

Mitigating Network Latency in QKD via Asynchronous Syndrome-Projected Cascade

Sanil A. Raut

rautsa@rbunagpur.in

Ramdeobaba University

Research Article

Keywords: Quantum Key Distribution, Quantum Bit Error Rate , Information reconciliation, Cascade, Low-Density Parity-Check, Optical transmission

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11047435/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Mitigating Network Latency in QKD via Asynchronous Syndrome-Projected Cascade

Sanil A. Raut^{1*}

^{1*}Department of CSE, Ramdeobaba University, Katol Road, Nagpur, 440013, Maharashtra, India.

Corresponding author(s). E-mail(s): rautsa@rbunagpur.in;

Abstract

Information reconciliation in Quantum Key Distribution (QKD) is limited by a dual-latency paradox: traditional Cascade protocols suffer severe network propagation delays, while Low-Density Parity-Check (LDPC) codes incur prohibitive computational overhead. We propose the Asynchronous Syndrome-Projected Cascade (ASP-Cascade) to overcome this limitation. ASP-Cascade replaces interactive parity queries with deterministic Galois Field $\mathbf{GF}(2)$ syndrome projections, achieving $\mathcal{O}(1)$ algebraic error localization. By introducing an asynchronous trace-back mechanism, cascading errors are resolved entirely within local CPU memory, completely bypassing iterative network requests.

We rigorously evaluate ASP-Cascade over a realistic 1550 nm optical channel, explicitly modeling SMF-28 fiber attenuation and Single-Photon Avalanche Diode (SPAD) dark counts. Results demonstrate that, despite a higher classical bandwidth requirement, ASP-Cascade reduces communication complexity to a single one-shot transmission burst. This completely decouples reconciliation from network latency, maintaining a 0% residual Quantum Bit Error Rate (QBER) even against the exponential dark count noise barrier. Furthermore, CPU micro-benchmarking confirms ASP-Cascade reduces local execution time by two orders of magnitude compared to LDPC belief propagation, yielding a significantly higher Devetak-Winter Secret Key Rate (SKR) over long-distance optical networks.

Keywords: Quantum Key Distribution, Quantum Bit Error Rate, Information reconciliation, Cascade, Low-Density Parity-Check, Optical transmission

1 Introduction

Quantum Key Distribution (QKD) provides theoretically unconditional security for cryptographic key exchange by exploiting the fundamental principles of quantum mechanics. However, practical deployments of discrete-variable QKD, such as the widely implemented decoy-state BB84 protocol, inevitably suffer from quantum channel noise, photon loss, and detector imperfections. Consequently, the raw key strings generated by the sender (Alice) and receiver (Bob) contain a non-negligible Quantum Bit Error Rate (QBER). To distill a perfectly identical and secure cryptographic key, a classical post-processing phase known as information reconciliation is strictly required.

Since its foundational origins in the BB84 [1] and E91 [2] protocols, QKD has evolved dramatically. Rigorous security proofs [3] have paved the way for highly advanced modalities. Today, the field extends beyond standard fiber optics to encompass Measurement-Device-Independent (MDI) architectures [4], Twin-Field QKD (TF-QKD) capable of overcoming the linear rate-transmittance bound [5, 6], and Continuous-Variable systems (CV-QKD) [7]. Furthermore, QKD has successfully broken the constraints of terrestrial fiber through free-space satellite networks over thousands of kilometers [8, 9]. Yet, as these physical limits are pushed and the fundamental PLOB capacity bounds are approached [10], the bottleneck of QKD has notably shifted from optical physics to classical post-processing [11].

The scalability of information reconciliation is currently bounded by a dual-latency paradox: the fundamental trade-off between communication round complexity and local computational overhead. The Cascade protocol, originally introduced by Brassard and Salvail [12, 13], remains a ubiquitous standard due to its deterministic convergence, virtually guaranteeing a 0% residual error rate. However, Cascade relies heavily on interactive binary searches. Over long-distance fiber-optic networks, the physical speed-of-light propagation delay forces local processing units to idle during multiple network round-trips, creating a severe network bottleneck that chokes overall system throughput. Conversely, forward error correction architectures such as Low-Density Parity-Check (LDPC) codes [14, 15] attempt to circumvent this network latency via one-shot matrix transmissions. While highly efficient over the network, LDPC introduces severe computational bottlenecks through complex floating-point Belief Propagation iterations and often suffers from probabilistic error floors, failing to correct the final few bit errors.

While Cascade and LDPC [16, 17] remain the industry standards, researchers have continually sought to optimize the reconciliation phase. The Winnow protocol [18] was introduced to reduce the interactivity of Cascade via Hamming codes. More recently, Polar codes [19] have gained significant traction in QKD environments due to their deterministic capacity to approach the Shannon limit [20]. However, all of these traditional and contemporary solutions ultimately face a trade-off between algorithmic latency and computational intensity.

Recent literature has attempted to bridge this performance gap. For instance, Hu et al. [21] proposed a Modified Cascade protocol that mathematically optimizes initial block sizes and introduces trace-back flag bits. While this modification successfully minimizes wasted parity disclosures (information leakage), it fails to escape the underlying interactive architecture that dictates its network latency. Similarly, while

various LDPC matrix adaptations have been proposed to lower computational floors in continuous-variable systems, they struggle to scale efficiently when confronted with the exponential noise spikes characteristic of realistic optical hardware [22–24]. Currently, no classical post-processing algorithm efficiently minimizes both network-induced propagation delays and CPU-bound computational bottlenecks simultaneously under practical telecom constraints.

In response to this dual-latency paradox, high-speed multi-megabit QKD implementations have increasingly relied on massive parallelization [25, 26]. Modern scalable networks frequently attempt to bypass software limitations by deploying dedicated hardware acceleration, such as Field Programmable Gate Arrays (FPGAs), to handle the intense computational load of LDPC belief propagation [27, 28]. Despite these brute-force hardware optimizations, a protocol designed specifically to execute locally in $\mathcal{O}(1)$ time without relying on iterative floating-point arrays is highly desirable for minimizing deployment costs and maximizing throughput.

To overcome this limitation, this paper proposes a novel hybrid software architecture: the **Asynchronous Syndrome-Projected Cascade (ASP-Cascade)**. By fusing the definitive multi-round convergence of Cascade with the one-shot transmission efficiency of forward error correction, ASP-Cascade decouples reconciliation throughput from network latency. The proposed protocol abandons interactive parity queries in favor of deterministic Galois Field $GF(2)$ syndrome projections, achieving $\mathcal{O}(1)$ algebraic error localization. Furthermore, it introduces a zero-latency asynchronous trace-back mechanism, resolving multi-round cascading errors entirely within local CPU memory without initiating iterative optical transmissions.

To rigorously validate this algorithmic breakthrough, we employ a dual-model evaluation framework. First, the algorithm is benchmarked across theoretical 3D parameter topologies using an idealized quantum noise model. Second, the protocol is subjected to a realistic 1550 nm telecommunications stress test, explicitly incorporating standard single-mode fiber (SMF-28) attenuation and Single-Photon Avalanche Diode (SPAD) dark count mechanics. Ultimately, our benchmarking confirms that by eliminating interactive network round-trips, ASP-Cascade achieves effective throughput rates orders of magnitude higher than Base and Modified Cascade, while completely bypassing the CPU floating-point bottlenecks that throttle LDPC codes under high-noise conditions.

The primary contributions of this paper are summarized as follows:

- **Novel Protocol Architecture (ASP-Cascade):** We propose the Asynchronous Syndrome-Projected Cascade protocol, which replaces multi-round network parity queries with one-shot $\mathcal{O}(1)$ Galois Field $GF(2)$ matrix projections, fundamentally eliminating the network round-trip time (RTT) bottlenecks of traditional interactive reconciliation.
- **Hardware-Optimized Local Execution:** By substituting the heavy, iterative floating-point arithmetic of Min-Sum Belief Propagation (LDPC) with native bitwise XOR operations, we theoretically and empirically demonstrate a two-order-of-magnitude reduction in localized CPU execution time.
- **Realistic Physical Modeling:** We develop a robust Monte Carlo simulation engine modeling a 1550 nm Decoy-State BB84 channel, incorporating exact optical

attenuation and Single-Photon Avalanche Diode (SPAD) dark count probabilities to evaluate algorithms against real-world optical noise limits.

- **Empirical Cryptographic Validation:** We stress-test the protocol against this channel, demonstrating that ASP-Cascade delivers an optimal cryptographic yield—measured via the Devetak-Winter Secret Key Rate (SKR) bound—over long-distance optical fiber, successfully bypassing the cascading waterfall failures characteristic of fixed-rate codes.

The remainder of this manuscript is organized as follows. Section 2 defines the physical system models, establishing the fundamental noise limitations of the optical channel. Section 3 details the mathematical formulation and architecture of the ASP-Cascade protocol. Section 4 presents our performance benchmarking, comparing ASP-Cascade against Base Cascade, Modified Cascade, and LDPC across both information leakage and effective throughput metrics, CPU micro-benchmarking, and final Secret Key Rate (SKR) yields. Finally, Section 5 provides concluding remarks.

2 System Models and Physical Channel Limitations

To comprehensively evaluate the reconciliation algorithms, we model the quantum communication environment under two distinct paradigms: an idealized theoretical channel for baseline algorithmic validation, and a highly realistic optical channel that enforces strict physical hardware limitations.

2.1 Idealized Quantum Noise Model

In the theoretical baseline model, quantum channel noise is treated purely as a mathematical abstraction. Errors are introduced via simulated amplitude damping (photon loss) and linear depolarizing drift (phase errors), with the detection probability assumed to be perfect. This idealized environment allows for the isolation of the algorithmic logic, demonstrating how the protocols perform when scaling against a strictly fixed number of errors, n , independent of propagation distance.

2.2 Realistic 1550 nm Optical Channel

Practical QKD deployments, such as the Decoy-State BB84 protocol operating at the 1550 nm telecommunications wavelength, are constrained by the physical properties of optical fiber and semiconductor detectors [29].

First, the transmittance of the quantum channel, t_{channel} , decays exponentially with distance L (in kilometers) due to the inherent attenuation of standard single-mode optical fiber (SMF-28):

$$t_{\text{channel}} = 10^{-\frac{\alpha L}{10}} \quad (1)$$

where α is the fiber attenuation coefficient, typically 0.2 dB/km.

At the receiver’s end, the quantum signal is measured using Single-Photon Avalanche Diodes (SPADs). The total signal yield Y is a combination of the true transmitted photons and the thermal noise inherent to the detector, known as the

dark count probability (P_{dark}):

$$Y = \eta_{\text{det}} \cdot t_{\text{channel}} + P_{\text{dark}} \quad (2)$$

where η_{det} represents the detection efficiency of the SPAD.

Consequently, the realistic Quantum Bit Error Rate (QBER) as a function of distance is not linear, but rather governed by the interplay between optical misalignment and detector noise:

$$\text{QBER}(L) = \frac{e_{\text{opt}} \cdot \eta_{\text{det}} \cdot t_{\text{channel}} + \frac{1}{2}P_{\text{dark}}}{Y} \quad (3)$$

Here, e_{opt} represents the baseline optical misalignment error. The multiplier $\frac{1}{2}$ accounts for the fact that a random dark count has a 50% probability of correctly guessing the intended bit. As distance L increases, t_{channel} approaches zero, causing the fixed P_{dark} variable to dominate the denominator. This physical phenomenon triggers an exponential noise spike known as the “dark count wall,” which acts as the ultimate stress test for post-processing algorithms.

2.3 Effective System Throughput Formulation

To accurately model the dual-latency paradox, the performance of a reconciliation protocol must account for both local CPU execution time and optical network propagation delays. The effective throughput, T_{eff} (in bits per second), for a key of length N is defined as:

$$T_{\text{eff}} = \frac{N}{t_{\text{cpu}} + (I \times \text{RTT})} \quad (4)$$

where t_{cpu} is the total local computational time, I is the total number of network interactions required by the protocol, and RTT is the Round-Trip Time of the optical network. RTT is physically bounded by $\text{RTT} = \frac{2L}{c_{\text{fiber}}}$, with $c_{\text{fiber}} \approx 2 \times 10^5$ km/s representing the speed of light within a silica glass core. In our empirical benchmarking, t_{cpu} is captured dynamically via high-resolution hardware timers to accurately reflect the computational complexity of each protocol—ranging from the heavy floating-point iterations of LDPC Belief Propagation to the highly efficient native bitwise XOR operations utilized in ASP-Cascade’s $GF(2)$ projections.

3 Proposed Methodology: ASP-Cascade

The fundamental limitation of traditional interactive information reconciliation is the reliance on iterative binary searches, which strictly bind algorithmic progression to optical network latency. To shatter this physical bottleneck, we propose the Asynchronous Syndrome-Projected Cascade (ASP-Cascade) protocol. ASP-Cascade is a hybrid architecture designed to inherit the 0% residual error guarantee of multi-round Cascade while leveraging the one-shot transmission efficiency of forward error correction, completely decoupling error resolution from network round-trip times.

3.1 Block Partitioning and Initialization

Let Alice and Bob share a raw key of length N with an estimated quantum bit error rate p . Following the multi-round structure of standard Cascade, the key undergoes a random permutation at the start of each round to disperse burst errors. To optimize the probability of isolating single errors within a block, ASP-Cascade inherits the optimized partitioning parameters established by Hu et al. [21, 30]. The block size for the first round, k_1 , and the second round, k_2 , are defined as:

$$k_1 = \max \left(2, \left\lfloor \frac{0.58}{p} \right\rfloor \right), \quad k_2 = \max (2, \lfloor 2.5 \cdot k_1 \rfloor) \quad (5)$$

Subsequent rounds utilize block sizes of $N/2$ to rapidly expose hidden errors.

3.2 One-Shot Syndrome Projection Burst

In standard Cascade, Alice and Bob exchange a single scalar parity bit per block, revealing only the presence of an odd number of errors. This necessitates a subsequent $\mathcal{O}(\log_2 k)$ binary search across the network [12]. ASP-Cascade replaces this process by projecting the block into a shortened Hamming space over the Galois Field $GF(2)$.

For a block of size k , Alice and Bob deterministically generate a parity-check matrix $\mathbf{H} \in \{0, 1\}^{m \times k}$. The minimum syndrome length m is dictated by the Hamming bound for single-error correction:

$$2^m - 1 \geq k \implies m = \lceil \log_2(k + 1) \rceil \quad (6)$$

The columns of $\mathbf{H}$ are populated with the binary representations of integers from 1 to k . Let Alice's block be represented as a column vector $\mathbf{x}_A \in \{0, 1\}^k$. Instead of computing a single parity bit, Alice computes an m -bit syndrome vector:

$$\mathbf{S}_A = \mathbf{H}\mathbf{x}_A \pmod{2} \quad (7)$$

To achieve total network decoupling, ASP-Cascade operates as a ‘‘burst’’ protocol. Alice computes $\mathbf{S}_A$ for *all blocks across all scheduled rounds* simultaneously and transmits this global state to Bob in a single, one-shot optical transmission, thereby reducing the communication round complexity to $\mathcal{O}(1)$.

3.3 Algebraic $\mathcal{O}(1)$ Error Localization

Upon receiving Alice's global syndrome state, Bob computes the local syndrome for his corresponding block, $\mathbf{S}_B = \mathbf{H}\mathbf{x}_B \pmod{2}$, and evaluates the syndrome differential:

$$\mathbf{S}_\Delta = \mathbf{S}_A \oplus \mathbf{S}_B \quad (8)$$

where $\oplus$ denotes bitwise exclusive-OR (XOR) addition. Because quantum channel errors manifest as linear bit-flips, if the block contains exactly one error, the vector $\mathbf{S}_\Delta$

will be mathematically identical to exactly one column vector $\mathbf{h}_e$ in $\mathbf{H}$. The integer representation of $\mathbf{S}_\Delta$ yields the exact physical index e of the erroneous bit.

Because this matrix projection utilizes native $GF(2)$ XOR logic rather than floating-point arithmetic, Bob’s local CPU executes the localization in $\mathcal{O}(1)$ time, completely bypassing both the network RTT and the Belief Propagation bottlenecks associated with LDPC.

3.4 Zero-Latency Asynchronous Trace-Back

The most severe latency penalty in standard Cascade occurs during the trace-back phase. When Bob corrects an error in round $i \geq 2$, it alters the parity of a block in a previous round $j < i$, exposing a hidden error. Traditional algorithms must halt local processing and initiate a new interactive network query to resolve this exposed error.

ASP-Cascade introduces an asynchronous trace-back mechanism that severs this dependency. Because Bob already possesses Alice’s global syndrome state in his local memory, he can perform trace-backs without communicating with Alice. When Bob flips an erroneous bit at global index e during round i , he determines which block that bit belonged to in previous round j . Bob then updates his *local* round j syndrome mathematically:

$$\mathbf{S}_{B,j}^{(\text{new})} = \mathbf{S}_{B,j}^{(\text{old})} \oplus \mathbf{h}_{e,j} \quad (9)$$

where $\mathbf{h}_{e,j}$ is the specific column vector in the round j parity-check matrix corresponding to the flipped bit.

Bob immediately recalculates the differential $\mathbf{S}_\Delta = \mathbf{S}_{A,j} \oplus \mathbf{S}_{B,j}^{(\text{new})}$. If $\mathbf{S}_\Delta \neq 0$, the hidden error is instantly localized and corrected. This recursive chain reaction cascades through all prior rounds instantaneously within Bob’s L1/L2 CPU cache. By substituting optical network queries with localized $GF(2)$ vector updates, ASP-Cascade achieves a zero-latency trace-back, securing a 0% QBER while operating at hardware-limited throughput speeds.

4 Performance Benchmarking and Results

To evaluate the efficacy of the ASP-Cascade protocol, we conducted extensive Monte Carlo simulations benchmarking it against Base Cascade, Modified Cascade, and a Min-Sum Belief Propagation LDPC code (rate 0.6). The evaluation is split into two domains: isolated algorithmic scaling over theoretical 3D topologies, and real-world performance against physical telecom hardware constraints.

4.1 Algorithmic Scaling and 3D Topologies

To decouple computational latency from optical network delays, the algorithms were first evaluated across a 3D parameter space, mapping Information Leakage and Effective Throughput simultaneously against Fiber Distance (L) and Error Count (n).

As observed in the topological surfaces (Fig. 1), information leakage is strictly an algorithmic byproduct of the error count (Y-axis), with distance (X-axis) exerting zero mathematical influence. LDPC maintains a rigid leakage plane, while both traditional Cascade protocols scale logarithmically. ASP-Cascade exhibits the steepest leakage




Fig. 1 3D topological scaling of information leakage (left) and effective throughput (right) across fiber distance and exact error counts.

topography, a direct mathematical consequence of broadcasting full m -bit syndromes for every block.

However, the effective throughput topology (plotted on a $\log_{10}$ scale) reveals the critical advantage of this bandwidth trade-off. The surfaces for Base and Modified Cascade collapse into a severe latency valley as both distance (speed-of-light delay) and error count (interactive trace-back queries) increase. LDPC throughput remains low but stable, bottlenecked exclusively by the CPU-intensive floating-point operations of Belief Propagation [22]. By avoiding both network round-trips and floating-point math, the ASP-Cascade throughput surface dominates the high-altitude canopy, scaling purely at the hardware limit of the local processor.

4.2 Real-World Stress Test: The Dark Count Wall

To transition from theoretical mathematics to practical engineering, the protocols were subjected to the 1550 nm Decoy-State BB84 channel model detailed in Section 2. We simulated fiber distances from 10 km to 130 km, integrating an attenuation of 0.2 dB/km, a baseline optical misalignment of 1.5%, and a SPAD efficiency of 10% with a 5×10^{-5} dark count probability.

As depicted in Fig. 2, the raw quantum bit error rate exhibits a distinct bipartite behavior. Between 10 km and 60 km, the signal-to-noise ratio is strong, and errors are dominated strictly by the flat 1.5% optical misalignment baseline. However, as distance surpasses 70 km, fiber attenuation exponentially depletes the true photon signal. At this critical inflection point, the constant thermal noise of the SPAD detectors begins to dominate the denominator of the error equation, triggering the “dark count wall.”

Under this extreme exponential noise gradient, the fixed-rate LDPC protocol experiences a catastrophic waterfall failure, rendering it incapable of resolving the errors.




Fig. 2 The Dark Count Wall: Residual error rates of post-processing protocols plotted against realistic 1550 nm optical fiber degradation.

In stark contrast, ASP-Cascade successfully navigates the exponential spike. By leveraging the deterministic structure of the multi-round Cascade trace-back, ASP-Cascade matches traditional interactive algorithms, successfully maintaining a 0% residual error rate even at distances exceeding 100 km.

4.3 Classical Bandwidth Cost and Information Leakage

Information leakage represents the percentage of classical bits that must be disclosed over the public channel to correct the quantum key. In practical deployments, this metric dictates the requisite classical bandwidth overhead.

Figure 3 illustrates the bandwidth trade-offs inherent to each algorithmic architecture. The Base and Modified Cascade protocols exhibit highly optimized leakage profiles, as they dynamically request single parity bits only when an error is suspected. Consequently, their leakage curves perfectly mirror the exponential climb of the dark count wall. LDPC maintains a rigidly flat leakage profile dictated by its pre-defined matrix rate (e.g., $1 - 0.6 = 40\%$).

ASP-Cascade requires a distinctly higher percentage of classical bandwidth across all distances. This is a direct, deliberate consequence of its mathematical architecture: to achieve one-shot transmission, Alice must broadcast the full m -bit syndrome for every block simultaneously, regardless of whether a block contains an error. While this inflates the classical bandwidth requirement, modern Dense Wavelength-Division




Fig. 3 Information Leakage: The classical bandwidth cost required to achieve convergence across varying fiber distances.

Multiplexing (DWDM) networks possess vast reserves of classical bandwidth, making this a highly favorable and realistic engineering compromise.

4.4 Effective Throughput and the Latency Paradox

The true superiority of ASP-Cascade is isolated in the effective throughput analysis, which aggregates both the local CPU execution time and the network propagation delay.

As demonstrated in Fig. 4 (plotted on a logarithmic scale), traditional interactive algorithms are fundamentally crippled by the speed of light. Even though Base and Modified Cascade utilize minimal CPU resources, their reliance on iterative network queries means the local processor spends the vast majority of its time idling, waiting for parity bits to traverse the fiber. Consequently, their throughput crashes by orders of magnitude as distance scales. LDPC bypasses this network bottleneck but is severely throttled by its own CPU execution time, specifically the heavy floating-point operations required for continuous Belief Propagation iterations.

ASP-Cascade successfully navigates both limitations. By broadcasting all syndromes upfront, it reduces network interactions to $\mathcal{O}(1)$, entirely negating the speed-of-light penalty that destroys standard Cascade. Concurrently, by processing its trace-backs using strictly local $GF(2)$ bitwise XOR operations, it bypasses the floating-point CPU bottleneck of LDPC. As a result, ASP-Cascade tracks seamlessly




Fig. 4 Effective System Throughput: A logarithmic representation of the total protocol speeds, incorporating both CPU bottlenecks and optical network round-trip delays.

just beneath the hardware clock ceiling, securing a 0% QBER while operating at maximum physical speeds.

4.5 Ultimate Cryptographic Yield: Secret Key Rate

While throughput and information leakage act as critical intermediary metrics, the definitive benchmark of any QKD protocol is the final Secret Key Rate (SKR). The SKR dictates the actual volume of usable, unconditionally secure bits generated per second after privacy amplification. We evaluate the SKR utilizing the standard Devetak-Winter bound [31]:

$$\text{SKR} = T_{\text{eff}} \times \max(0, 1 - f_{\text{leak}} - H(e)) \quad (10)$$

where T_{eff} is the effective system throughput, f_{leak} is the fraction of information leaked during reconciliation, and $H(e)$ is the binary Shannon entropy of the QBER.

As illustrated in Fig. 5, standard Cascade protocols generate a moderate SKR at short distances but suffer a catastrophic logarithmic collapse as network RTT suffocates T_{eff} . LDPC generates a minimal SKR until it abruptly hits its waterfall failure near 70 km. In contrast, ASP-Cascade dominates the cryptographic yield, maintaining nearly 10^2 kbps securely. Despite sacrificing a larger leakage fraction than Modified Cascade, ASP-Cascade’s ability to maintain maximum throughput over the optical




Fig. 5 Secret Key Rate demonstrating final cryptographic yield over distance. ASP-Cascade dominates throughput until the dark count wall forces an algorithmic leakage limit at 110 km.

fiber vastly outweighs the classical bandwidth penalty. Notably, at 110 km, the violent spike in channel noise forces ASP-Cascade to leak nearly the entire key to maintain zero latency, correctly driving the SKR to zero and validating the physical limits of the dark count wall.

4.6 Hardware Validation: CPU Micro-Benchmarking

To empirically validate the theoretical $\mathcal{O}(1)$ complexity claims of the ASP-Cascade architecture, we conducted a localized CPU micro-benchmark. This test isolates the pure hardware execution time at a simulated 50 km transmission distance, strictly eliminating optical network RTT delays to observe local processing limits.

As depicted in Fig. 6 on a logarithmic scale, the Belief Propagation required for LDPC consumes over 3000 ms of continuous floating-point CPU cycles, empirically verifying it as an extreme hardware bottleneck. While Base and Modified Cascade reduce this localized execution time to roughly 15–25 ms, they still require hundreds of subsequent network interactions to complete reconciliation due to their interactive binary searches. ASP-Cascade completes its localized processing in approximately 22.5 ms. By replacing iterative floating-point arithmetic with native Galois Field $GF(2)$ bitwise XOR operations, ASP-Cascade matches the highly efficient CPU footprint of classical Cascade while requiring only a single, one-shot network transmission, proving its zero-latency algorithmic design.




Fig. 6 Log-scale CPU micro-benchmark isolating pure local hardware execution time at a 50 km transmission distance.

5 Conclusion

The practical deployment of high-speed Quantum Key Distribution (QKD) has historically been bounded by a dual-latency paradox: interactive algorithms like Cascade are throttled by optical network round-trip times (RTT), while forward error correction codes like LDPC are bottlenecked by CPU-intensive floating-point operations. In this paper, we introduced the Asynchronous Syndrome-Projected Cascade (ASP-Cascade) protocol to fundamentally sever these dependencies.

By substituting multi-round network parity queries with one-shot $\mathcal{O}(1)$ Galois Field $GF(2)$ matrix projections, ASP-Cascade executes recursive error trace-backs instantaneously within the local hardware cache. Extensive benchmarking against a realistic 1550 nm telecommunications model confirms that this asynchronous architecture successfully absorbs the exponential noise of the physical “dark count wall” without suffering the cascading waterfall failures characteristic of fixed-rate LDPC codes. Furthermore, empirical CPU micro-benchmarking demonstrates that ASP-Cascade reduces local hardware execution time by over two orders of magnitude compared to Min-Sum Belief Propagation.

While this zero-latency design necessitates a fractional increase in classical bandwidth consumption, the trade-off fundamentally eliminates the network RTT bottleneck. Ultimately, ASP-Cascade delivers a secure cryptographic yield—measured via the Devetak-Winter Secret Key Rate (SKR) bound—that is orders of magnitude higher than traditional protocols over long-distance optical fiber. By operating efficiently at the hardware clock ceiling, ASP-Cascade provides a robust, highly scalable post-processing architecture tailored for next-generation, high-throughput quantum cryptographic networks.

Declarations

Funding: No funding was received for conducting this study.

Competing Interests: The authors have no competing interests to declare that are relevant to the content of this article.

Data Availability: The datasets generated during the current study are available from the corresponding author on reasonable request.

References

- [1] Bennett, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. In: Proceedings of IEEE International Conference on Computers, Systems and Signal Processing, pp. 175–179 (1984)
- [2] Ekert, A.K.: Quantum cryptography based on bell’s theorem. *Physical Review Letters* **67**(6), 661 (1991)
- [3] Shor, P.W., Preskill, J.: Simple proof of security of the bb84 quantum key distribution protocol. *Physical Review Letters* **85**(2), 441 (2000)
- [4] Lo, H.-K., Curty, M., Qi, B.: Measurement-device-independent quantum key distribution. *Physical Review Letters* **108**(13), 130503 (2012)
- [5] Lucamarini, M., Yuan, Z.L., Dynes, J.F., Shields, A.J.: Overcoming the rate–distance limit of quantum key distribution without quantum repeaters. *Nature* **557**(7705), 400–403 (2018)
- [6] Fang, Y.-T., *et al.*: Implementation of quantum key distribution surpassing the linear rate-transmittance bound. *Nature Photonics* **14**(7), 422–425 (2020)
- [7] Jouguet, P., Kunz-Jacques, S., Leverrier, A., Grangier, P., Diamanti, E.: Experimental demonstration of long-distance continuous-variable quantum key distribution. *Nature Photonics* **7**(5), 378–381 (2013)
- [8] Liao, S.-K., *et al.*: Satellite-to-ground quantum key distribution. *Nature* **549**(7670), 43–47 (2017)
- [9] Yin, J., *et al.*: Entanglement-based secure quantum cryptography over 1,120 kilometres. *Nature* **582**(7813), 501–505 (2020)
- [10] Pirandola, S., Laurenza, R., Ottaviani, C., Banchi, L.: Fundamental limits of repeaterless quantum communications. *Nature Communications* **8**(1), 15043 (2017)
- [11] Boaron, A., *et al.*: Secure quantum key distribution over 421 km of optical fiber. *Physical Review Letters* **121**(19), 190502 (2018)
- [12] Brassard, G., Salvail, L.: Secret-key reconciliation by public discussion. *Advances*

in Cryptology—EUROCRYPT’93, 410–423 (1994)

- [13] Yan, H., Ren, T., *et al.*: Information reconciliation protocol in quantum key distribution system. In: Fourth International Conference on Natural Computation (2008). <https://doi.org/10.1109/ICNC.2008.755> . IEEE
- [14] Gallager, R.: Low-density parity-check codes. IRE Transactions on information theory **8**(1), 21–28 (1962)
- [15] Elkouss, D., Leverrier, A., Alleaume, R., Boutros, J.J.: Efficient rate-adaptive reconciliation for continuous-variable quantum key distribution. IEEE International Symposium on Information Theory, 1879–1883 (2009)
- [16] Martinez-Mateo, J., Elkouss, D., Martin, V.: Key reconciliation for high performance quantum key distribution. Scientific Reports **3**(1), 1576 (2013)
- [17] Elkouss, D., Leverrier, A., Alleaume, R., Boutros, J.J.: Information reconciliation for quantum key distribution. Quantum Information & Computation **11**(3-4), 226–238 (2011)
- [18] Buttler, W.T., Lamoreaux, S.K., Hughes, R.J., Morgan, G.L., Peterson, C.G.: Fast, efficient error reconciliation for quantum cryptography. Physical Review A **67**(5), 052303 (2003)
- [19] Arikan, E.: Channel polarization: A method for constructing capacity-achieving codes for symmetric memoryless channels. IEEE Transactions on Information Theory **55**(7), 3051–3073 (2009)
- [20] Jouguet, P., Kunz-Jacques, S.: High performance error correction for quantum key distribution using polar codes. Quantum Information & Computation **14**(3-4), 329–338 (2014)
- [21] Hu, X., *et al.*: An improved cascade error reconciliation protocol for quantum key distribution. Quantum Information Processing **18**(10), 313 (2019)
- [22] Mueller, R., De Lazzari, C., Chirici, F., Vagniluca, I., Oxenløwe, L.K., Forchhammer, S., Zavatta, A., Bacco, D.: Performance of cascade and ldpc codes for information reconciliation on industrial quantum key distribution systems. IET Quantum Communication (2024) <https://doi.org/10.1049/qtc2.70003>
- [23] Tupkary, D., *et al.*: Using Cascade in Quantum Key Distribution (2023)
- [24] Pedersen, T.B., Toyran, M.: High Performance Information Reconciliation for QKD with CASCADE (2013)
- [25] Dixon, A.R., Dynes, J.F., Lucamarini, M., Fröhlich, B., Sharpe, A.W., Plews, A., Yuan, Z.L., Shields, A.J.: High speed continuous variable quantum key distribution with a pulsed external cavity diode laser. Optics Express **23**(6), 7583–7592

(2015)

- [26] Yuan, Z.L., *et al.*: 10-mb/s quantum key distribution. *Journal of Lightwave Technology* **36**(16), 3427–3433 (2018)
- [27] Walenta, N., *et al.*: A fast and versatile quantum key distribution system with hardware key routing and wavelength multiplexing. *New Journal of Physics* **16**(1), 013047 (2014)
- [28] Kiktenko, E.O., *et al.*: Scalable quantum cryptography network with conditionally free-space optical links. *IEEE Access* **8**, 13110–13123 (2020)
- [29] Lo, H.-K., Ma, X., Chen, K.: Decoy state quantum key distribution. *Physical Review Letters* **94**(23), 230504 (2005)
- [30] Hu, L., Liu, H., Lin, Y.: Parameter optimization of cascade in quantum key distribution. *Optik* **181**, 156–162 (2019) <https://doi.org/10.1016/j.ijleo.2018.12.023>
- [31] Devetak, I., Winter, A.: Distillation of secret key and entanglement from quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **461**(2053), 207–235 (2005)