds Bob's randomly generated raw-key bits $\mathbf{c}_g$ into the constant-norm vector $\mathbf{u}_g$. Bob can then construct an orthogonal transformation that maps his normalized continuous observation to $\mathbf{u}_g$. After the transformation is communicated, Alice applies it to her correlated data to obtain a noisy virtual observation of the key-bearing vector $\mathbf{u}_g$, thereby forming a binary-input virtual channel for likelihood calculation and LDPC-based error correction. The normalized observation block $\mathbf{y}_g/\|\mathbf{y}_g\|$ is then mapped to the target spherical vector $\mathbf{u}_g$ through a block-dependent orthogonal matrix $\mathbf{M}_g \in \mathbb{R}^{D\times D}$ satisfying

$$\frac{\mathbf{M}_g \mathbf{y}_g}{\|\mathbf{y}_g\|} = \mathbf{u}_g. \tag{15}$$

Following the conventional multidimensional rotation construction [25,29], $\mathbf{M}_g$ is expressed as

$$\mathbf{M}_g = \sum_{k=1}^{D} \alpha_{k,g} \mathbf{A}_k, \tag{16}$$

where $\alpha_{k,g} \triangleq \mathbf{u}_g^{\mathrm{T}} \mathbf{A}_k \mathbf{y}_g / \|\mathbf{y}_g\|$ is the k-th mapping coefficient and $\mathbf{A}_k \in \mathbb{R}^{D\times D}$ is a fixed orthogonal-design matrix defined in the Appendix of [29]. Clearly, $\mathbf{M}_g$ is orthogonal and satisfies $\mathbf{M}_g^{\mathrm{T}} \mathbf{M}_g = \mathbf{I}_D$. Bob only discloses the scalar coefficients $\{\alpha_{k,g}\}_{k=1}^{D}$ and the norm $\|\mathbf{y}_g\|$ for each block, together with the syndrome $\mathbf{s}$, over the authenticated classical channel.

Upon receiving this side information, Alice reconstructs $\mathbf{M}_g$ locally. The resulting phase-dependent virtual observation is

$$\mathbf{v}_g(\phi) = \frac{\mathbf{M}_g \mathbf{x}_g(\phi)}{\|\mathbf{y}_g\|} = \frac{\mathbf{M}_g \mathbf{y}_g}{\|\mathbf{y}_g\|} - \frac{\mathbf{M}_g \mathbf{z}_g}{\|\mathbf{y}_g\|} = \mathbf{u}_g + \tilde{\mathbf{z}}_g, \tag{17}$$

where $\tilde{\mathbf{z}}_g \triangleq -\mathbf{M}_g \mathbf{z}_g / \|\mathbf{y}_g\|$ denotes the equivalent noise vector. Under the isotropic Gaussian noise assumption, the orthogonal mapping preserves the Gaussian form of the noise, while the normalization by $\|\mathbf{y}_g\|$ scales its variance. Consequently, the effective noise variance for the g-th block is $\sigma_g^2 = \sigma^2 / \|\mathbf{y}_g\|^2$. Based on the phase-dependent virtual channel in (17) and the effective noise variance σ_g^2 defined above, the conditional probability density function of the virtual observation sequence given the spherical vector sequence can be expressed as

$$\begin{aligned} p(\mathbf{v}(\phi)|\mathbf{u};\phi) &= \prod_{g=1}^{G} p(\mathbf{v}_g(\phi)|\mathbf{u}_g;\phi) \\ &= \prod_{g=1}^{G} \prod_{k=1}^{D} \mathcal{N}\left(v_{k,g}(\phi) \middle| u_{k,g}, \sigma_g^2\right) \end{aligned} \tag{18}$$

where $v_{k,g}(\phi)$ and $u_{k,g}$ denote the k-th element of the virtual observation vector $\mathbf{v}_g(\phi)$ and the spherical vector $\mathbf{u}_g$, respectively. Thus, Equation (17) characterizes the virtual-channel relation in terms of the unknown residual phase ϕ. At iteration m, $\mathbf{v}_g(\phi^{(m)})$ is deterministically reconstructed using the fixed mapping information and is not an additional physical observation; Equation (18) specifies the corresponding Gaussian likelihood under the adopted virtual-channel model. Under the deterministic binary-to-spherical mapping in (14), the conditional distribution of the spherical vector sequence $\mathbf{u}$ given Bob's binary vector $\mathbf{c}$ factorizes as

$$p(\mathbf{u}|\mathbf{c}) = \prod_{g=1}^{G} p(\mathbf{u}_g|\mathbf{c}_g) = \prod_{g=1}^{G}\prod_{k=1}^{D} p(u_{k,g}|c_{k,g}), \tag{19}$$

where the element-wise conditional density is

$$p(u_{k,g}|c_{k,g}) = \delta\left(u_{k,g} - \frac{(-1)^{c_{k,g}}}{\sqrt{D}}\right), \tag{20}$$

where $\delta(\cdot)$ denotes the Dirac delta function. The one-to-one mapping also gives the inverse conditional probability density function

$$p(\mathbf{c}|\mathbf{u}) = \prod_{g=1}^{G} p(\mathbf{c}_g|\mathbf{u}_g) = \prod_{g=1}^{G}\prod_{k=1}^{D} p(c_{k,g}|u_{k,g}), \tag{21}$$

where

$$p(c_{k,g}|u_{k,g}) = \delta\left(c_{k,g} - \frac{1-\sqrt{D}u_{k,g}}{2}\right). \tag{22}$$

Note that $p(\mathbf{u}|\mathbf{c})$ and $p(\mathbf{c}|\mathbf{u})$ encode the same deterministic mapping constraint in opposite directions.

Assuming that Bob's random bits are uniformly distributed, the prior distribution of the binary vector $\mathbf{c}$ can be factorized as

$$p(\mathbf{c}) = \prod_{g=1}^{G} p(\mathbf{c}_g) = \prod_{g=1}^{G}\prod_{k=1}^{D} p(c_{k,g}), \tag{23}$$

where each bit satisfies $p(c_{k,g}=0) = p(c_{k,g}=1) = \frac{1}{2}$. Equivalently, the prior distribution of the spherical vector sequence is

$$p(\mathbf{u}) = \prod_{g=1}^{G} p(\mathbf{u}_g) = \prod_{g=1}^{G}\prod_{k=1}^{D} p(u_{k,g}), \tag{24}$$

with $p\left(u_{k,g}=\frac{1}{\sqrt{D}}\right) = p\left(u_{k,g}=-\frac{1}{\sqrt{D}}\right) = \frac{1}{2}$. The LDPC syndrome imposes a parity-check constraint on Bob's binary vector $\mathbf{c}$. Given the syndrome $\mathbf{s} = \mathbf{Hc} \bmod 2$ transmitted by Bob, this constraint can be expressed as

$$p(\mathbf{s}|\mathbf{c}) = \delta(\mathbf{s} - (\mathbf{Hc} \bmod 2)), \tag{25}$$

Combining the spherical-vector prior, the deterministic spherical-to-binary mapping factor, the phase-dependent virtual-channel likelihood, and the syndrome constraint, the joint probabilistic model can be factorized as

$$\begin{aligned} p(\mathbf{v}(\phi), \mathbf{u}, \mathbf{c}, \mathbf{s}; \phi) &= p(\mathbf{c})p(\mathbf{u}|\mathbf{c})p(\mathbf{v}(\phi)|\mathbf{u};\phi)p(\mathbf{s}|\mathbf{c}) \\ &= p(\mathbf{u})p(\mathbf{c}|\mathbf{u})p(\mathbf{v}(\phi)|\mathbf{u};\phi)p(\mathbf{s}|\mathbf{c}). \end{aligned} \tag{26}$$

https://doi.org/10.3390/info17090861

The corresponding factor graph is shown in Figure 2, where blank circles and black squares represent variable nodes and factor nodes, respectively. The variable nodes $\mathcal{C}_g$ and $\mathcal{U}_g$ correspond to the binary block $\mathbf{c}_g$ and the spherical vector $\mathbf{u}_g$, respectively. The factor node $\mathcal{C}_g\mathcal{U}_g$ represents the deterministic mapping $p(\mathbf{c}_g|\mathbf{u}_g)$, while the syndrome-check node $\mathcal{SC}$ represents the constraint factor $p(\mathbf{s}|\mathbf{c})$. Each spherical variable node is connected to its prior factor $p(\mathbf{u}_g)$ and the phase-dependent virtual-channel likelihood $p(\mathbf{v}_g(\phi)|\mathbf{u}_g;\phi)$. The posterior information propagated through the graph is subsequently used by the EM module to refine the residual phase parameter ϕ.

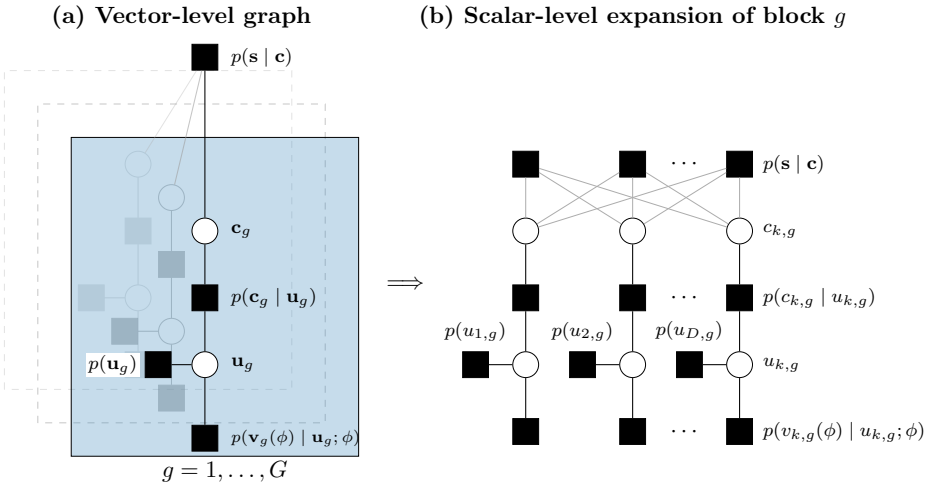


Figure 2. Factor graph corresponding to the probabilistic model in (26).

3.2. Phase-Aware Syndrome-Based LDPC Decoding

Message passing provides a systematic means of performing probabilistic inference on a factor graph by exchanging posterior information between variable and factor nodes [31,32]. Multidimensional reconciliation can couple LDPC decoding with parameter estimation, so that the decoder output refines the unknown parameters and the updated parameters, in turn, improve the decoding likelihoods [25]. The derivation below follows a message-passing procedure similar to that in [25], with likelihood information propagated from the spherical variables through the deterministic mapping to the binary variables and LDPC extrinsic information propagated in the reverse direction. In the present model, however, Bob's pulse-dependent basis-selection phase is explicitly incorporated into $\mathbf{v}_g(\phi)$. Consequently, the virtual observations must be reconstructed using the current residual-phase estimate at each joint iteration, and the corresponding virtual-channel likelihoods and LLRs must be updated accordingly.

Denote the message passed from node a to node b by $\Delta_{a\to b}(\cdot)$, and use m to index the joint LDPC–EM iteration, with $\phi^{(m)}$ representing the residual-phase estimate at that iteration. At iteration m, the current estimate $\phi^{(m)}$ is used to construct $\mathbf{v}_g(\phi^{(m)})$ and evaluate the corresponding virtual-channel likelihood. According to the factor graph in Figure 2, the message passed from the spherical variable node $\mathcal{U}_g$ to the mapping factor node $\mathcal{U}_g\mathcal{C}_g$ therefore combines the spherical-vector prior with this likelihood:

$$\begin{aligned}\Delta_{\mathcal{U}_g\to\mathcal{U}_g\mathcal{C}_g}(\mathbf{u}_g) &\propto p(\mathbf{u}_g)\,p\Big(\mathbf{v}_g(\phi^{(m)})|\mathbf{u}_g;\phi^{(m)}\Big)\\ &= \prod_{k=1}^{D}\left[\frac{1}{2}\mathcal{N}\Big(v_{k,g}(\phi^{(m)}) \mid \frac{1}{\sqrt{D}},\sigma_g^2\Big)\delta\Big(u_{k,g}-\frac{1}{\sqrt{D}}\Big)\right.\\ &\qquad\left.+\frac{1}{2}\mathcal{N}\Big(v_{k,g}(\phi^{(m)}) \mid -\frac{1}{\sqrt{D}},\sigma_g^2\Big)\delta\Big(u_{k,g}+\frac{1}{\sqrt{D}}\Big)\right].\end{aligned} \tag{27}$$

The message passed from the mapping factor node $\mathcal{U}_g\mathcal{C}_g$ to the binary variable node $\mathcal{C}_g$ is obtained by marginalizing over $\mathbf{u}_g$:

$$\begin{aligned}\Delta_{\mathcal{U}_g\mathcal{C}_g\to\mathcal{C}_g}(\mathbf{c}_g) &\propto \sum_{\mathbf{u}_g} p(\mathbf{c}_g|\mathbf{u}_g)\Delta_{\mathcal{U}_g\to\mathcal{U}_g\mathcal{C}_g}(\mathbf{u}_g)\\ &\propto \prod_{k=1}^{D}\left[\delta(c_{k,g})\mathcal{N}\left(v_{k,g}(\phi^{(m)}) \mid \frac{1}{\sqrt{D}}, \sigma_g^2\right)\right.\\ &\qquad \left.+\,\delta(c_{k,g}-1)\mathcal{N}\left(v_{k,g}(\phi^{(m)}) \mid -\frac{1}{\sqrt{D}}, \sigma_g^2\right)\right],\end{aligned} \tag{28}$$

which yields the initial LLR delivered to the LDPC decoder for bit $c_{k,g}$ as

$$\mathrm{LLR}_{k,g}(\phi^{(m)}) = \ln\frac{\mathcal{N}\left(v_{k,g}(\phi^{(m)}) \mid 1/\sqrt{D}, \sigma_g^2\right)}{\mathcal{N}\left(v_{k,g}(\phi^{(m)}) \mid -1/\sqrt{D}, \sigma_g^2\right)} = \frac{2}{\sqrt{D}\sigma_g^2}v_{k,g}(\phi^{(m)}). \tag{29}$$

The phase-dependent LLRs are supplied to the LDPC decoder, where the sum-product algorithm performs conventional syndrome-based message passing. Here, $p_{k,g}^0$ and $p_{k,g}^1$ denote the extrinsic probabilities produced by the parity-check constraints for $c_{k,g} = 0$ and $c_{k,g} = 1$, respectively. The subsequent messages from the syndrome-check nodes to the binary variables and then through the deterministic mapping to the spherical variables follow the standard sum-product updates in [25]. We therefore omit these intermediate steps and directly define the resulting factorized belief representation of the spherical variables as

$$b^{(m)}\left(\mathbf{u} \mid \mathbf{v}(\phi^{(m)}), \mathbf{s}; \phi^{(m)}\right) \triangleq \prod_{g=1}^{G}\prod_{k=1}^{D}\left[p_{k,g}^{(m)}\delta\left(u_{k,g} - \frac{1}{\sqrt{D}}\right) + \left(1 - p_{k,g}^{(m)}\right)\delta\left(u_{k,g} + \frac{1}{\sqrt{D}}\right)\right], \tag{30}$$

where

$$p_{k,g}^{(m)} = \frac{p_{k,g}^0 \exp\left(\mathrm{LLR}_{k,g}(\phi^{(m)})\right)}{p_{k,g}^0 \exp\left(\mathrm{LLR}_{k,g}(\phi^{(m)})\right) + p_{k,g}^1}.$$

The corresponding component-wise posterior mean and variance are

$$\mathbb{E}_m[u_{k,g}] = \frac{2p_{k,g}^{(m)} - 1}{\sqrt{D}}, \tag{31}$$

$$\mathrm{Var}_m(u_{k,g}) = \frac{4p_{k,g}^{(m)}\left(1 - p_{k,g}^{(m)}\right)}{D}. \tag{32}$$

The posterior mean replaces the latent spherical symbol in the maximization step, while the posterior variance contributes to the expected squared-error objective.

3.3. EM-Driven Residual-Phase Refinement

The physical observation relation in Equation (9), the multidimensional mapping in Equations (15)–(17), and the Gaussian virtual-channel likelihood in Equation (18) follow directly from the adopted signal and virtual-channel models. The EM objective is then formed using this likelihood and the posterior moments provided by the LDPC decoder.

Since the residual phase ϕ is unknown in practice, the EM algorithm is used to estimate it iteratively. Taking the sequence $\mathbf{u}$ (and its underlying binary sequence $\mathbf{c}$) as the hidden variables, the unobservable terms in the log-likelihood function are replaced by their conditional expectations derived in the expectation step. Consequently, the unknown

phase drift is updated at the $(m+1)$-th iteration by maximizing the expected complete-data log-likelihood:

$$\begin{aligned}\phi^{(m+1)} &= \arg\max_{\phi} \mathbb{E}_m\{\ln p(\mathbf{v}(\phi), \mathbf{u}, \mathbf{c}, \mathbf{s}; \phi)\} \\ &= \arg\min_{\phi} \sum_{g=1}^{G} \frac{1}{2\sigma_g^2}\Big(\|\mathbf{v}_g(\phi) - \mathbb{E}_m\{\mathbf{u}_g\}\|^2 + \mathrm{tr}(\mathrm{Cov}_m(\mathbf{u}_g))\Big),\end{aligned} \tag{33}$$

where $\mathrm{Cov}_m(\cdot)$ denotes the covariance with respect to the decoder belief distribution in (30) under the current phase estimate $\phi^{(m)}$. Since $\mathrm{tr}\big(\mathrm{Cov}_m(\mathbf{u}_g)\big)$ is independent of ϕ, it can be discarded during the maximization step. The resulting objective is

$$J^{(m)}(\phi) = \sum_{g=1}^{G} \frac{1}{2\sigma_g^2}\big\|\mathbf{v}_g(\phi) - \mathbb{E}_m\{\mathbf{u}_g\}\big\|^2. \tag{34}$$

Taking the derivative of the objective function in (34) with respect to ϕ yields

$$\frac{\partial J^{(m)}(\phi)}{\partial \phi} = \sum_{g=1}^{G} \frac{1}{\sigma_g^2}\Big(\mathbf{v}_g'(\phi)\Big)^{\mathrm{T}}\big[\mathbf{v}_g(\phi) - \mathbb{E}_m\{\mathbf{u}_g\}\big], \tag{35}$$

where $\mathbf{v}_g'(\phi) \triangleq \partial \mathbf{v}_g(\phi)/\partial\phi$.

According to the observation model in (9), the random quadrature-selection phase $\theta_{\mathrm{PM},i}$ is incorporated into $\mathbf{v}_g(\phi)$ and hence also into $\mathbf{v}_g'(\phi)$. Since $\mathbf{v}_g(\phi)$ is nonlinear in ϕ, we use a first-order approximation around the current estimate [33], i.e.,

$$\mathbf{v}_g(\phi^{(m)} + \Delta\phi^{(m)}) \approx \mathbf{v}_g(\phi^{(m)}) + \mathbf{v}_g'(\phi^{(m)})\Delta\phi^{(m)}. \tag{36}$$

Equation (36) is the only explicit analytical approximation in the phase-refinement step. It yields the phase increment in Equation (37) and the update in Equation (38). Substituting (36) into the linearized objective and setting the derivative with respect to $\Delta\phi^{(m)}$ to zero gives

$$\Delta\phi^{(m)} = \frac{\sum_{g=1}^{G} \frac{1}{\sigma_g^2}\Big(\mathbf{v}_g'(\phi^{(m)})\Big)^{\mathrm{T}}\Big[\mathbb{E}_m\{\mathbf{u}_g\} - \mathbf{v}_g(\phi^{(m)})\Big]}{\sum_{g=1}^{G} \frac{1}{\sigma_g^2}\Big\|\mathbf{v}_g'(\phi^{(m)})\Big\|^2}. \tag{37}$$

The residual phase is then updated as

$$\phi^{(m+1)} = \phi^{(m)} + \Delta\phi^{(m)}. \tag{38}$$

As the successive updates in Equation (38) refine $\phi^{(m+1)}$ toward the true residual phase ϕ under the adopted model, the expansion point in Equation (36) moves closer to the true value and the local first-order approximation becomes increasingly accurate. This alternating procedure follows the same iterative parameter-refinement principle as the off-grid refinement in [33]: the current parameter estimate is used to evaluate the model and posterior statistics, and the refined parameter is fed back to reconstruct the model for the next iteration. The phase estimate is iteratively refined until the syndrome constraint is satisfied or the phase increment falls below the prescribed threshold.

The proposed joint message-passing algorithm is initialized with $\phi^{(0)} = 0$. At the m-th iteration, the current phase estimate is used to reconstruct the virtual observations and calculate the phase-dependent LLRs in (29). Syndrome-based LDPC decoding then produces extrinsic information, from which the posterior probabilities defined above and the posterior mean in (31) are evaluated. The posterior mean is subsequently passed to

the EM module to compute the phase increment in (37) and update the residual phase according to (38). The refined phase estimate is fed back to update the virtual observations and LLRs for the next iteration, while the disclosed mapping coefficients, block norms, syndrome, and effective noise variances remain fixed. This alternating procedure continues until the syndrome constraint is satisfied or the phase increment becomes sufficiently small.

3.4. Information Leakage and Complexity Analysis

The phase-aware extension changes the local inference procedure but does not alter the side information required by conventional multidimensional reverse reconciliation. The public transcript consists of Bob's basis-setting sequence, the syndrome $\mathbf{s}$, the block norms $\{\|\mathbf{y}_g\|\}$, and the mapping coefficients $\{\alpha_{k,g}\}$. Bob's basis-setting sequence is generated independently of the binary vector $\mathbf{c}$, and the block norms depend only on his Gaussian observations. Moreover, under the isotropic Gaussian model and the orthogonal construction in (16), the mapping information is statistically independent of the spherical vector $\mathbf{u}$, as established for conventional multidimensional reconciliation [25,29]. Consequently, the basis settings, block norms, and mapping coefficients introduce no additional leakage about the reconciled binary sequence. The syndrome, by contrast, reveals parity information about $\mathbf{c}$. Its leakage is bounded by the number of transmitted syndrome bits, $L - K$, and is accounted for during privacy amplification in the same manner as in conventional syndrome-based reconciliation. Starting from $\phi^{(0)} = 0$, the residual-phase refinement is performed locally, at Alice's side, using the phase-dependent virtual observations and the posterior probabilities generated during LDPC decoding. The joint LDPC–EM iterations therefore require no additional data disclosure for phase estimation and introduce no public communication beyond the standard reconciliation transcript.

Each joint iteration reconstructs the phase-dependent virtual observations, updates the LLRs in (29), performs syndrome-based LDPC message passing, evaluates the posterior mean in (31), and updates the residual phase according to (38). Let E_{H} denote the number of edges in the Tanner graph of $\mathbf{H}$ and let I_{iter} denote the number of joint iterations. Computing the phase-dependent signal samples and the posterior mean requires $\mathcal{O}(L)$ operations per iteration. The products $\mathbf{M}_g\mathbf{x}_g(\phi)$ require $\mathcal{O}(GD^2) = \mathcal{O}(LD)$ operations, while the vector inner products involved in the scalar phase update require $\mathcal{O}(L)$ operations. Syndrome-based LDPC message passing requires $\mathcal{O}(E_{\mathrm{H}})$ operations per iteration. Hence, the overall complexity is

$$\mathcal{O}(I_{\mathrm{iter}}(E_{\mathrm{H}} + LD)).$$

Since the reconciliation dimension D is fixed and small, this reduces to $\mathcal{O}(I_{\mathrm{iter}}(E_{\mathrm{H}} + L))$ and retains the same asymptotic order as conventional multidimensional reconciliation. The additional phase-refinement operations are linear in L, require neither large-scale matrix inversion nor high-dimensional numerical optimization, and are typically dominated by LDPC decoding.

4. Numerical Results and Discussion

4.1. Simulation Setup and Performance Metrics

For conciseness, the proposed joint expectation maximization and low-density parity-check decoding method is denoted by EM–LDPC throughout this section. Following the evaluation framework in [25], we assess both estimation accuracy and reconciliation performance using the residual-phase root-mean-square error (RMSE), bit error rate (BER), and frame error rate (FER). The simulations use the discrete-time homodyne model in (10). The residual signal–LO phase is constant within each data block and independently drawn as $\phi \sim \mathcal{U}[-\pi/6, \pi/6]$ between blocks. An Advanced Television Systems Committee (ATSC)

3.0 LDPC code of length $L = 64{,}800$ and rate $R_c = 0.2$ is used with eight-dimensional reconciliation. Each operating point is evaluated over $N_{MC} = 500$ independent Monte Carlo trials. The EM–LDPC receiver starts from $\phi^{(0)} = 0$. The no-compensation baseline fixes $\widehat{\phi} = 0$, whereas the perfect-compensation reference is supplied with the true residual phase. For the CCPE benchmark in [24], the two parties publicly disclose N_p additional samples from the same phase-stable frame, and these samples do not enter LDPC decoding or key extraction. Thus, every receiver decodes the same L data symbols, but CCPE requires a total of $L + N_p$ transmitted symbols per block, whereas EM–LDPC requires only the L data symbols. The principal simulation settings are summarized in Table 1.

To ensure a controlled comparison, all receivers process the same generated L-symbol data block in each Monte Carlo frame, including the same residual phase, noisy channel realization, multidimensional mapping information, syndrome, LDPC code, code rate, SNR, and maximum iteration parameter. The phase-estimation information is method dependent by design. For CCPE, N_p additional samples are independently generated under the same SNR and residual phase and are publicly disclosed exclusively for phase estimation; they do not enter LDPC decoding or key extraction. EM–LDPC uses $N_p = 0$ and performs phase refinement locally using Alice's modulation data, the standard reconciliation information associated with the same L-symbol block, and decoder-derived posterior spherical-symbol moments. Thus, the comparison holds the reconciliation and decoding conditions fixed while explicitly accounting for the different phase-estimation resources required by the two methods. For $N_p = 100$, 1000, and 10,000, CCPE requires 64,900, 65,800, and 74,800 transmitted symbols, corresponding to overheads of 0.1543%, 1.5432%, and 15.4321% relative to the L-symbol reconciliation block, respectively. The no-compensation receiver is included as a diagnostic baseline, whereas perfect compensation is supplied with the true residual phase. Perfect compensation uses the true residual phase as an oracle reference.

Table 1. Principal settings used in the numerical evaluation.

Parameter	Setting
Reconciliation dimension	$D = 8$
LDPC code family	ATSC 3.0
LDPC block length	$L = 64{,}800$
LDPC code rate	$R_c = 0.2$
Maximum iteration parameter (all receivers)	200
Residual-phase distribution	$\mathcal{U}[-\pi/6, \pi/6]$
Initial phase estimate	$\phi^{(0)} = 0$
CCPE disclosed phase-estimation samples	$N_p \in \{100, 1000, 10{,}000\}$
EM–LDPC disclosed phase-estimation samples	$N_p = 0$
Independent blocks per operating point	$N_{MC} = 500$
SNR range for phase-estimation RMSE	-4 to 12 dB

For the normalized observation model, the linear signal-to-noise ratio is $\gamma = \mathbb{E}[x_i^2(\phi)]/\sigma^2$, and $\mathrm{SNR} = 10\log_{10}\gamma$ is reported in decibels. The phase-estimation accuracy is measured by

$$\mathrm{RMSE}_\phi = \sqrt{\frac{1}{N_{MC}}\sum_{n=1}^{N_{MC}}\left(\widehat{\phi}_n - \phi_n\right)^2}, \tag{39}$$

where $\widehat{\phi}_n$ and ϕ_n denote the estimated and true residual phases in the n-th Monte Carlo trial, respectively. The BER is the number of erroneous decoded bits divided by the total number of decoded bits, and the FER is the fraction of decoded blocks containing at least one bit error. With $N_{MC} = 500$, a zero FER means that no decoding failure was observed and

should be interpreted subject to the corresponding Monte Carlo resolution. As in [30,34], the reconciliation efficiency is

$$\beta = \frac{R_c}{C(\gamma)}, \qquad C(\gamma) = \frac{1}{2}\log_2(1+\gamma), \tag{40}$$

where $C(\gamma)$ is the capacity per real dimension of the corresponding Gaussian channel. For the fixed code rate R_c, the SNR and β are therefore in one-to-one correspondence: a larger β places the code closer to the channel limit and generally increases the decoding-failure probability.

4.2. Residual-Phase Estimation Accuracy

Figure 3 compares the residual-phase estimation RMSE over the SNR range from -4 to 12 dB under the static phase model. The RMSE decreases with increasing SNR for every method because the phase information contained in each observation becomes more reliable. The CCPE curves also show a strong dependence on the number of publicly disclosed samples: increasing N_p from 100 to 10,000 substantially improves the cross-correlation estimate by reducing the finite-sample estimation error. This behavior is consistent with the limitation identified in Section 2.2 and in [24], namely that the accuracy of correlation-based phase estimation depends jointly on the SNR and the number of disclosed samples.

EM–LDPC achieves the lowest RMSE throughout the simulated SNR range without requiring additional data disclosure for phase estimation. At -4 dB, its RMSE is approximately 6.58×10^{-3} rad, compared with 1.90×10^{-2} rad for CCPE even with $N_p = 10{,}000$. At 12 dB, the corresponding values are approximately 1.04×10^{-3} and 1.06×10^{-2} rad, respectively. The gain comes from the different use of the available observations. CCPE forms a single estimate from the disclosed samples before reconciliation and keeps that estimate fixed, whereas EM–LDPC uses the posterior spherical-symbol moments generated during decoding to refine ϕ over the entire length-L data block. The updated phase then improves the virtual observations and likelihoods used in the next iteration. This iterative feedback mitigates the effect of errors introduced during a separate phase-estimation stage and extracts phase information from samples that already contribute to reconciliation.

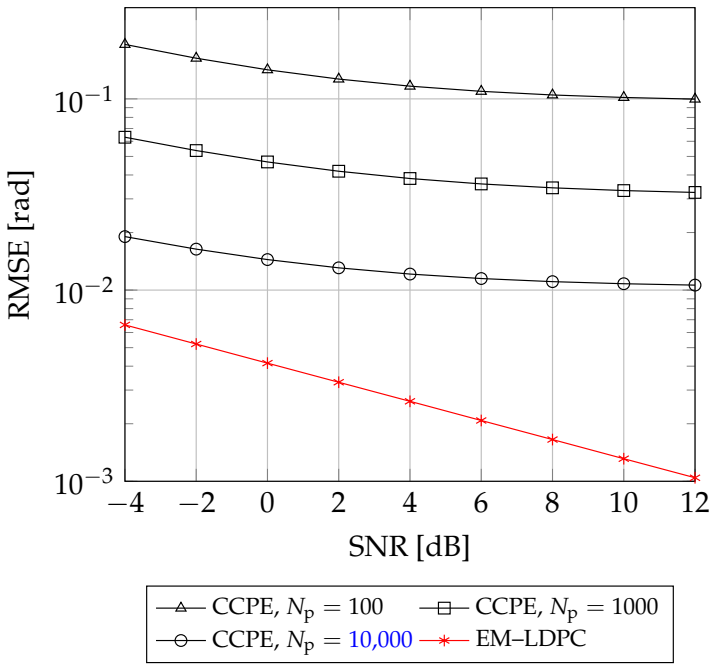
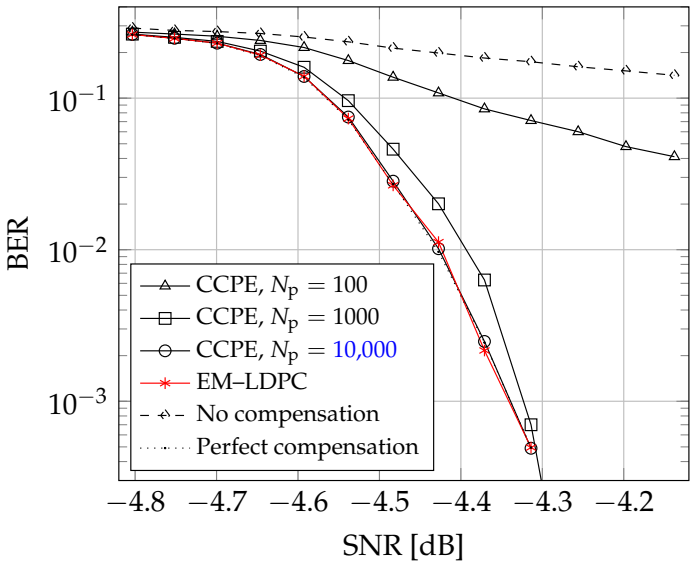


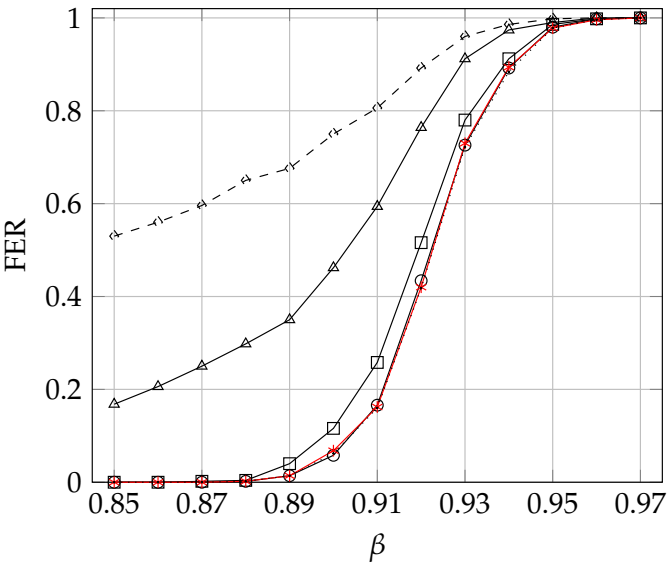
Figure 3. Residual-phase estimation RMSE versus SNR for EM–LDPC and CCPE with different numbers of disclosed samples.

4.3. Reconciliation Performance

Figure 4a compares the BER in the low-SNR waterfall region. Without phase compensation, the decoder does not exhibit a clear waterfall because the residual phase produces mismatched virtual observations and likelihoods. CCPE progressively restores the waterfall as N_p increases, confirming that increasing the number of publicly disclosed samples can significantly improve reconciliation performance. However, the curves obtained with $N_p = 100$ and 1000 remain visibly separated from the perfect-compensation reference. CCPE approaches the reference only when N_p is increased to 10,000. By contrast, EM–LDPC requires no additional data disclosure for phase estimation and nearly overlaps the perfect-compensation curve, yielding the best overall BER performance among the practical receivers.



(**a**)



(**b**)

Figure 4. Reconciliation performance under different phase-compensation strategies: (**a**) BER versus SNR and (**b**) FER versus reconciliation efficiency β.

Figure 4b reports FER as a function of the reconciliation efficiency β. As expected from (40), the FER rises as β approaches unity because the fixed-rate LDPC code operates

closer to the effective-channel capacity. The same ordering observed in the BER results is maintained: no compensation gives the highest FER, CCPE improves as more data samples are publicly disclosed, and EM–LDPC closely follows perfect compensation throughout the transition region. Hence, decoder-aided phase refinement preserves the high-efficiency operating region without allocating additional transmission resources to phase estimation.

The comparison also clarifies the cost of closing the gap with conventional CCPE. Its strongest result requires $N_{\mathrm{p}} = 10{,}000$ additional publicly disclosed data samples for a data block of length $L = 64{,}800$, an overhead equal to approximately 15.4% of the data-block length. These disclosed samples do not contribute to reconciliation or key extraction, and the resulting phase estimate cannot be refined using the posterior symbol probabilities subsequently generated by the LDPC decoder. EM–LDPC avoids both limitations: it requires no additional data disclosure for phase estimation and continually updates the phase estimate within the decoding loop. Although the BER and FER curves of EM–LDPC and CCPE with a large number of disclosed samples can become nearly indistinguishable once both estimates are sufficiently accurate for successful decoding, Figure 3 shows that EM–LDPC retains a clear advantage in phase-estimation accuracy. Its principal benefit is therefore not only near-oracle decoding performance, but also improved symbol efficiency and greater robustness to an inaccurate initial phase estimate.

The BER and FER results therefore constitute conditional decoding comparisons for the common L-symbol reconciliation block. The additional CCPE samples are not included in the BER or FER denominators because they do not enter decoding.

4.4. Robustness to Intra-Block Phase Drift

The preceding results are based on a blockwise-static residual-phase model. The residual phase is assumed to remain approximately constant within one reconciliation block, but it may vary between different blocks. This approximation is intended to represent operating conditions in which the residual phase evolves slowly relative to the duration of one reconciliation block. In practical CV-QKD implementations, however, time-dependent phase noise and drift may introduce finite intra-block phase variation.

To examine the robustness of this approximation under finite intra-block phase variation, we introduce a controlled linear phase drift within each reconciliation block. Specifically, the true residual phase of the i-th symbol is generated as

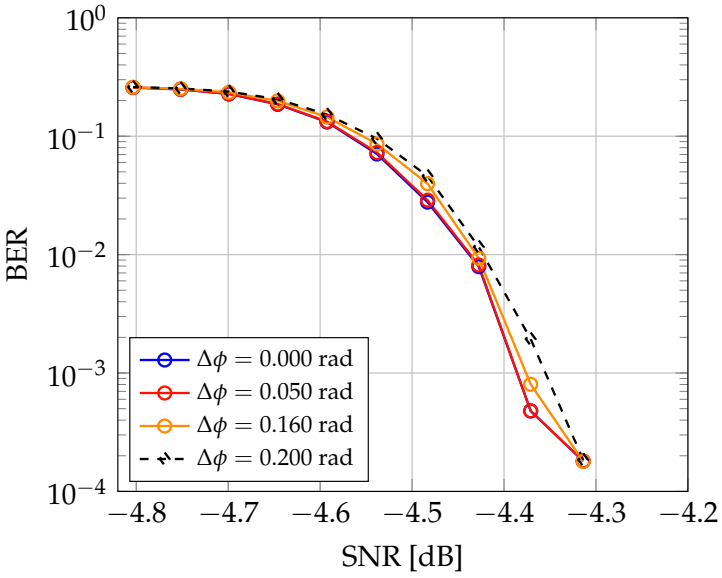
$$\phi_i = \phi_0 + \frac{i-1}{L-1}\Delta\phi, \qquad i = 1, \ldots, L, \tag{41}$$

where $\Delta\phi$ denotes the accumulated phase drift from the first to the last symbol of the block, and $\Delta\phi = 0$ corresponds to the zero-drift case. The robustness sweep is performed as an independent Monte Carlo run from the blockwise-static evaluation in Figure 4. Consequently, small numerical differences between the $\Delta\phi = 0$ curve and the main baseline, such as the FER values of 0.410 and 0.420, respectively, at $\beta = 0.92$, reflect the finite Monte Carlo variability of the two independent 500-frame evaluations rather than a difference in the underlying phase model. The received observations are generated using the symbol-dependent phases $\{\phi_i\}$. The EM–LDPC receiver remains unchanged and estimates one scalar residual phase for the entire block. Thus, this experiment constitutes a controlled model-mismatch test rather than a dynamic phase-tracking implementation. We evaluate $\Delta\phi \in \{0, 0.05, 0.10, 0.12, 0.14, 0.15, 0.16, 0.20\}$ rad over 500 frames at each operating point. For clarity, the figure below displays the zero-drift baseline, a small-drift case, the upper end of the dense drift sweep, and a larger model-mismatch case.

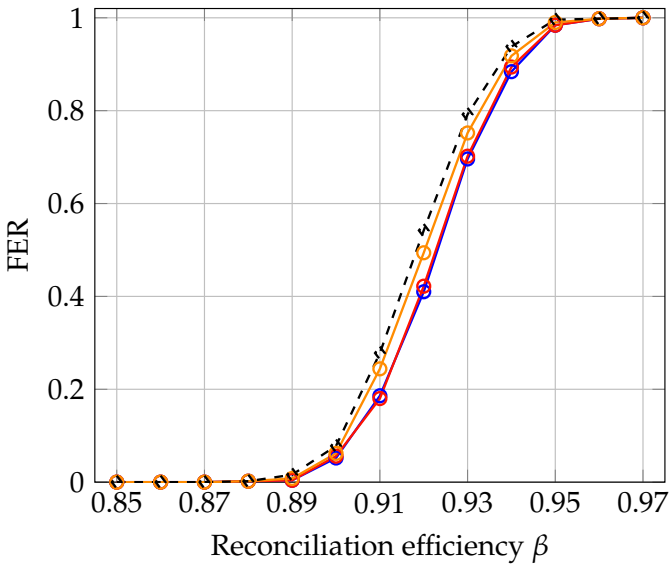
Figure 5 compares the BER and FER of the proposed EM–LDPC method under representative intra-block phase-drift conditions. The BER is shown versus SNR, while the

corresponding FER is shown versus reconciliation efficiency β, following the same representation used in the main decoding-performance results.

Across the tested drift levels, the BER and FER exhibit an overall progressive degradation as the accumulated intra-block phase drift increases. The effect is most visible near the LDPC decoding waterfall region. For example, at $\beta = 0.92$, the BER values for $\Delta\phi = 0$, 0.05, 0.16, and 0.20 rad are 0.0707, 0.0729, 0.0861, and 0.0953, respectively. The corresponding FER values are 0.410, 0.422, 0.494, and 0.542. The intermediate drift values $\Delta\phi = 0.10$, 0.12, 0.14, and 0.15 rad yield FERs of 0.442, 0.460, 0.472, and 0.482 at the same operating point. These values confirm the overall progressive trend. Small crossings between adjacent drift levels at a few operating points are consistent with the finite Monte Carlo resolution of 500 frames. The results therefore do not indicate a distinct critical phase-drift threshold.



(**a**)



(**b**)

Figure 5. Robustness of the proposed EM–LDPC method to intra-block linear phase drift: (**a**) BER and (**b**) FER for representative accumulated phase drifts of $\Delta\phi = 0$, 0.05, 0.16, and 0.20 rad.

For the small accumulated drift of $\Delta\phi = 0.05$ rad, the decoding performance remains close to the baseline case. Thus, limited intra-block variation introduces only a modest mismatch under the tested conditions. As the accumulated drift increases, however, the receiver's single scalar phase becomes progressively less representative of the actual intra-block phase trajectory. The larger degradation observed at $\Delta\phi = 0.20$ rad confirms that an evaluation restricted to the static baseline can become increasingly optimistic when time-dependent intra-block variation is present. The adopted phase model is therefore most appropriate when the residual phase evolves slowly relative to the reconciliation-block duration. The results show progressive degradation in reconciliation performance as this approximation is increasingly violated.

The linear drift considered here is a controlled robustness test. It is not a complete temporal model of practical phase noise. Moreover, the tested values should not be interpreted as typical drift levels of a particular CV-QKD implementation. The adopted receiver model includes the pulse-dependent quadrature-selection phase and additive Gaussian channel and detection noise. However, the simulations do not separately model the temporal behavior of laser and LO phase noise. Residual frequency and timing errors, gain-calibration uncertainty, detector imbalance, and detector-bandwidth effects are also not explicitly included. Faster or more irregular intra-block phase dynamics would require an extended dynamic phase model. Assessing these implementation effects and validating the proposed reconciliation procedure on an optical CV-QKD platform are left for future work.

4.5. Secret Key Rate Performance

To examine the end-to-end impact of phase estimation on key generation, we evaluate the finite-size secret key rate (SKR) following [35]. For a reconciliation block of L data symbols, the SKR per transmitted symbol is calculated as

$$K = \frac{L}{N_{\mathrm{PE}} + L + N_{\mathrm{p}}}(1 - \mathrm{FER})[\beta I(A;B) - \chi^{\epsilon}(B;E) - \Delta(L)], \tag{42}$$

where N_{PE} is the number of samples used for parameter estimation, N_{p} is the number of publicly disclosed data samples used for phase estimation, $I(A;B)$ is the mutual information between the legitimate parties, $\chi^{\epsilon}(B;E)$ is the finite-size upper bound on the Holevo information available to an eavesdropper, and $\Delta(L)$ denotes the finite-size privacy-amplification penalty. In this evaluation, $N_{\mathrm{PE}} = L = 64{,}800$ and the security parameter is $\epsilon = 10^{-10}$. The channel transmittance at distance ℓ is modeled as $T(\ell) = 10^{-\alpha\ell/10}$ with fiber loss $\alpha = 0.2$ dB/km. The excess noise, homodyne-detection efficiency, and electronic noise are set to 0.01 shot-noise units (SNU), 0.606, and 0.041 SNU, respectively. For each value of β, the modulation variance is adjusted according to

$$V_A = \left(2^{2R_{\mathrm{c}}/\beta} - 1\right)(1 + \chi_{\mathrm{tot}}), \tag{43}$$

where χ_{tot} is the total channel and detection noise referred to the channel input. This choice maintains the SNR required by (40). The FER values are taken from Figure 4b. EM–LDPC uses $N_{\mathrm{p}} = 0$, whereas the CCPE benchmark uses $N_{\mathrm{p}} = 1000$. Following the conventional finite-size model in [35], the parameter-estimation uncertainty is incorporated through the worst-case estimates of the channel transmittance and noise variance in $\chi^{\epsilon}(B;E)$, together with the finite-size correction $\Delta(L)$. At matched distance and reconciliation efficiency, these security terms are identical for EM–LDPC and CCPE; only the measured FER and additional disclosed phase-estimation samples differ. The sample-accounting factor in Equation (42) normalizes by the total transmitted samples. Thus, under the adopted finite-

size security model, the reported gain reflects the improvement in average secret key yield resulting from the proposed reconciliation method.

The finite-size SKR represents the net secure key yield per transmitted symbol after accounting for reconciliation efficiency, decoding failures, Eve's accessible information, and finite-size effects. Figure 6 shows the finite-size SKR for $\beta \in \{0.87, 0.90, 0.92, 0.94, 0.95, 0.96\}$.

The dependence of the SKR on β reflects a tradeoff between reconciliation efficiency and decoding reliability. Increasing β enlarges the positive reconciliation contribution $\beta I(A;B)$ in (42), allowing a larger fraction of the Alice–Bob correlation to contribute to secret key generation and thereby extending the secure transmission distance. However, for the fixed LDPC code rate used here, a larger β places the operating point closer to the effective-channel capacity, reduces the decoding margin, and can increase the FER. Since the SKR is multiplied by the successful-reconciliation factor $(1 - \mathrm{FER})$, the resulting loss in decoding reliability can offset part of the gain provided by the larger reconciliation efficiency.

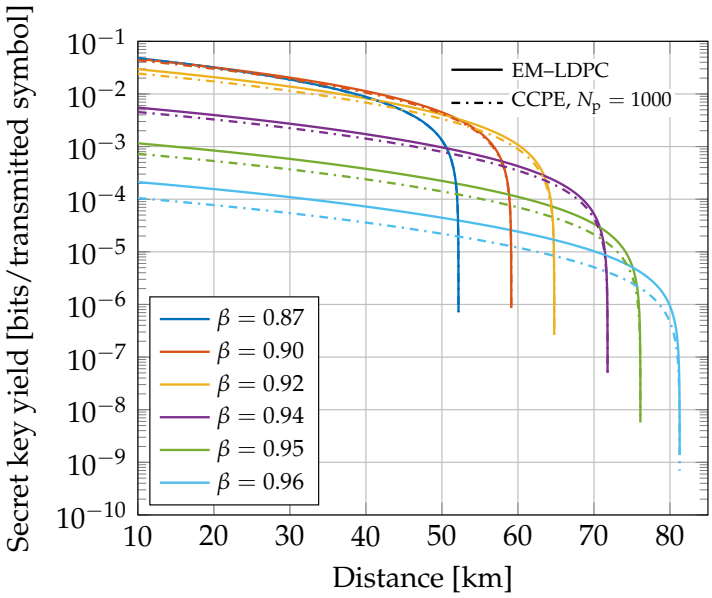


Figure 6. Finite-size secret key rate versus transmission distance for EM–LDPC and CCPE, with $N_p = 1000$ at different reconciliation efficiencies.

This tradeoff is visible in Figure 6. Increasing β extends the maximum transmission distance from approximately 52.2 km at $\beta = 0.87$ to 81.3 km at $\beta = 0.96$, whereas a larger β may give a lower SKR at shorter distances because of the increased FER. Therefore, the appropriate choice of β depends on both the transmission distance and the corresponding decoding reliability. For a fixed β, EM–LDPC provides a higher SKR than CCPE because it achieves a lower FER and avoids the additional phase-estimation disclosure overhead associated with $N_p = 1000$.

The above results provide a numerical evaluation of the proposed method under the considered CV-QKD channel and receiver model. The present study remains simulation-based, while full experimental verification would require an optical CV-QKD platform incorporating coherent-state transmission, local-oscillator control, homodyne detection, residual-phase evolution, and the corresponding classical reconciliation processing.

5. Conclusions

We developed a phase-aware multidimensional reconciliation method for homodyne-based CV-QKD with residual signal–LO phase mismatch. The receiver model separates Bob's known pulse-dependent quadrature-selection phase from the unknown residual

phase. The residual phase is included in the multidimensional virtual observations and their decoding likelihoods. The LDPC decoder calculates posterior spherical-symbol moments. These moments are used in a locally linearized EM-based phase update. After each update, the refined phase is used to reconstruct the virtual observations and LLRs. This feedback forms a closed loop between LDPC decoding and phase refinement. It does not require additional publicly disclosed samples for phase estimation.

Under the considered simulation conditions, EM–LDPC achieved a lower residual-phase-estimation RMSE than the compared CCPE receivers throughout the tested SNR range. Its BER and FER remained close to the perfect-compensation reference. It also outperformed the practical CCPE benchmarks under the same reconciliation and decoding conditions. In the finite-size evaluation, EM–LDPC achieved a higher SKR than the compared CCPE benchmark. This gain resulted from the lower FER and the absence of additional phase-estimation sample overhead. The reported improvement is obtained within the adopted finite-size security model and represents an increase in reconciliation-level secret key yield.

The main simulations assume that the residual phase is constant within each reconciliation block. We also performed a controlled test with linear intra-block phase drift. Small drifts caused only a modest performance loss under the tested conditions. Larger drifts caused progressive BER and FER degradation. This test demonstrates the robustness of the unchanged receiver to limited model mismatch and quantifies the progressive performance change as the drift increases.

The reported numerical gains apply to the considered Gaussian channel, detection noise model, parameter range, and reconciliation settings. Future work will extend the evaluation to more general stochastic phase dynamics and practical receiver impairments and will validate the proposed method on an optical CV-QKD platform.

Author Contributions: Conceptualization, J.D. and X.-Q.J.; methodology, Y.L. and J.D.; software, Y.L.; validation, Y.L. and H.H.; formal analysis, Y.L.; investigation, Y.L., H.H. and X.-Q.J.; writing—original draft preparation, Y.L.; writing—review and editing, J.D., H.H. and X.-Q.J.; visualization, Y.L.; supervision, J.D.; project administration, X.-Q.J. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Natural Science Foundation of Shanghai (Grants No. 25ZR1401002 and No. 24ZR1401000), the National Science and Technology Major Project in Quantum Science and Technology (Grant No. 2021ZD0300703), and the Shanghai Municipal Science and Technology Major Project (Grant No. 2019SHZDZX01).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors acknowledge the use of OpenAI Codex (version 0.153.0-alpha.5) for assistance with language editing. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

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

References

1. Scarani, V.; Bechmann-Pasquinucci, H.; Cerf, N.J.; Dušek, M.; Lütkenhaus, N.; Peev, M. The Security of Practical Quantum Key Distribution. *Rev. Mod. Phys.* **2009**, *81*, 1301–1350. [CrossRef]
2. Weedbrook, C.; Pirandola, S.; Garcia-Patron, R.; Cerf, N.J.; Ralph, T.C.; Shapiro, J.H.; Lloyd, S. Gaussian quantum information. *Rev. Mod. Phys.* **2012**, *84*, 621–669. [CrossRef]

3. Grosshans, F.; Van Assche, G.; Wenger, J.; Brouri, R.; Cerf, N.J.; Grangier, P. Quantum key distribution using gaussian-modulated coherent states. *Nature* **2003**, *421*, 238–241. [CrossRef] [PubMed]
4. Laudenbach, F.; Pacher, C.; Fung, C.H.F.; Poppe, A.; Peev, M.; Schrenk, B.; Hentschel, M.; Walther, P.; Hübel, H. Continuous-Variable Quantum Key Distribution with Gaussian Modulation—The Theory of Practical Implementations. *Adv. Quantum Technol.* **2018**, *1*, 1800011. [CrossRef]
5. Zhang, Y.; Bian, Y.; Li, Z.; Yu, S.; Guo, H. Continuous-Variable Quantum Key Distribution System: Past, Present, and Future. *Appl. Phys. Rev.* **2024**, *11*, 011318. [CrossRef]
6. Jouguet, P.; Kunz-Jacques, S.; Leverrier, A. Long-Distance Continuous-Variable Quantum Key Distribution with a Gaussian Modulation. *Phys. Rev. A* **2011**, *84*, 062317. [CrossRef]
7. Fu, H.; Dai, J.; Feng, Y.; Hai, H.; Ge, H.; Huang, P.; Jiang, X.Q. Rate-Adaptive Information Reconciliation for CV-QKD Systems at Low Signal-to-Noise Ratios. *Entropy* **2026**, *28*, 10. [CrossRef] [PubMed]
8. Yang, H.; Liu, S.; Yang, S.; Lu, Z.; Li, Y.; Li, Y. High-Efficiency Rate-Adaptive Reconciliation in Continuous-Variable Quantum Key Distribution. *Phys. Rev. A* **2024**, *109*, 012604. [CrossRef]
9. Jia, S.; Hai, H.; Jiang, X.Q.; Mumtaz, S.; Markakis, E.K.; Huang, P. High-Efficiency Reconciliation for Continuous-Variable Quantum Key Distribution Using MGC-LDPC Codes. *Quantum Inf. Process.* **2025**, *24*, 7. [CrossRef]
10. Chen, Z.; Wang, X.; Yu, S.; Li, Z.; Guo, H. Continuous-Mode Quantum Key Distribution with Digital Signal Processing. *npj Quantum Inf.* **2023**, *9*, 28. [CrossRef]
11. Gehring, T.; Händchen, V.; Duhme, J.; Furrer, F.; Franz, T.; Pacher, C.; Werner, R.F.; Schnabel, R. Implementation of Continuous-Variable Quantum Key Distribution with Composable and One-Sided-Device-Independent Security against Coherent Attacks. *Nat. Commun.* **2015**, *6*, 8795. [CrossRef] [PubMed]
12. Qi, B.; Lougovski, P.; Pooser, R.; Grice, W.; Bobrek, M. Generating the Local Oscillator "Locally" in Continuous-Variable Quantum Key Distribution Based on Coherent Detection. *Phys. Rev. X* **2015**, *5*, 041009. [CrossRef]
13. Soh, D.B.S.; Brif, C.; Coles, P.J.; Lütkenhaus, N.; Camacho, R.M.; Urayama, J.; Sarovar, M. Self-Referenced Continuous-Variable Quantum Key Distribution Protocol. *Phys. Rev. X* **2015**, *5*, 041010. [CrossRef]
14. Laudenbach, F.; Schrenk, B.; Pacher, C.; Hentschel, M.; Fung, C.H.F.; Karinou, F.; Poppe, A.; Peev, M.; Hübel, H. Pilot-Assisted Intradyne Reception for High-Speed Continuous-Variable Quantum Key Distribution with True Local Oscillator. *Quantum* **2019**, *3*, 193. [CrossRef]
15. Zhao, W.; Guo, Y.; Zhang, L.; Huang, D. Phase Noise Estimation Using Bayesian Inference for Continuous-Variable Quantum Key Distribution. *Opt. Express* **2019**, *27*, 1838–1853. [CrossRef] [PubMed]
16. Kleis, S.; Schaeffer, C.G. Improving the Secret Key Rate of Coherent Quantum Key Distribution with Bayesian Inference. *J. Light. Technol.* **2019**, *37*, 722–728. [CrossRef]
17. Su, Y.; Guo, Y.; Huang, D. Kalman Filter-Based Phase Estimation of Continuous-Variable Quantum Key Distribution without Sending Local Oscillator. *Phys. Lett. A* **2019**, *383*, 2394–2399. [CrossRef]
18. Huang, B.; Huang, Y.; Peng, Z. Tracking Reference Phase with a Kalman Filter in Continuous-Variable Quantum Key Distribution. *Opt. Express* **2020**, *28*, 28727–28739. [CrossRef] [PubMed]
19. Shen, T.; Qi, D.; Wang, X.; Zheng, D. Enhancing Noise Suppression of Phase Estimation in Continuous-Variable Quantum Key Distribution with Discrete Modulation. *Opt. Express* **2025**, *33*, 215–227. [CrossRef] [PubMed]
20. Liao, J.; Fan, W.; Huang, J.; Xiao, H.; Wang, D. Simple Phase Recovery Using Multi-Pilot Tones and Wiener Filters for Continuous Variable Quantum Key Distribution. *Opt. Express* **2025**, *33*, 35999–36014. [CrossRef] [PubMed]
21. Chin, H.M.; Jain, N.; Zibar, D.; Andersen, U.L.; Gehring, T. Machine Learning Aided Carrier Recovery in Continuous-Variable Quantum Key Distribution. *npj Quantum Inf.* **2021**, *7*, 20. [CrossRef]
22. Long, N.K.; Malaney, R.; Grant, K.J. A Survey of Machine Learning Assisted Continuous-Variable Quantum Key Distribution. *Information* **2023**, *14*, 553. [CrossRef]
23. Huang, D.; Huang, P.; Lin, D.; Zeng, G. Long-Distance Continuous-Variable Quantum Key Distribution by Controlling Excess Noise. *Sci. Rep.* **2016**, *6*, 19201. [CrossRef] [PubMed]
24. Wang, S.; Huang, P.; Liu, M.; Wang, T.; Wang, P.; Zeng, G. Phase Compensation for Free-Space Continuous-Variable Quantum Key Distribution. *Opt. Express* **2020**, *28*, 10737–10745. [CrossRef] [PubMed]
25. Dai, J.; Jiang, X.Q.; Huang, P.; Wang, T.; Zeng, G. Joint Parameter Estimation and Multidimensional Reconciliation for Continuous-Variable Quantum Key Distribution. *Phys. Rev. A* **2026**, *114*, 022429. [CrossRef]
26. Wang, X.Y.; Guo, X.B.; Jia, Y.X.; Zhang, Y.; Lu, Z.G.; Liu, J.Q.; Li, Y.M. Accurate Shot-Noise-Limited Calibration of a Time-Domain Balanced Homodyne Detector for Continuous-Variable Quantum Key Distribution. *J. Light. Technol.* **2023**, *41*, 5518–5528. [CrossRef]
27. Beguin, J.B.; Bookjans, E.M.; Christensen, S.L.; Sorensen, H.L.; Muller, J.H.; Appel, J.; Polzik, E.S. Generation and Detection of a Sub-Poissonian Atom Number Distribution in a One-Dimensional Optical Lattice. *Phys. Rev. Lett.* **2014**, *113*, 263603. [CrossRef] [PubMed]

28. Roudas, I. Coherent Optical Communication Systems. In *WDM Systems and Networks*; Optical Networks; Springer: Berlin/Heidelberg, Germany, 2012; pp. 373–417. [CrossRef]
29. Leverrier, A.; Alléaume, R.; Boutros, J.; Zémor, G.; Grangier, P. Multidimensional Reconciliation for a Continuous-Variable Quantum Key Distribution. *Phys. Rev. A* **2008**, *77*, 042325. [CrossRef]
30. Dai, J.; Jiang, X.Q.; Wang, T.; Huang, P.; Zeng, G. Arbitrarily-High-Dimensional Reconciliation via Cross-Rotation for Continuous-Variable Quantum Key Distribution. *Phys. Rev. A* **2025**, *112*, 022610. [CrossRef]
31. Kschischang, F.R.; Frey, B.J.; Loeliger, H.A. Factor Graphs and the Sum-Product Algorithm. *IEEE Trans. Inf. Theory* **2001**, *47*, 498–519. [CrossRef]
32. Loeliger, H.A. An Introduction to Factor Graphs. *IEEE Signal Process. Mag.* **2004**, *21*, 28–41. [CrossRef]
33. Dai, J.; So, H.C. Real-Valued Sparse Bayesian Learning for DOA Estimation With Arbitrary Linear Arrays. *IEEE Trans. Signal Process.* **2021**, *69*, 4977–4990. [CrossRef]
34. Milicevic, M.; Feng, C.; Zhang, L.M.; Gulak, P.G. Quasi-Cyclic Multi-Edge LDPC Codes for Long-Distance Quantum Cryptography. *npj Quantum Inf.* **2018**, *4*, 21. [CrossRef]
35. Leverrier, A.; Grosshans, F.; Grangier, P. Finite-Size Analysis of a Continuous-Variable Quantum Key Distribution. *Phys. Rev. A* **2010**, *81*, 062343. [CrossRef]

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.