Article

Not peer-reviewed version

Topology-Oblivious Random-Walk Key Relaying in Quantum Key Distribution Networks

Krišjānis Petručeņa *, Sergejs Kozlovičs, Juris Vīksna, Elīna Kalniņa, Reinis Isaks, Edgars Celms, Lelde Lāce, Edgars Rencis

Posted Date: 8 April 2026

doi: 10.20944/preprints202604.0533.v1

Keywords: quantum key distribution; QKD networks; trusted-node relaying; random walks; topology-oblivious routing; privacy amplification

 doi:10.20944/preprints202604.0533.v1

Article

Topology-Oblivious Random-Walk Key Relaying in Quantum Key Distribution Networks

Krišjānis Petručeņa *, Sergejs Kozlovičs, Juris Vīksna, Elīna Kalniņa, Reinis Isaks, Edgars Celms, Lelde Lāce and Edgars Rencis

Institute of Mathematics and Computer Science, University of Latvia, Latvia
* Correspondence: krisjanis.petrucena@lumii.lv

Abstract

Quantum key distribution (QKD) networks require relaying when distant key management entities share no direct quantum link. Most relay strategies, however, rely on centralized control or globally maintained routing state. This paper asks whether useful security and efficiency can still be obtained with topology-oblivious stochastic forwarding. It studies the security-overhead trade-off in a model in which fragmented key material is relayed via random-walk variants and reconstructed under privacy amplification. Under a restricted model with at most one compromised relay, the analysis asks whether strictly local forwarding can retain useful information-theoretic security. Evaluation on the GÉANT topology, representing a European academic backbone network, shows clear differences between random-walk variants. The proposed highest-score-neighbor local path-diversification heuristic reduces the risk that relayed key material passes through a compromised node. The evaluation also shows that a preliminary loop-erasure step significantly shortens sampled routes and improves throughput in the model. These findings position topology-oblivious stochastic forwarding as a decentralized alternative to global-state maintenance or centralized orchestration in QKD networks.

Keywords: quantum key distribution; QKD networks; trusted-node relaying; random walks; topology-oblivious routing; privacy amplification

1. Introduction

Quantum key distribution (QKD) can provide symmetric keys with *information-theoretic security* (ITS), but present-day systems are limited in distance and secret key rate. To address distance limitations, deployments are composed of multiple QKD links and a key management layer. The key management layer is responsible for authorization and key forwarding and delivery. Ultimately, any two QKD network nodes - key management entities (KMEs) - can establish a shared secret key. If no node is compromised and we use one-time pad (OTP) encryption for hop-by-hop relaying, we can reclaim ITS.

Most QKD-network proposals rely on centralized or SDN-assisted orchestration, or on topology-aware path selection [1]. In this paper, we ask whether useful *efficiency* (security-overhead trade-off) can be achieved when forwarding is topology-oblivious. In particular, we analyze the case in which a single relay node has been compromised. If viable, this approach could further simplify the decentralized key-management layer, especially in growing networks with nodes added over time.

By contrast, in well-known distributed link-state routing protocols such as OSPF [2], routers flood link-state advertisements (LSAs), spreading local link information in a gossip-like manner so that all routers can learn the topology. Afterwards, nodes can compute end-to-end paths using the learned topology. In this paper, we propose an alternative random-walk construction and avoid gossip protocols.

We study a topology-oblivious stochastic relay scheme that we call a *random flow*. In one random flow from source s to destination t, the source splits material into M fragments and emits those

 doi:10.20944/preprints202604.0533.v1

fragments in parallel. Each fragment is then forwarded independently using a random-walk *variant*. The destination reconstructs the recovered material and applies seeded privacy amplification. In this setting, routing behavior directly determines the security-overhead trade-off through proxy metrics such as exposure χ (worst-case probability of traversing a malicious node) and expected hop count $h_{s,t}$. Our contributions are the following:

1. A topology-oblivious random-flow relaying model in which fragmented key material is forwarded by stochastic rules that avoid global adjacency knowledge, link-state maintenance, and end-to-end path computation.
2. A highest-score neighbor local diversification heuristic that improves worst-case exposure without requiring global topology knowledge or even the node count.
3. A scouting-based loop-erasure mechanism that shortens realized payload routes, reduces queueing pressure, and eliminates self-induced cyclic waiting in the model.
4. A formal entropy-overhead connection based on a leftover-hash-style bound, together with a separate heuristic estimator used in the simplified evaluation model.
5. A simulation study on reconstructed and synthetic topologies that compares the evaluated random-walk variants under common exposure, hop-count, efficiency, and throughput proxies.

The remainder of the paper is organized as follows. Section 2 reviews the QKD-network setting and related work. Section 3 introduces the random-flow model, random-walk variants, privacy-amplification view, and loop-erasure mechanism. Section 4 presents the simulation setup and comparative results. Section 5 concludes the paper.

We study this approach under the restricted threat model of Section 3.1 and the evaluation setup of Section 4. In short, we assume exactly one compromised relay, focus on biconnected source–target pairs (s, t), and target information-theoretic security.

2. Background and Related Work

BB84 remains the canonical example of a QKD protocol. In that setting, classical bits are encoded into quantum states and sent over an optical channel one photon at a time. After sifting and reconciliation, the endpoints can detect eavesdropping. In terrestrial deployments, distance is a limitation: fiber loss and other implementation imperfections cause the secret-key rate to drop quickly as the span grows. Outside experimental satellite settings [3], practical links are typically limited to roughly $\approx$150 km [4].

Because of this range limit, larger deployments use relay nodes that pass key material from one hop to the next until it reaches the destination. Taken together, the QKD links and relay-capable nodes form a graph-like QKD network. The tasks of relaying, buffering, and authorizing access to key material are handled by the key-management layer.

At that transport layer [5], two nodes α, β are adjacent when they share a direct QKD link. Such neighbors continuously obtain a common sequence of link keys $\{\ell_{\alpha,\beta}\}$, usually identified by UUIDs. These keys may be buffered and then consumed for OTP encryption of data sent over an authenticated classical channel. To establish an end-to-end key between non-adjacent nodes s and t, the source can generate a fresh final key K_f, choose a relay path, and send $K_f \oplus \ell_{s,\eta}$ to its first-hop neighbor η . Each relay on the path can decrypt, recover K_f, and re-encrypt it for the next hop. This explains why the security model of trusted-node QKD networks depends critically on intermediate relays remaining uncompromised.

At the application interface, ETSI GS QKD 014 [6] defines a RESTful HTTPS API through which applications obtain QKD-derived keys as JSON containers containing pairs of $(\text{UUID}, \text{key})$. The standard separates Secure Application Entities (SAEs) from Key Management Entities (KMEs), with the KMEs placed inside Trusted Nodes together with one or more QKD Entities (QKDEs) that terminate the physical QKD links. Mutual certificate-based authentication between SAE and KME is part of the model, but inter-KME relay protocols are left unspecified. Figure 1 sketches this separation between the standardized application-facing API and the non-standard relay logic behind it.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 April 2026 doi:10.20944/preprints202604.0533.v1

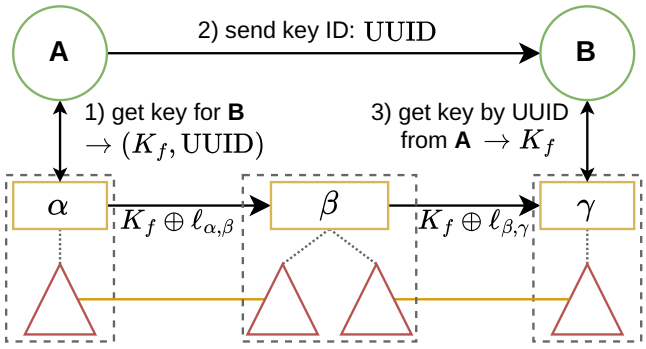

Figure 1. Illustration of ETSI 014-style key delivery. SAE A requests a fresh key from KME α and passes the returned key identifier to SAE B, which later retrieves the same key from KME γ. The corresponding KMEs are embedded in trusted nodes with QKD link endpoints, while the underlying Q-KMS transfers the final key K_f hop-by-hop using OTP protection with $\ell_{\alpha,\beta}$ and $\ell_{\beta,\gamma}$.

To reduce dependence on any single intermediate "trusted" node, the end-to-end key K_f may be decomposed and sent over several strictly or partially node-disjoint routes. This family of approaches is usually called multi-path QKD. In the simplest case, the destination reconstructs K_f from independently delivered pieces, for example by combining $K_1, \ldots, K_n$ as $K_f \leftarrow K_1 \oplus \cdots \oplus K_n$. The adversary then has to compromise at least one relay on every relevant path to recover the final key. More general (t, n) secret-sharing constructions offer tolerance against partial share leakage, albeit with additional overhead [7,8]. Such techniques are a standard way to reason about partially trusted QKD networks.

The software stack that performs relaying, buffering, and authorization is often referred to as a Quantum KMS (Q-KMS). This differs from a conventional KMS, whose primary role is usually key lifecycle management: key generation, access control, auditing, and cryptographic service exposure, frequently via hardware security modules. In practice, vendors such as ID Quantique, Toshiba, and LuxQuanta typically ship proprietary Q-KMS software alongside their hardware, which can limit interoperability across different vendors.

Most routing proposals for trusted-relay QKD networks are topology-aware and assume some form of global or near-global state, often maintained through link-state-style control exchange [9,10]. Recent work has also proposed an explicitly OSPF-based distributed key-relay architecture for QKD networks [11]. A recent systems-level comparison between decentralized key pre-distribution and on-demand relaying is given in [12]; unlike our work, it proactively builds end-to-end key pools rather than using local stochastic forwarding. Partially trusted approaches instead split key material across multiple paths or shares [7,8,13,14], while stochastic routing has also been explored before [15,16]. Recent zero-trust proposals pursue a stronger goal than the one studied here [17].

3. Random Flow

At a high level, we replace deterministic routing by stochastic forwarding of fragmented key material: each of the M fragments is routed by a random walk *variant* from source s towards target t. Node t collects all fragments, and derives the final key K via seeded privacy amplification. We find that the random walk (RW) variant plays a significant role in *exposure*, the max (worst-case) probability of traversing the malicious node.

Building on the QKD transport model from the background section, we now switch to the algorithmic abstraction. We represent the QKD network as a simple, undirected graph $G = (V, E)$, where V denotes the set of trusted nodes and E denotes the set of QKD links . The per-link keys $\ell_{u,v}$ introduced earlier are treated here simply as the hop-by-hop OTP resource available on edge (u, v). To simplify the notation in the analysis, we assume that both the link-key blocks and the derived end-to-end key have length 256 bits.

We assume a reactive (rather than a proactive) model. In reactive mode, key *allocation* request triggers a *transmission*, and transmissions are not initiated ahead of time. A transmission (node $s \rightarrow$

 doi:10.20944/preprints202604.0533.v1

node t) is a key-relaying session by which a *block* of final keys $K_1, \ldots, K_\theta$ is established. In the base model (without loop erasure), the source s emits M raw fragments $f_1, \ldots, f_M$, which eventually reach destination t by means of random walks. We can define the ratio ρ between the final block size (full-entropy key count) θ and emitted fragment count M as $\rho := \theta \div M$.

The motivation behind emitting multiple fragments is clear: a node may be compromised at some point in time, and, by increasing M, we decrease the probability that all or most of them will traverse the malicious node. We later denote exposure by $\chi_{s,t}$ - the worst-case probability that a fragment in transmission $s \to t$ traverses the malicious node v, maximized over all choices of $v \in V \setminus \{s, t\}$ available to the adversary.

3.1. Threat Model and Objectives

For confidentiality analysis, we assume at most one compromised relay $m \in V \setminus \{s, t\}$. Compromise reveals the relay's full internal state and the plaintext of every fragment it forwards. Public-key wrapping of fragments does not satisfy the information-theoretic goal of QKD and remains vulnerable to harvest-now decrypt-later attacks.

The compromised relay may drop, misroute, or inject traffic. We do not attempt to address availability. For confidentiality, a fragment is treated as compromised once it passes through m. Traffic injection is handled separately by standard authentication mechanisms.

We assume honest endpoints (s and t), authenticated classical channels based on PKI and mutual TLS, and trusted QKD hardware, excluding side-channel attacks [18].

The design goals are:

1. Information-theoretic security.
2. Fail to establish a key when ITS is no longer possible.
3. Tolerance of one compromised relay for biconnected s, t pairs.
4. Compatibility with ETSI GS QKD 014.
5. Independent fragment forwarding without global adjacency knowledge, link-state exchange, or end-to-end path computation.

Forwarding decisions use neighbor information and variant-specific token state; in the NC variant, the source additionally knows the global node-identifier universe. The topology itself may change over time. We study what confidentiality and overhead can be achieved under the assumption of exactly one malicious relay.

3.2. Random Walk Notation

We represent each in-flight (travelling from s to t) fragment f_i by a *token*. Concretely, transmission token i is a tuple consisting of id - transmission identifier, source and target pair (s, t), i - token index, f_i - fragment itself, and σ_i - variant-specific local state. State σ_i is the only element that may change over time.

$$\mathsf{token}_i = (\mathsf{id}, s, t, i, f_i, \sigma_i),$$

Each token i is forwarded by a discrete-time walk. The trajectory of the token i during the sampled walk is $X_0, X_1, \ldots, X_h$, with $X_0 = s$, $X_{k+1} \in N(X_k)$, $X_h = t$, where $N(X_k)$ is the neighborhood of X_k. The whole trajectory is a sampled random sequence $(X_k)_{k=0}^{h}$, where h is the hop count. The probability distribution from which X_{k+1} is sampled is P_{variant}.

$$X_{k+1} \sim P_{\mathsf{variant}}(\,\cdot \mid X_k, \sigma_k),$$

It depends only on the current token i and variant-specific state σ (vertices are stateless under this objective and under our interpretation of the random-walk setting). The dot in $P_{\mathsf{variant}}(\,\cdot \mid \ldots)$ is a placeholder, and altogether the expression means "distribution over possible next nodes". The walk terminates when it first hits the target.

 doi:10.20944/preprints202604.0533.v1

For readability, we will adopt the notation $P(v \to u \mid \star)$ instead of $P(X_{k+1} \mid X_k, \sigma_k)$, where $v = X_k$, $u = X_{k+1}$ and $\star$ is syntactically substituted for some predicate that can be evaluated from σ_k. It is the probability to go from v to u given position X_k and state σ_k.

3.3. Base Random Walk Variants

Simple random walk (R). The simple random walk variant R is memoryless: at node v, the token chooses the next hop uniformly at random. R induces a Markov chain on V.

$$P_{\mathrm{R}}(v \to u) = \begin{cases} 1/|N(v)|, & u \in N(v), \\ 0, & \text{otherwise.} \end{cases}$$

Non-backtracking random walk (NB). Non-backtracking NB suppresses immediate return. The token state σ carries a single field $\mathsf{prev} \in V \cup \{\mathsf{null}\}$ with the previous node (or null at the start). Let $N'(v) = N(v) \setminus \{p\}$, where $p = \mathsf{prev}$. Then

$$P_{\mathrm{NB}}(v \to u \mid p) = \begin{cases} 1 \,/\, |N'(v)|, & u \in N'(v), \\ 1, & u = p \text{ and } d(v) = 1, \\ 0, & \text{otherwise.} \end{cases}$$

After choosing u, the token updates $\mathsf{prev} \leftarrow v$. Equivalently, NB is a first-order Markov chain on directed edges. On regular expander graphs, NB can "mix" provably faster than R [19]. Informally, an *expander* is a sparse graph with strong connectivity.

Least-recently-visited walk (LRV). LRV biases the walk away from recently visited vertices. We use an LRV-*vertex* rule: token i maintains timestamps $\mathsf{last} : V \to \mathbb{N}_0$, where $\mathsf{last}[x]$ is the most recent time at which the token visited x. We initialize $\mathsf{last}[s] = 1$, and return $\mathsf{last}[x] = 0$ by default. At step k with $X_k = v$,

$$X_{k+1} \in \arg\min_{u \in N(v)} \mathsf{last}[u],$$

breaking ties uniformly. Unvisited neighbors (with value 0) are preferred. Local LRV-type policies are well studied in graph exploration; they can improve practical coverage [20].

3.4. Privacy Amplification

Because a compromised node may eavesdrop raw key material, and in the security analysis we assume at most one compromised relay, we can and *should* divide the source material into M fragments and send them via at least partially disjoint paths.

Let $\chi_{s,t}$ or *exposure* be the max probability of the eavesdropper positioned on the "worst" intermediate relay to observe a fragment travelling $s \to t$. To be precise, define $p_v^{(s,t)}$ as the probability of visiting v and let $\chi_{s,t} = \max_{v \in G \setminus \{s,t\}} p_v^{(s,t)}$. Because nodes do not know the topology in advance, we may need to assume a χ value in advance when choosing M (fragment count). A safe χ value greatly depends on the random walk variant.

Assume we transmit M fragments and let G denote the number of "good" fragments, i.e., fragments not observed by the compromised relay. Then $\Pr[G \geq g]$ follows a cumulative binomial distribution (see Figure 2) and can be calculated as follows:

$$\Pr[G \geq g] = \sum_{k=g}^{M} \binom{M}{k} (1-\chi)^k \chi^{M-k}$$

Let us define $g_\alpha^\star(M, \chi)$ as the largest g such that $\Pr[G \geq g]$ is at least, e.g., $\alpha = 99.99\%$. It is the maximum number of fragments out of M that we can guarantee to be safe (contain full entropy) in transmission $s \to t$, for a given value of $\chi_{s,t}$.

 doi:10.20944/preprints202604.0533.v1

$$g^{\star}_{\alpha}(M,\chi) \;=\; \max\{\, g \in \{0,\dots,M\} : \Pr[G \geq g] \geq \alpha \,\}$$

See Table 1 for $g^{\star}_{\alpha}$ values corresponding to different assumed χ values. For $M = 1024$ and $\chi = 95\%$, the table gives $g^{\star}_{99.99\%}(1024, 95\%) = 27$. The choice $\chi = 95\%$ is not arbitrary: later measurements in Section 4.2 on GÉANT and on the synthetic graph show worst-case exposures for the better-performing variants still concentrated around 93–96%, making 95% a useful conservative design point for illustration.

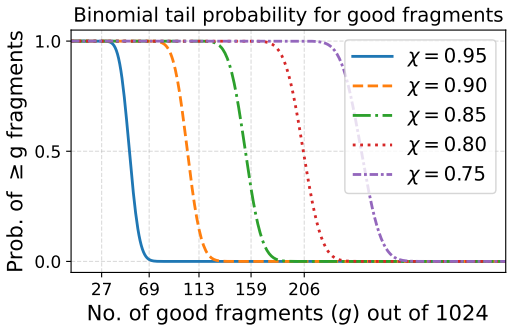

Figure 2. Prob of receiving $\geq g$ good fragments.

Table 1. Maximum g values such that $P[\# \geq g] > \alpha$ holds for various χ (exposure) values. The table helps determine the maximum number of "good" fragments that we can guarantee will arrive safely at the target.

	$g^{\star}_{99\%}(M,\chi)$					$g^{\star}_{99.99\%}(M,\chi)$				
M	95%	90%	85%	80%	75%	95%	90%	85%	80%	75%
32	0	0	1	2	3	0	0	0	0	0
64	0	2	4	6	8	0	0	1	3	5
128	1	6	10	16	21	0	2	6	10	15
256	5	15	26	37	48	2	10	19	29	39
512	15	36	59	82	106	9	28	48	70	93
1024	36	81	128	175	224	27	69	113	159	206

Let $\theta_{s,t}$ denote the number of final 256-bit keys extracted after privacy amplification for transmission $s \to t$. Then $g^{\star}_{\alpha}$ should be understood as a lower bound on the number of good fragments, not directly on $\theta_{s,t}$.

To keep the formal result separate from the simplified evaluation model, define the *estimated extracted count*

$$\hat{\theta}_{s,t} \;:=\; g^{\star}_{\alpha}(M, \chi_{s,t}).$$

This estimator is used later as a heuristic proxy for extracted output under the simplified assumptions of the evaluation, but it is not itself the formal privacy-amplification guarantee. More formally, let each good fragment contribute at least μ bits of conditional min-entropy. By the leftover hash lemma / privacy-amplification bound [21,22], for privacy-amplification error $\varepsilon_{\mathrm{pa}}$, the extractor output must satisfy

$$\theta_{s,t} \cdot 256 \;\leq\; \mu\, g^{\star}_{\alpha}(M, \chi_{s,t}) - 2\log_2(1/\varepsilon_{\mathrm{pa}}).$$

Thus, $\hat{\theta}_{s,t} = g^{\star}_{\alpha}(M, \chi_{s,t})$ should be read only as a heuristic extracted-count estimator for the special case in which $\mu \approx 256$ and the extractor overhead is neglected. The formal guarantee remains the leftover-hash-style inequality above.

For later use, define the *estimated yield* as the estimated extracted count normalized by the fragment count. For a fixed ordered pair (s,t), the corresponding pair-specific estimator is

$$\rho^{\mathrm{est}}_{s,t} \;:=\; \frac{\hat{\theta}_{s,t}}{M}.$$

 doi:10.20944/preprints202604.0533.v1

If, instead, an implementation fixes a conservative design exposure χ_{variant} for an entire RW variant, we define the corresponding *assumed* yield

$$\rho^{\text{assm}} := \frac{\hat{\theta}^{\text{assm}}}{M} = \frac{g^{\star}_{\alpha}(M, \chi_{\text{variant}})}{M}$$

3.5. Exposure Reduced RW Variants

The experimental results in Section 4.2 show that the aforementioned variants R, NB, and LRV have high $\max_{s,t} \chi_{s,t}$ values. On each evaluated graph, there exists a pair s, t such that $\chi_{s,t} \geq 96.4\%$ for each of these walk variants.

See Figure 3 for intuition about why high χ can arise. There is often a longer bypass path connecting s and t, but the walk is likely to be diverted back into the more "congested" direct connection. To address this, we can temporarily color nodes one by one, or otherwise assign them lower priority throughout the walk. Different walks within the same transmission should penalize different nodes. This motivates the following two constructions.

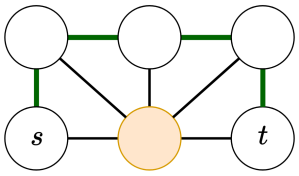

Assume the yellow node at the bottom, between s and t, is compromised. There exists one "good" path from s to t. Along this path, at each v, the correct neighbor must be chosen.

Figure 3. High χ construction example.

The following constructions reduce max χ from 96.4% to 93.3%. In the paper's simplified exposure-based intuition, the fraction of fragment bits that avoid the compromised relay therefore increases by $(1 - 0.933) \div (1 - 0.964) = 86\%$.

One-by-one node-coloring (NC). If s and t are biconnected, we may color nodes one-by-one, effectively removing them from the graph. This requires knowledge of the global node-identifier universe. Under the paper's terminology, NC remains topology-oblivious because it does not require global adjacency knowledge, link-state exchange, or path computation, but it is less local than R, NB, LRV, and HS. If $M > |V|$, we can choose the vertex identifier in circular order. Otherwise, when choosing the next hop, we apply the LRV walk strategy.

As a side note, this method allows for a particularly easy construction of additive (XOR) secret sharing. For each end-to-end key $K \in \{0,1\}^{256}$, the source samples $f_1, \ldots, f_{M-1} \xleftarrow{\$} \{0,1\}^{256}$ and sets the final fragment f_M so that all fragments together XOR to K

$$f_M := K \oplus \left(\bigoplus_{i=1}^{M-1} f_i \right), \qquad \text{so that} \quad K = \bigoplus_{i=1}^{M} f_i.$$

Any subset of at most $M - 1$ shares is uniformly distributed and independent of K, and hence reveals no information about K in the information-theoretic sense. This is known as additive (XOR) secret sharing [23]. The malicious party must observe all f_i to get K.

Although additive (XOR) secret sharing is easy to implement and reason about, we instead quantify χ and apply entropy extraction because this yields a slightly higher ρ.

 doi:10.20944/preprints202604.0533.v1

Highest-score neighbor (HS). HS is a seed-based diversification heuristic. For each fragment token i, the source s samples a fresh random seed $\zeta_i \xleftarrow{\$} \{0,1\}^\lambda$ at the start of the walk. This seed defines a deterministic per-walk score for each vertex $u \in V$

$$\mathsf{score}_i(u) = h(u, \zeta_i),$$

where h is a deterministic mixing function of the vertex identifier u and seed ζ_i. Different ζ_i induce different vertex rankings without requiring topology knowledge. Intuitively, a previously very probable relay will receive a low score in some walks. HS chooses a neighbor with the largest assigned score.

$$X_{k+1}^{(i)} \in \arg\max_{u \in U_i(v)} \mathsf{score}_i(u)$$

If all neighbors of v have already been visited, HS falls back to the LRV rule. Token i also maintains visit timestamps - last : $V \to \mathbb{N}_0$ (just like in LRV). Let $U(v) \subseteq N(v)$ be the set of unvisited vertices. If a neighbor $u \notin U(v)$, then we override $\mathsf{score}_i(u)$ to 0.

3.6. Efficiency and Throughput

For a fixed source–destination pair (s,t), let $H_{s,t}$ denote the random hop count of one fragment walk, and let $h_{s,t} := \mathbb{E}[H_{s,t}]$ be the expected hop count. As previously assumed, the fragment size is 256 bits and matches link key size. One delivered fragment therefore consumes, in expectation, $h_{s,t}$ link keys. A transmission consumes $M \cdot h_{s,t}$ link keys.

Suppose a transmission uses M fragments. Recall from Section 3.4 that

$$\rho_{s,t}^{\mathsf{est}} = \frac{\hat{\theta}_{s,t}}{M}, \qquad \rho^{\mathsf{assm}} = \frac{\hat{\theta}^{\mathsf{assm}}}{M} = \frac{g_\alpha^\star(M, \chi_{\mathrm{variant}})}{M}.$$

Because QKD-derived link keys are a scarce resource, we define the variant-specific *model-based extracted efficiency* as the estimated extracted output bits per consumed QKD-derived link-key bit:

$$\eta_{s,t}^{\mathsf{est}} := \frac{\hat{\theta}_{s,t}}{M} \frac{1}{h_{s,t}} = \rho_{s,t}^{\mathsf{est}} h_{s,t}^{-1}$$

When reporting a topology-wide metric over a set $\mathcal{P}$ of ordered (s,t) pairs, we use

$$\overline{\eta} := \frac{1}{|\mathcal{P}|} \sum_{(s,t)\in\mathcal{P}} \eta_{s,t}, \quad \overline{\rho} := \frac{1}{|\mathcal{P}|} \sum_{(s,t)\in\mathcal{P}} \rho_{s,t}, \quad \overline{h^{-1}} := \frac{1}{|\mathcal{P}|} \sum_{(s,t)\in\mathcal{P}} \frac{1}{h_{s,t}}.$$

For the pair-specific estimator, the average efficiency admits the exact decomposition

$$\overline{\eta^{\mathsf{est}}} = \overline{\rho^{\mathsf{est}}}\,\overline{h^{-1}} + \mathrm{Cov}(\rho^{\mathsf{est}}, h^{-1}),$$

where the covariance is taken over $(s,t) \in \mathcal{P}$. This is just the standard identity [24] $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + \mathrm{Cov}(X,Y)$ applied to $X = \rho_{s,t}^{\mathsf{est}}$ and $Y = h_{s,t}^{-1}$. This decomposition is useful because it separates three effects: the average security yield ρ, the average routing cost $h_{s,t}$ (equivalently, $h_{s,t}^{-1}$ in the formula), and the extent to which $\rho_{s,t}^{\mathsf{est}}$ and $h_{s,t}^{-1}$ are correlated across (s,t) pairs. For the assumed quantity, ρ^{assm} is constant within a topology/variant, because it is computed from a single conservative design exposure χ_{variant} rather than the pair-specific exposure $\chi_{s,t}$. Therefore, the covariance term is identically zero.

Next, define the raw fragment throughput $T_{s,t}$ as the long-run average rate, in bits/s, at which fragment bits arrive at the destination t. We assume that only one ordered pair (s,t) is active at

 doi:10.20944/preprints202604.0533.v1

a time, that QKD links continuously generate and buffer keys (starting with empty buffers). The corresponding model-based extracted throughput is

$$R^{\text{est}}_{s,t} \;:=\; T_{s,t}\,\rho^{\text{est}}_{s,t}, \quad \text{equivalently } R^{\text{est}}_{s,t} = T_{s,t}\,\eta^{\text{est}}_{s,t}\,h_{s,t}$$

Thus, $\eta^{\text{est}}_{s,t}$ captures model-based extracted efficiency, while $R^{\text{est}}_{s,t}$ captures the corresponding extracted-throughput proxy under the simplified entropy model.

3.7. Loop Erasure

We separate route discovery from payload transmission. In the first phase, the source emits a lightweight *scouting token* over the authenticated classical channel only. The scouting token carries no secret fragment material and consumes no QKD-derived link-key bits. It is forwarded according to the same stochastic forwarding rule as the corresponding payload walk and records the visited history

$$W = (v_0 = s, v_1, \ldots, v_\tau = t).$$

When we later report exposure for loop-erased variants, we still measure it on this original scouting walk as a conservative upper bound on literal payload exposure.

During the walk a time-bounded link key reservation is applied to traversed links. Ultimately, if some relay along this sampled walk cannot admit the request, the session may be rejected before any key material is transmitted; in a REST realization, this can be surfaced as an implementation-level failure such as HTTP `503`.

After the scouting token reaches t, the recorded walk W is converted into a realized payload route by *chronological loop erasure* [25,26]. Concretely, one scans W from left to right and, whenever a vertex reappears, deletes the closed subwalk between the earlier occurrence and the repeat, while retaining a single copy of the repeated vertex. Let $\mathrm{LE}(W)$ denote the resulting simple s–t path. In the second phase, the actual fragment material is forwarded hop by hop along $\mathrm{LE}(W)$, with each hop protected by fresh QKD-derived link-key bits used as a one-time pad.

Strictly speaking, for the NB, LRV, NC, and HS variants, $\mathrm{LE}(W)$ is not the *classical* loop-erased random walk distribution from probability theory; here loop erasure is used only as a path-simplification operator applied to the sampled scouting history.

This two-phase construction provides two operational benefits. First, removing loops shortens the realized payload route, thereby reducing expected hop count, QKD key consumption, and queueing pressure. Second, because the realized payload route is simple, self-induced cyclic waiting caused by one fragment revisiting the same relays is eliminated. Rejection, if necessary, occurs before secret material enters the relay path.

The interpretation of exposure under loop erasure requires care. A compromised relay that appears only on an erased segment of the scouting walk does not directly observe fragment plaintext, because the scouting phase is classical only. However, such a relay may still influence the realized route by biasing the scouting trajectory or by affecting admission decisions. Therefore, when reporting $\chi_{s,t}$ for loop-erased variants, we count scouting hits as a *conservative upper bound* on literal payload exposure. This preserves comparability with the no-loop-erasure case while avoiding an optimistic estimate of adversarial influence.

3.8. Implementation Considerations

If the destination t is an adjacent neighbor of the current relay, i.e. $t \in N(v)$, the relay should bypass the stochastic rule and forward the fragment directly to t. This is the model we assume in the following evaluation section.

At the destination, the recovered fragment material is concatenated into Z and processed by seeded privacy amplification,

$$K = \mathrm{Ext}_S(Z),$$

 doi:10.20944/preprints202604.0533.v1

where the seed S is public but authenticated end to end. In practical QKD post-processing, privacy amplification is a standard final step, and universal-hash implementations such as Toeplitz hashing are commonly used [27]. The output length and entropy constraints for this step are those already defined in Section 3.4.

Each fragment walk should use fresh independent randomness, including independent diversification seeds for HS, so that the assumptions behind the extraction step remain valid. Among the diversification variants, NC is less deployment-friendly than HS because it requires a globally known set of node identifiers. Under the paper's terminology, this still fits the topology-oblivious setting, but it is weaker in locality than HS.

4. Experimental Evaluation

We focus on three questions: (i) *exposure*—how large the worst-case relay hit probability $\chi_{s,t}$ can be for some source–target pair, and what conservative design exposure this suggests when choosing the target extracted block size θ and, where applicable, the fragment count M; (ii) *efficiency*—what extracted-yield estimators and model-based output per consumed QKD-derived link-key bit the evaluated variants achieve under the paper's simplified entropy model; and (iii) *scalability*—how these quantities and the routing cost evolve with network size. We address these questions on topologies resembling existing quantum network deployments.

4.1. Simulated Topologies

Trusted-node QKD deployments are currently small (typically tens of nodes) due to the cost of dark fiber and QKD equipment. For context, the largest openly reported QKD deployment is the China Quantum Communication Network (CN-QCN), spanning 145 nodes [28]; however, CN-QCN uses centralized network management and exhibits a hierarchical topology with many articulation points. For our experiments, we selected three topologies spanning increasing size and structural complexity: 6, 14, and 43 nodes (SECOQC, NSFNET, GÉANT). See Figures 4–6. All graphs are visualized in Gephi using a geographic layout based on approximate site latitude/longitude, and node colors denote Gephi modularity communities with no physical or administrative meaning. We deliberately focus on graphs with a nontrivial biconnected structure, and, whenever we report ITS-oriented exposure, yield, efficiency, or throughput quantities, we restrict attention to biconnected ordered pairs (s, t).

- **SECOQC.** The metro-scale "SEcure COmmunication based on Quantum Cryptography" testbed (Figure 4) with 6 nodes. SECOQC was a major European research initiative involving 41 research and industrial organizations from the EU, Switzerland, and Russia [29]. It ran from April 2004 to October 2008 and employed heterogeneous QKD link technologies (e.g., entanglement-based, decoy-state BB84, continuous-variable), with short physical spans typical of early field deployments.
- **NSFNET.** The NSFNET T1 backbone from 1991 with 14 nodes [30] (Figure 5). Although NSFNET is a classical network, it serves as a medium-sized, historically relevant analog: in early wide-area optical transport, capacity was scarce and expensive, motivating sparse topologies and careful end-to-end provisioning. Similar constraints reappear in QKD networks.
- **GÉANT.** The GÉANT GN4 Phase 3 backbone (GN4-3N) is a large, well-connected topology. Our variant (Figure 6), after pruning links longer than 1000 km to improve topology perceptibility in the graph diagram, has 43 nodes and 59 edges. GÉANT is also engaged in Europe's "ultra-secure" communications direction (including EuroQCI-oriented efforts that consider QKD overlays), making this backbone a plausible substrate for a future QKD overlay [31,32].

In practice, the GÉANT and NSFNET topologies would require many additional relay sites because their longest links exceed the $\approx$ 150–300 km range typical of current commercial devices. We do not introduce such relays here; these two graphs are treated as structural proxies.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 April 2026 doi:10.20944/preprints202604.0533.v1

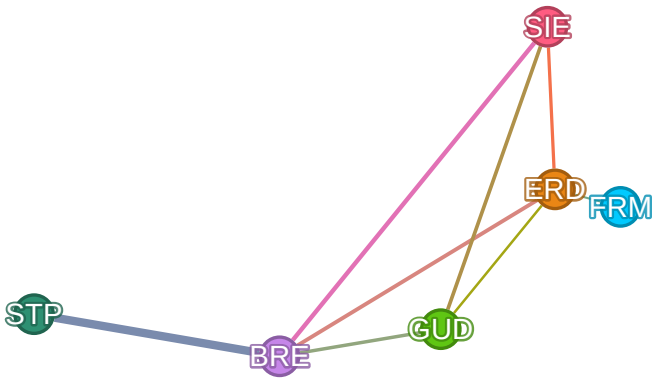

Figure 4. SECOQC Vienna QKD network.

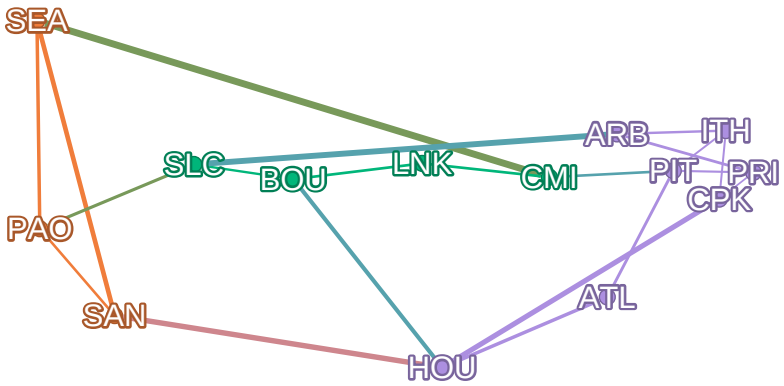

Figure 5. Reconstructed NSFNET T1 topology.

To study scaling trends, we generate a synthetic 2-vertex-connected graph with 99 nodes and 143 edges by placing clusters of 3 nodes at randomly chosen nearby coordinates and connecting each cluster to 2–3 of its 6 closest neighbors; generation is retried if the graph is not 2-vertex-connected. See Figure 7. In that figure, darker nodes have higher sequence numbers (they appear later in snapshots), and edge directions indicate the insertion step when the edge was added. The graph has an average degree of ≈ 2.9, comparable to the real topologies above. Furthermore, we bias edge creation toward geographically closer nodes. Each node is also assigned a sequence number so we can take snapshots at $n = 3, 6, \ldots, 99$ such that the snapshots maintain the same properties.

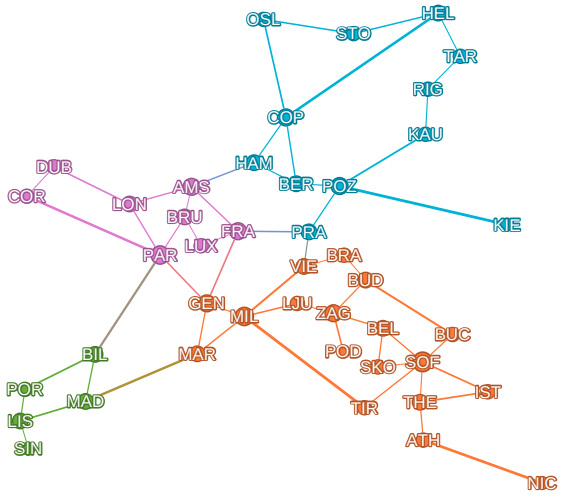

Figure 6. Adapted version of GÉANT GN4-3N.

 doi:10.20944/preprints202604.0533.v1

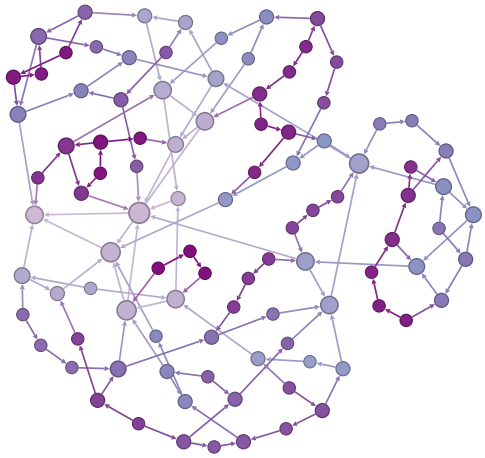

Figure 7. Synthetic graph with 99 nodes.

Table 2 summarizes basic structural metrics of the evaluated graphs.

Table 2. Topology overview. ASP is the average shortest-path length. The biconnected core fraction is the fraction of nodes in the largest biconnected component. Max betweenness is the maximum node betweenness centrality (normalized). The 2-connected column shows the percent of (s, t) pairs for which no node removal would leave s and t in different components

Graph	Nodes	Edges	Diam.	Avg. deg.	ASP	2-connected	Max betw.
GÉANT	43	59	12	2.74	4.682	83.7%	0.4150
NSFNET	14	21	3	3.00	2.143	100%	0.2201
SECOQC	6	8	3	2.67	1.533	66.7%	0.4000
Generated	99	143	10	2.89	4.784	100%	0.2896

4.2. Random Walk Security

To evaluate the security-relevant exposure, we estimate $\chi_{s,t}$ (max intermediate-node hit probability) for a given graph and ordered pair (s, t) as follows. Let W be the number of completed sampled walks and $X^{(i)}$ be the path taken by the i-th walk. The empirically determined $\hat{\chi}$ is

$$\hat{\chi}_{s,t} = \max_{v \in G \setminus \{s,t\}} p_v^{(s,t)}, \text{ where } \quad p_v^{(s,t)} = \frac{1}{W} \sum_{i=1}^{W} \left[v \in X^{(i)} \right]$$

In each graph, we then identify $\max_{s,t} \hat{\chi}$ for each RW variant. This is the conservative design exposure used when dimensioning the target extracted block size θ and, when needed, the fragment count M, via the lower-bound quantities $g_\alpha^\star(M, \chi)$, $\theta_{s,t}$, and $\rho_{s,t}^{\text{est}}$. For many ordered pairs, this value is 1 because an articulation point is unavoidable. We therefore compute $\hat{\chi}$ only over biconnected (s, t) pairs. For loop-erased variants, $X^{(i)}$ still denotes the original sampled/scouting walk rather than only the realized payload route, so the reported $\hat{\chi}$ remains a conservative upper bound on literal payload exposure. Table 3 reports the detailed GÉANT breakdown, including the maximizing (s, t, v) triple and the average $\hat{\chi}$, while Table 4 summarizes per-graph max and median $\hat{\chi}$ across variants. SECOQC is omitted from the exposure overview because it has no indirect biconnected source–target pair under this threat model.

Table 3. RW variant $\max_{s,t} \hat{\chi}$ values (in %) on GÉANT and corresponding intermediate vertex v.

Variant	Max $\hat{\chi}$	s	t	v	Avg $\hat{\chi}$	Median $\hat{\chi}$
R	99.2	TIR	LIS	MIL	79.4	89.7
NB	96.4	TIR	LIS	MIL	74.1	83.5
LRV	96.1	TIR	LIS	MIL	72.5	81.1
NC	93.7	POR	COR	PAR	71.8	80.4
HS	92.6	MAD	COR	PAR	69.4	77.2

 doi:10.20944/preprints202604.0533.v1

Table 4. Exposure overview by graph and RW variant. For each graph and variant, we report the worst-case pair exposure $\max_{s,t} \hat{\chi}$ and the median $\hat{\chi}$ over biconnected (s,t) pairs.

Graph	Max exposure $\hat{\chi}$ [%]					Median exposure $\hat{\chi}$ [%]				
	R	NB	LRV	NC	HS	R	NB	LRV	NC	HS
NSFNET	81.0	78.5	76.2	72.1	70.1	60.5	53.8	53.4	52.0	52.0
GÉANT	99.2	96.4	96.1	93.7	92.6	89.7	83.5	81.1	80.4	77.2
Generated (99)	98.5	97.0	96.4	95.2	93.3	84.8	78.5	77.2	76.7	72.8

As reported in Tables 3 and 4, HS attains the lowest worst-case exposure among the evaluated variants, with NC generally close behind despite its stronger global-node knowledge requirement. If we were to assume $\chi \leq 77\%$ when fixing *random flow* (HS) fragment count, we would materially overestimate $g^{\star}_{\alpha}(M,\chi)$ and the resulting extracted yield for about half of biconnected GÉANT pairs. This is further illustrated in Figure 8; it shows the fraction of biconnected (s,t) pairs whose true HS exposure lies below an assumed χ threshold on the x-axis. If we assume a lower $\chi_{s,t}$ than the true one, we do not merely misestimate *a few* pairs, but overstate the guaranteed number of good fragments and extracted final keys for a substantial fraction of pairs, as suggested by the sigmoidal shape in Figure 8.

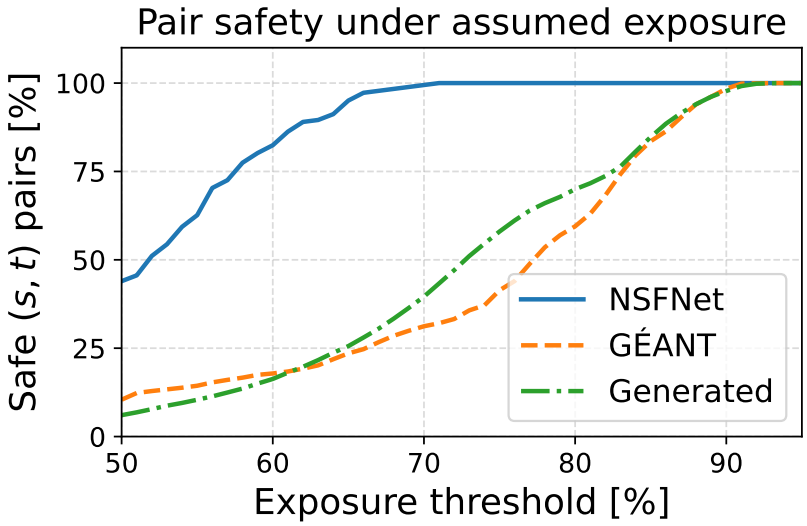

Figure 8. Fraction of biconnected (s,t) pairs whose $\chi_{s,t}$ (HS) is below an assumed threshold.

To evaluate whether worst-case exposure $\max_{s,t} \chi_{s,t}$ grows with network size, we measured it on successive snapshots of the synthetic graph and report the results in Figure 9. Although the worst-case exposure remains very high throughout, we do not observe a clear monotonic increase with $|V|$. Moreover, for the non-backtracking variants, the **maximum is not attained at the final 99-node snapshot**: NB peaks at 97.3% on the 69-node snapshot, while LRV, NC, and HS peak at 97.4%, 96.5%, and 96.1%, respectively, on the 78-node snapshot. This suggests that, at least for our generated graph family, worst-case exposure is driven more by structural bottlenecks than by network size alone.

These results also highlight a negative finding: worst-case exposure remains high even for HS, at 92.6% on GÉANT and 93.3% on the final 99-node synthetic graph, with a peak of 96.1% on the 78-node snapshot. HS helps, but the trusted-relay problem remains unsolved in this model. Among the evaluated variants, HS gives the best worst-case exposure. NC remains close, but it assumes the globally known node set, whereas HS preserves the stronger locality of using only neighbor information plus token state.

 doi:10.20944/preprints202604.0533.v1

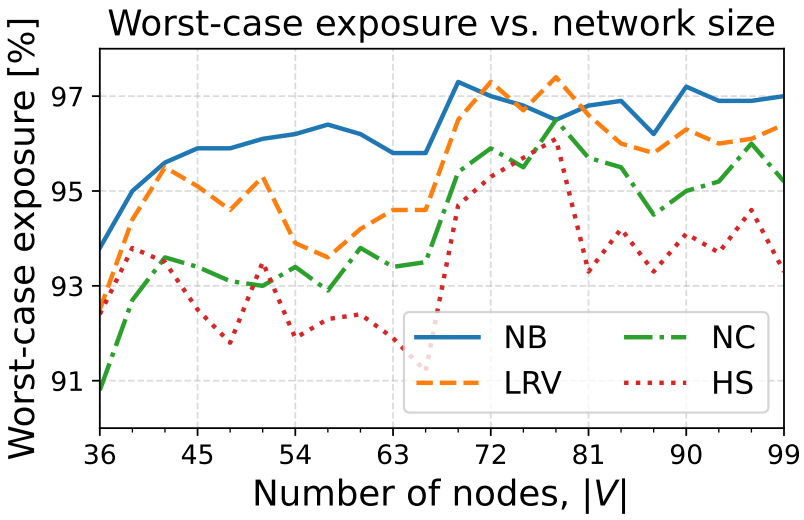

Figure 9. Worst-case exposure on synthetic graph snapshots as the network size grows.

4.3. Expected Hops and Efficiency

As shown in Figures 10 and 11, the average expected hop count over all (s, t) pairs grows linearly with network size.

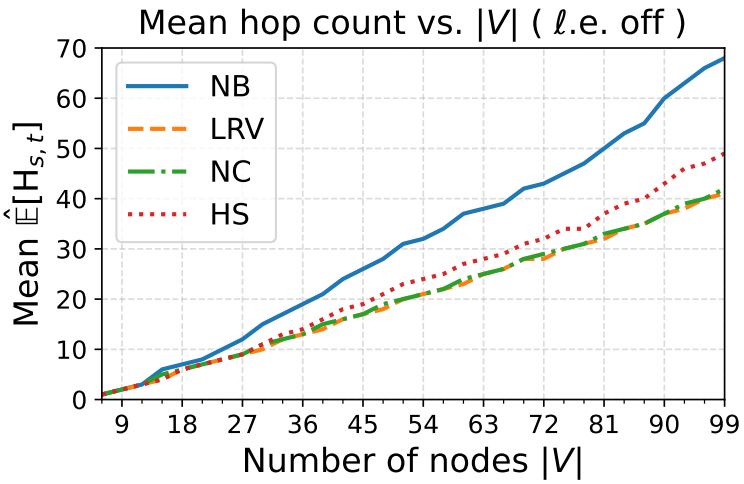

Figure 10. Average expected hop count by graph and RW variant without loop-erasure.

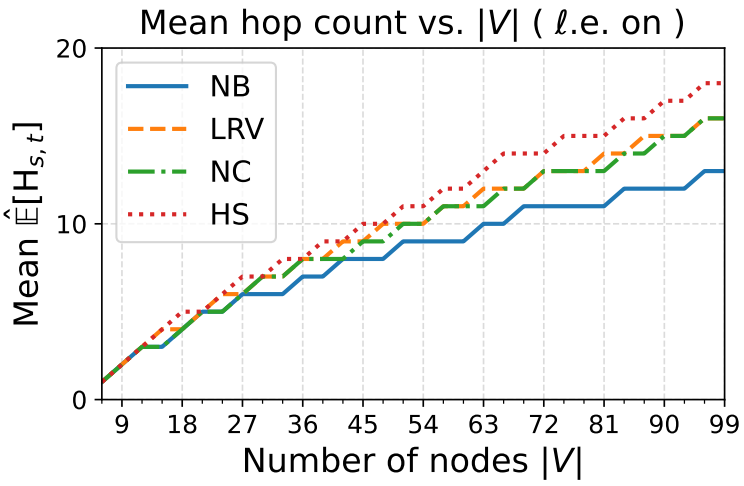

Figure 11. Average expected hop count by graph and RW variant with loop-erasure.

 doi:10.20944/preprints202604.0533.v1

To put these numbers into context, they remain well above the corresponding average shortest-path lengths from Table 2. For the generated graph, for example, the average shortest-path length over 9702 ordered pairs is 4.8 (as seen in Table 2), whereas the average mean hop count is 10–18 with loop-erasure and 41–128 without it. The average length of the shortest path when original shortest path nodes are removed is 7.5, still significantly below 10–18 post-loop-erasure.

The longest mean path after loop-erasure is attained by HS. This aligns with its original path-diversification motivation.

From Section 3.6, we factor pairwise efficiency as

$$\eta_{s,t}^{\mathsf{est}} = \rho_{s,t}^{\mathsf{est}} \frac{1}{h_{s,t}}, \qquad \rho_{s,t}^{\mathsf{est}} = \frac{g_\alpha^\star(M, \chi_{s,t})}{M}.$$

We further distinguish *pair-specific estimated* efficiency, where the exact pair-specific exposure $\chi_{s,t}$ is used inside the heuristic estimator, from *assumed* efficiency, where the implementation fixes a conservative variant-specific design exposure taken from Section 4.2. Concretely, we use $\chi_{\mathrm{NB}} = 97.3\%$, $\chi_{\mathrm{LRV}} = 97.4\%$, $\chi_{\mathrm{NC}} = 96.5\%$, and $\chi_{\mathrm{HS}} = 96.1\%$, corresponding to the peak worst-case exposures reported there. Hence,

$$\eta_{s,t}^{\mathsf{assm}} = \rho^{\mathsf{assm}} \frac{1}{h_{s,t}}, \qquad \rho^{\mathsf{assm}} = \frac{g_\alpha^\star(1024, \chi_{\mathrm{variant}})}{1024}.$$

In the experiments below, we report topology-wide averages $\overline{\eta}^{\mathsf{assm}}$ and $\overline{\eta}^{\mathsf{est}}$ over all biconnected ordered pairs. Since exposure is still computed from the original walk, loop-erasure affects only the denominator through the expected hop count.

In Tables 6 and 7, we omit the simple random walk R and focus on NB, LRV, NC, and HS, since R is already dominated on the evaluated graphs in both exposure and expected hop count.

Table 6 shows that the conservative fixed-exposure assumption substantially understates achievable efficiency on all evaluated topologies: $\overline{\eta}^{\mathsf{est}}$ is consistently far above $\overline{\eta}^{\mathsf{assm}}$. The efficiency gain is, however, noticeably smaller than the corresponding reduction in expected hop count from Table 5. This is not a contradiction. Although loop-erasure reduces average hop count dramatically, the gain in average efficiency is more modest because efficiency scales with the reciprocal hop count. Hence loop-erasure primarily removes a heavy tail of excessively long walks, rather than uniformly shortening all source–target pairs. In other words, it has a lower impact on $\frac{1}{|\mathcal{P}|}\sum_{(s,t)\in\mathcal{P}} h_{s,t}^{-1}$ than on $\frac{1}{|\mathcal{P}|}\sum_{(s,t)\in\mathcal{P}} h_{s,t}$.

Table 5. Estimated expected hop count by graph and RW variant. Let $H_{s,t}$ be the hop count in the sampled walk from s to t. We report the average of mean, median, and maximum of $H_{s,t}$ over all biconnected (s,t) pairs, and distinguish whether loop-erasure is toggled on.

Graph	Metric	Without loop-erasure					With loop-erasure				
		R	NB	LRV	NC	HS	R	NB	LRV	NC	HS
NSFNET	Mean	6	4	3	4	3	3	3	3	3	3
	Median	4	3	3	3	3	2	3	3	3	3
	Max	41	22	9	11	10	7	7	7	7	7
GÉANT	Mean	64	32	18	18	25	6	7	8	8	8
	Median	41	21	15	15	16	5	6	7	7	7
	Max	535	234	70	82	236	17	21	23	23	23
Generated	Mean	128	68	41	42	49	10	13	16	16	18
	Median	85	46	34	35	36	9	11	14	14	16
	Max	1027	516	163	181	357	33	42	54	53	58

 doi:10.20944/preprints202604.0533.v1

More importantly, the near-equality of $\overline{\eta}^{\text{est}}$ across RW variants should *not* be interpreted as a simple cancellation in which lower exposure is paid for by systematically longer walks. The covariance decomposition

$$\overline{\eta}^{\text{est}} = \overline{\rho}^{\text{est}}\,\overline{h^{-1}} + \text{Cov}\big(\rho^{\text{est}}, h^{-1}\big)$$

shows instead that pair-specific security and routing cost are strongly aligned: pairs with lower exposure also tend to have shorter walks, while difficult pairs suffer from both higher exposure and longer routes. Consequently, the covariance term is large and positive.

This alignment is already visible in the completed simulations. On NSFNET, $\text{Corr}(\rho^{\text{est}}, h^{-1})$ is approximately 0.97 for all evaluated variants, both with and without loop-erasure. On GÉANT, it remains similarly high, around 0.94–0.96. For example, on NSFNET-NB without loop-erasure, $\overline{\rho}^{\text{est}}\,\overline{h^{-1}} \approx 0.198$ and $\text{Cov}(\rho^{\text{est}}, h^{-1}) \approx 0.091$, yielding $\overline{\eta}^{\text{est}} \approx 0.289$. On GÉANT-NB without loop-erasure, the product term is only about 0.028, whereas the covariance term contributes about 0.066, again showing that the pairwise alignment is substantial.

Therefore, the similarity of $\overline{\eta}^{\text{est}}$ across variants is better understood as follows: on these topologies, the evaluated variants induce similar averages of ρ^{est}, h^{-1}, and their covariance, rather than a clean trade-off between improved exposure and worse path length.

Finally, the distribution of pairwise efficiencies is strongly right-skewed, especially on GÉANT. For instance, for NB without loop-erasure on GÉANT, the mean pair-specific estimated efficiency is 9.4% whereas the median is only 0.40%; with loop-erasure, the corresponding values are 11.0% and 1.56%. Thus, the mean efficiency should be read as a topology-wide average, not as a typical per-pair value.

Table 6. Topology-wide mean RW efficiency by graph and variant. For each graph and variant, we report separate rows for $\overline{\eta}^{\text{assm}}$ and $\overline{\eta}^{\text{est}}$, expressed in percent. The left block uses hop counts from the original walk, while the right block uses hop counts after loop-erasure. In both blocks, the exposure term is computed from the original walk and $M = 1024$.

		Mean η *without* loop-erasure [%]				Mean η *with* loop-erasure [%]			
Graph	Metric	NB	LRV	NC	HS	NB	LRV	NC	HS
NSFNET	Assm	0.41	0.39	0.61	0.74	0.45	0.40	0.64	0.75
	Est.	29	30	30	30	30	30	31	30
GÉANT	Assm	0.14	0.14	0.23	0.25	0.24	0.20	0.32	0.36
	Est.	9.4	9.8	9.8	9.7	11	11	11	11
Generated	Assm	0.05	0.06	0.09	0.10	0.12	0.09	0.15	0.16
	Est.	3.5	3.7	3.7	3.7	4.6	4.4	4.4	4.4

Table 7. Mean throughput by graph and RW variant. The top block reports raw fragment throughput $T_{s,t}$ in kbit/s. The bottom block reports mean model-based extracted throughput $R^{\text{est}}_{s,t} = T_{s,t}\rho^{\text{est}}_{s,t}$ in kbit/s under the simplified entropy model.

Metric	Graph	NB	LRV	NC	HS
$T_{s,t}$	NSFNET	2.03	2.02	2.05	2.03
	GÉANT	1.70	1.70	1.71	1.77
	Generated	1.97	1.97	1.97	2.01
$R^{\text{est}}_{s,t}$	NSFNET	0.93	0.95	0.98	0.98
	GÉANT	0.34	0.37	0.38	0.43
	Generated	0.41	0.43	0.44	0.49

4.4. Throughput

To evaluate *throughput* for a fixed (s, t) pair, we implement the random-walk key-relaying scheme as a discrete-event simulation driven by a min-heap of timestamped events. To remain consistent with the two-phase protocol model, route discovery is treated as a preceding classical control phase

 doi:10.20944/preprints202604.0533.v1

that determines the realized payload route (with loop-erasure enabled, this route is $\mathrm{LE}(W)$). The throughput metric $T_{s,t}$ counts only delivered payload fragment bits; scouting/control traffic is omitted from $T_{s,t}$ because it uses the classical channel and consumes no QKD-derived OTP key material. Each hop of a packet from node u to node v is split into: (i) *OTP key reservation* on link (u, v), which may incur waiting time computed from the link secret-key rate (SKR) and current key balance; (ii) a fixed-latency classical channel; and (iii) *receiver admission* into a finite relay buffer with FIFO backpressure. The simulator repeatedly pops the earliest event, advances the simulation clock, updates link/node state, and schedules successor events.

Each undirected link $e \in E$ generates QKD key material at a constant secret-key rate $g_e = 1\,\mathrm{kbit/s}$ for all links. Key material is consumed in 256 bit units. Each hop spends 256 bit for one-time-pad (OTP) encryption on the traversed QKD link. We use a default inter-node classical latency of 5 ms.

Simulations start with empty link buffers. We do not discard a warm-up period when calculating throughput; given a fixed simulation duration of 1000 s, its impact is unlikely to be large. We ignore classical network throughput (goodput) and protocol overhead, as QKD is the dominant bottleneck (e.g., 1 kbit/s vs. typical 1 Gbit/s classical links).

Table 7 summarizes both the mean raw fragment throughput and the mean model-based extracted throughput across graphs and RW variants.

Raw fragment throughput varies only modestly across variants, but extracted throughput differs more because it inherits the pair-specific estimated yield ρ^{est}. On NSFNET, mean raw throughput is about 2.0 kbit/s for all evaluated variants, whereas mean model-based extracted throughput ranges from roughly 0.93 to 0.98 kbit/s. On GÉANT, mean raw throughput ranges from about 1.70 to 1.77 kbit/s, while mean model-based extracted throughput rises from about 0.34 kbit/s for NB to about 0.43 kbit/s for HS. On the Generated graph, mean raw throughput again stays nearly flat across variants, at about 1.97–2.01 kbit/s, but mean model-based extracted throughput increases from about 0.41 kbit/s for NB to about 0.49 kbit/s for HS. Under the conservative assumed exposure, the corresponding extracted-throughput proxy would collapse by roughly one to two orders of magnitude, mirroring the behavior already seen in Table 6.

5. Conclusions

This paper studies topology-oblivious QKD key relaying under an explicit and restricted threat model: honest endpoints, at most one compromised relay, and biconnected source–target pairs. In this regime, local stochastic forwarding can be a usable decentralized heuristic within the simplified model, but the forwarding rule has a clear effect on both exposure and routing cost.

Among the evaluated topology-oblivious random walk variants, HS achieved the lowest worst-case exposure. On GÉANT, the maximum exposure over biconnected ordered pairs decreased from 96.1% for LRV to 92.6% for HS. On the 99-node synthetic graph, the corresponding maxima were 97.0% (NB), 96.4% (LRV), 95.2% (NC), and 93.3% (HS). Thus, HS provided the best diversification among the evaluated topology-oblivious methods, while avoiding the global node-set knowledge required by NC.

The second main result concerns the effect of scouting-based loop erasure on routing cost. It sharply reduced expected hop counts on the larger graphs. On GÉANT, the topology-wide mean hop count fell from 32 to 7 for NB, from 18 to 8 for LRV, from 18 to 8 for NC, and from 25 to 8 for HS. On the 99-node synthetic graph, the corresponding reductions were from 68 to 13, from 41 to 16, from 42 to 16, and from 49 to 18. Therefore, loop erasure materially lowers QKD key consumption, reduces queueing pressure, and eliminates self-induced cyclic waiting in the model.

Under the paper's simplified entropy and throughput model, these routing improvements also translate into higher model-based extracted throughput. On GÉANT, mean extracted throughput increased from 0.34 kbit/s for NB to 0.37 kbit/s for LRV, 0.38 kbit/s for NC, and 0.43 kbit/s for HS. On the synthetic graph, the same ordering was observed: 0.41, 0.43, 0.44, and 0.49 kbit/s, respectively.

 doi:10.20944/preprints202604.0533.v1

On NSFNET, the variants were closer, ranging from 0.93 to 0.98 kbit/s. Under conservative assumed exposure, however, this throughput proxy can collapse by roughly one to two orders of magnitude.

A further empirical finding is that pair-specific estimated yield and routing efficiency are positively aligned rather than exhibiting a simple trade-off: pairs with lower exposure also tend to have shorter walks. This effect is strong, with $\mathrm{Corr}(\rho_{s,t}^{\mathsf{est}}, h_{s,t}^{-1})$ around 0.97 on NSFNET and about 0.94–0.96 on GÉANT.

Random flow does reduce the end-to-end rate obtained per consumed link-key bit, because one delivered fragment traverses multiple hops. Relative to shortest-path forwarding, the post-loop-erasure average hop-count overhead is about 1.5× for NB and 1.7× for LRV/NC/HS on GÉANT, and about 2.7×, 3.3×, 3.3×, and 3.8× on the synthetic graph. Without loop erasure, this overhead is substantially worse.

Overall, the paper identifies a usable middle ground between centralized, topology-aware routing and idealized disjoint-path constructions. Its main value is as a simple decentralized alternative to maintaining global routing state: HS improves local diversification without global path computation, and scouting-based loop erasure makes the resulting random-flow construction substantially cheaper in hop count. At the same time, worst-case exposure remains high even for the better variants: 92.6% on GÉANT, 93.3% on the final synthetic graph, and up to 96.1% on a synthetic snapshot. Under the terminology used here, topology-oblivious excludes global adjacency knowledge, link-state exchange, and path computation, but may still allow the node-identifier universe to be known in advance, as in NC; under a one-compromised-relay model, stochastic forwarding is therefore best read as a decentralized heuristic, not as a solution to the trusted-relay problem.

Future work should tighten the entropy analysis, extend the threat model beyond a single malicious relay, and study more realistic traffic regimes. Two important next steps are proactive pairwise buffer maintenance and relaxing the stateless-relay assumption to test whether limited node-local state can further improve diversification, routing efficiency, and congestion behavior.

Author Contributions: K.P. and S.K. developed the research concept and methodology. K.P. implemented the software, conducted the main evaluation, performed the analysis, and prepared the original draft. R.I. contributed to numerical evaluation and validation. J.V. provided supervision and project leadership. E.K., E.C., L.L., and E.R. contributed through discussion, review, and support within the broader project. All authors have read and agreed to the published version of the manuscript.

Funding: This work has been supported by the Latvian Quantum Initiative under European Union Recovery and Resilience Facility project no. 2.3.1.1.i.0/1/22/I/CFLA/001.

Data Availability Statement: The graph edge lists and experiment source code supporting the findings of this study are publicly available at https://github.com/LUMII-Syslab/random-walk-key-relaying.

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

References

1. Dervisevic, E.; Tankovic, A.; Fazel, E.; Kompella, R.; Fazio, P.; Voznak, M.; Mehic, M. Quantum Key Distribution Networks – Key Management: A Survey. *ACM Computing Surveys* **2025**, *57*, 257:1–257:39. https://doi.org/10.1145/3730575.
2. Moy, J. OSPF Version 2. RFC 2328, 1998. https://doi.org/10.17487/RFC2328.
3. Liao, S.K.; Cai, W.Q.; Liu, W.Y.; Zhang, L.; Li, Y.; Ren, J.G.; Yin, J.; Shen, Q.; Cao, Y.; Li, Z.P.; et al. Satellite-to-ground quantum key distribution. *Nature* **2017**, *549*, 43–47. https://doi.org/10.1038/nature23655.
4. Huang, D.; Huang, P.; Lin, D.; Zeng, G. Long-distance continuous-variable quantum key distribution by controlling excess noise. *Scientific Reports* **2016**, *6*, 19201. https://doi.org/10.1038/srep19201.
5. Elliott, C. Building the quantum network*. *New Journal of Physics* **2002**, *4*, 46–46. https://doi.org/10.1088/1367-2630/4/1/346.
6. ETSI Industry Specification Group (ISG) Quantum Key Distribution. ETSI GS QKD 014 V1.1.1: Quantum key distribution (QKD); protocol and data format of REST-based key delivery API, 2019. tex.howpublished: Group Specification GS QKD 014 V1.1.1.

 doi:10.20944/preprints202604.0533.v1

7. Beals, T.R.; Sanders, B.C. Distributed Relay Protocol for Probabilistic Information-Theoretic Security in a Randomly-Compromised Network **2008**. Version Number: 2, https://doi.org/10.48550/ARXIV.0803.2919.
8. Wen, H.; Han, Z.; Zhao, Y.; Guo, G.; Hong, P. Multiple stochastic paths scheme on partially-trusted relay quantum key distribution network. *Science in China Series F: Information Sciences* **2009**, *52*, 18–22. https://doi.org/10.1007/s11432-009-0001-4.
9. Kumar, P.; Kundu, N.K.; Kar, B. Quantum Key Distribution Routing Protocol in Quantum Networks: Overview and Challenges. *arXiv preprint arXiv:2407.13156* **2024**. https://doi.org/10.48550/arXiv.2407.13156.
10. Yao, J.; Wang, Y.; Li, Q.; Mao, H.; El-Latif, A.A.A.; Chen, N. An Efficient Routing Protocol for Quantum Key Distribution Networks. *Entropy* **2022**, *24*, 911. https://doi.org/10.3390/e24070911.
11. Drif, Y.; Bedhief, I.; Chatzinotas, S. Distributed Key Relay: OSPF for Effective QKD. *IEEE Communications Standards Magazine* **2026**, *10*, 154–161. https://doi.org/10.1109/MCOMSTD.2025.3572642.
12. Álvarez Roa, M.; Stan, C.; Verschoor, S.; Tafur Monroy, I.; Rommel, S. Decentralized Key Distribution versus On-Demand Relaying for QKD Networks. *Journal of Optical Communications and Networking* **2025**, *17*, 732–742. https://doi.org/10.1364/JOCN.547793.
13. Kiktenko, E.O.; Tayduganov, A.; Fedorov, A.K. Routing Algorithm Within the Multiple Non-Overlapping Paths' Approach for Quantum Key Distribution Networks. *Entropy* **2024**, *26*, 1102. https://doi.org/10.3390/e26121102.
14. Wang, M.; Li, J.; Xue, K.; Li, R.; Yu, N.; Li, Y.; Liu, Y.; Sun, Q.; Lu, J. A Segment-Based Multipath Distribution Method in Partially-Trusted Relay Quantum Networks. *IEEE Communications Magazine* **2023**, *61*, 184–190. https://doi.org/10.1109/MCOM.010.2200672.
15. Le, Q.C.; Bellot, P.; Demaille, A. Towards the World-Wide Quantum Network, 2008. Paper on partially compromised QKD networks and stochastic routing; introduces ADRA and discusses its experimental evaluation.
16. Le, Q.C.; Bellot, P.; Demaille, A. Stochastic Routing in Large Grid Shaped Quantum Networks. In Proceedings of the Proceedings of the 5th International Conference on Computer Sciences, Research, Innovation and Vision for the Future, Hanoi, Vietnam, 2007; pp. 166–174.
17. Ghourab, E.M.; Azab, M.; Gračanin, D. A Quantum Key Distribution Routing Scheme for a Zero-Trust QKD Network System: A Moving Target Defense Approach. *Big Data and Cognitive Computing* **2025**, *9*, 76. https://doi.org/10.3390/bdcc9040076.
18. Lydersen, L.; Wiechers, C.; Wittmann, C.; Elser, D.; Skaar, J.; Makarov, V. Hacking commercial quantum cryptography systems by tailored bright illumination. *Nature Photonics* **2010**, *4*, 686–689. https://doi.org/10.1038/nphoton.2010.214.
19. Alon, N.; Benjamini, I.; Lubetzky, E.; Sodin, S. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics* **2007**, *09*, 585–603. https://doi.org/10.1142/S0219199707002551.
20. Akash, A.K.; Fekete, S.; Lee, S.K.; López-Ortiz, A.; Maftuleac, D.; McLurkin, J. Lower Bounds for Graph Exploration Using Local Policies. *Journal of Graph Algorithms and Applications* **2017**, *21*, 371–387. https://doi.org/10.7155/jgaa.00421.
21. Renner, R.; Wolf, S. Simple and Tight Bounds for Information Reconciliation and Privacy Amplification. In *Advances in Cryptology - ASIACRYPT 2005*; Hutchison, D.; Kanade, T.; Kittler, J.; Kleinberg, J.M.; Mattern, F.; Mitchell, J.C.; Naor, M.; Nierstrasz, O.; Pandu Rangan, C.; Steffen, B.; et al., Eds.; Springer Berlin Heidelberg: Berlin, Heidelberg, 2005; Vol. 3788, pp. 199–216. Series Title: Lecture Notes in Computer Science, https://doi.org/10.1007/11593447_11.
22. Tomamichel, M.; Schaffner, C.; Smith, A.; Renner, R. Leftover Hashing Against Quantum Side Information. *IEEE Transactions on Information Theory* **2011**, *57*, 5524–5535. https://doi.org/10.1109/TIT.2011.2158473.
23. Stinson, D.R.; Paterson, M. *Cryptography: Theory and Practice*, 4 ed.; Chapman and Hall/CRC, 2018. https://doi.org/10.1201/9781315282497.
24. Ross, S.M. *Introduction to Probability Models*, 12 ed.; Academic Press, 2019.
25. Lawler, G.F.; Limic, V. *Random Walk: A Modern Introduction*; Vol. 123, *Cambridge Studies in Advanced Mathematics*, Cambridge University Press, 2010. https://doi.org/10.1017/CBO9780511750854.
26. Wilson, D.B. Generating Random Spanning Trees More Quickly than the Cover Time. In Proceedings of the Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. ACM, 1996, pp. 296–303. https://doi.org/10.1145/237814.237880.
27. Fung, C.H.F.; Ma, X.; Chau, H.F. Practical Issues in Quantum-Key-Distribution Postprocessing. *Physical Review A* **2010**, *81*, 012318. https://doi.org/10.1103/PhysRevA.81.012318.

 doi:10.20944/preprints202604.0533.v1

28. Mehic, M.; Niemiec, M.; Rass, S.; Ma, J.; Peev, M.; Aguado, A.; Martin, V.; Schauer, S.; Poppe, A.; Pacher, C.; et al. Quantum Key Distribution: A Networking Perspective. *ACM Computing Surveys* **2021**, *53*, 1–41. https://doi.org/10.1145/3402192.
29. Peev, M.; Pacher, C.; Alléaume, R.; Barreiro, C.; Bouda, J.; Boxleitner, W.; Debuisschert, T.; Diamanti, E.; Dianati, M.; Dynes, J.F.; et al. The SECOQC quantum key distribution network in Vienna. *New Journal of Physics* **2009**, *11*, 075001. https://doi.org/10.1088/1367-2630/11/7/075001.
30. Mukherjee, B.; Ramamurthy, S.; Banerjee, D.; Mukherjee, A. Some principles for designing a wide-area optical network. In Proceedings of the Proceedings of INFOCOM '94 Conference on Computer Communications, Toronto, Ont., Canada, 1994; pp. 110–119. https://doi.org/10.1109/INFCOM.1994.337626.
31. GÉANT. GN4-3N. Available online: https://network.geant.org/gn4-3n/ (accessed on 2026-02-04)., 2023. Web page describing the GN4-3N project (2019–2023); topology diagram consulted to reconstruct node/edge lists (links > 1000 km removed in our processed graph).
32. Xuereb, A. Towards an ultra-secure communication network for the EU. GÉANT CONNECT. Available online: https://connect.geant.org/2023/12/13/towards-an-ultra-secure-communication-network-for-the-eu (accessed on 2026-02-04)., 2023. Online article discussing EuroQCI-oriented ultra-secure networking efforts including QKD.