Measuring Quantum Key Distribution (QKD) Dependability via Wide Area Network Fault Injection

Audrey Feng
[audreyf@illinois.edu](mailto audreyf@illinois.edu)
National Center for Supercomputing Applications,
University of Illinois Urbana-Champaign
Urbana, Illinois, USA

Phuong Cao (Advisor)
[pcao3@illinois.edu](mailto pcao3@illinois.edu)
National Center for Supercomputing Applications,
University of Illinois Urbana-Champaign
Urbana, Illinois, USA

1 Background

Hybrid High-Performance Computing and Quantum Computing (HPC-QC) systems increasingly rely on quantum networks for secure communication, yet their dependability under real-world conditions has not been systematically studied. Government mandates for the transition to post-quantum cryptography (PQC) are accelerating deployment of QKD and PQC systems. Existing approaches use quantum network simulators such as SeQUeNCe [1] and NetSquid [2], which assume ideal classical channels and fail to account for interactions between cross-layer faults and the effects of real-world network faults on protocol behavior.

2 Methodology

This work presents a systematic fault injection study [3] of BB84 [4] QKD using QFabric [5], a quantum-network emulation platform built on the FABRIC testbed [6]. Unlike pure simulators such as SeQUeNCe and NetSquid, QFabric emulates the quantum channel using a P4/BMv2 programmable switch while routing classical sifting traffic over a real FABRIC WAN link. This allows genuine network conditions to affect protocol behavior and enables cross-layer fault injection.

We injected classical packet loss simultaneously with polarization fidelity degradation, and separately with detector efficiency degradation, on real FABRIC WAN infrastructure.

3 Results

We swept polarization fidelity $p_f \in \{0.8, 0.9, 1.0\}$ against packet loss $\in \{0, 1, 5, 15\}$% (target n=10 per cell; $n = 111$ usable runs total). A two-way ANOVA found that polarization fidelity overwhelmingly determines QBER ($F(2, 99) = 699.6$, $p = 3.9 \times 10^{-59}$; mean QBER 9.99% at $p_f = 0.8$, 5.18% at $p_f = 0.9$, 0% at $p_f = 1.0$), while packet loss ($p = 0.29$) and the fidelity×loss interaction ($p = 0.35$) were both non-significant. This matches the expected mechanism: BB84's classical sifting traffic rides over TCP, so packet loss is retransmitted rather than corrupting bits, while quantum-channel fidelity degradation directly injects bit errors.

Although the *mean* QBER at $p_f = 0.8$ (9.3–10.6% across loss levels) stays below the 11% failure threshold [7], individual runs do not: 8 of 31 runs (25.8%) at $p_f = 0.8$ exceeded 11% from ordinary run-to-run variance alone, with zero crossings at $p_f = 0.9$ or $p_f = 1.0$. This rate is not explained by loss level ($\chi^2 = 2.45$, $p = 0.48$ across loss conditions): crossings depend on how close a fidelity setting's mean QBER sits to the failure threshold, not on classical network conditions. This indicates that aggregate statistics can understate real deployment risk: a fidelity setting with a "safe" mean QBER can still fail a meaningful fraction of individual sessions.

In contrast, substituting detector efficiency $\eta \in \{0.5, 0.7, 0.85, 1.0\}$ for polarization fidelity (holding $p_f = 0.8$ fixed) showed no significant interaction (QBER: $p = 0.74$; key rate: $p = 0.55$) or main effects at $\alpha = 0.05$. This is consistent with efficiency acting purely as a photon-loss fault rather than a state-corruption fault: it degrades key rate independently of QBER and does not interact with classical packet loss the way polarization fidelity's QBER contribution does.

4 Related Work

Existing QKD testbeds and simulators each capture only part of the dependability picture this work addresses. Berlin OpenQKD [8] uses real quantum hardware and a real classical WAN, but its high hardware cost limits the scale of systematic fault-injection studies. StarBED [9] supports classical-network fault injection at moderate cost but has no real quantum channel, so it cannot study quantum-layer faults or their interaction with classical ones. Pure simulators such as SeQUeNCe [1] and NetSquid [2] support neither a real classical WAN link nor cross-layer fault injection, since both assume an idealized classical channel by design.

To our knowledge, QFabric is the first platform combining a real classical WAN link, quantum-channel emulation, and simultaneous cross-layer (classical, quantum, and detector) fault injection, at a fraction of the cost of physical QKD hardware.

5 Limitations

Our polarization-fidelity × packet-loss design is unbalanced: $p_f = 0.9$ and $p_f = 1.0$ reached a clean $n = 10$ per cell, while $p_f = 0.8$ retains gaps (n=6–10 across loss levels) from legacy missing data and pre-deployment infrastructure failures predating this campaign, which we retain as missing data rather than excluding post hoc. We deliberately selected a single FABRIC link pair to isolate the fault-domain mechanism from confounding network variability; the vulnerability we characterize is a property of the distributed protocol's endpoint computation rather than an artifact of this specific link's noise profile, but we have not validated this independence across multiple link pairs or geographic distances. The absence of a significant mean-level interaction should not be read as evidence that no cross-layer interaction exists in QKD generally; our conclusion is specific to the fault types, parameter ranges, and FABRIC link conditions tested here.

6 Conclusion and Future Work

Our results show that polarization fidelity, not classical packet loss, is the dominant driver of QBER at the conditions tested, with no significant mean-level interaction between classical and quantum-layer faults. However, individual-run variance at marginal fidelity

produces real failure-threshold crossings invisible to mean-based analysis alone, indicating that QKD dependability assessments should examine run-level distributions, not only averages. Future work includes testing a finer-grained fidelity sweep near the threshold boundary and additional classical fault types (jitter, asymmetric latency) to determine where this tail risk grows most sharply.

References

[1] Xiaoliang Wu, Alexander Kolar, et al. "SeQUeNCe: a customizable discrete-event simulator of quantum networks". In: *Quantum Science and Technology* 6.4 (2021). ISSN: 2058-9565. DOI: 10.1088/2058-9565/ac22f6.

[2] Tim Coopmans et al. "NetSquid, a NETwork Simulator for QUantum Information using Discrete events". In: *Communications Physics* 4.1 (2021). DOI: 10.1038/s42005-021-00647-8.

[3] Mei-Chen Hsueh, T.K. Tsai, and R.K. Iyer. "Fault injection techniques and tools". In: *Computer* 30.4 (1997), pp. 75–82. DOI: 10.1109/2.585157.

[4] Charles H. Bennett and Gilles Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Theoretical Computer Science* 560 (Dec. 2014), pp. 7–11. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2014.05.025.

[5] Komal Thareja. *QFabric: Quantum Network Emulation on FABRIC (BB84 QKD over P4)*. FABRIC Artifact Manager. 2026.

[6] Ilya Baldin, Anita Nikolich, et al. "FABRIC: A national-scale programmable experimental network infrastructure". In: *IEEE Internet Computing* 23.6 (2019), pp. 38–47.

[7] Peter W. Shor and John Preskill. "Simple Proof of Security of the BB84 Quantum Key Distribution Protocol". In: *Physical Review Letters* 85.2 (2000), pp. 441–444. DOI: 10.1103/PhysRevLett.85.441.

[8] Ralf-Peter Braun et al. "Berlin OpenQKD Testbed Evaluating Quantum Key Distribution in Provider Networks". In: *Proceedings of the 2023 8th International Conference on Systems, Control and Communications (ICSCC '23)*. ACM, 2024, pp. 41–51. DOI: 10.1145/3634865.3634872.

[9] Toshiyuki Miyachi, Ken-ichi Chinen, and Yoichi Shinoda. "StarBED and SpringOS: Large-Scale General Purpose Network Testbed and Supporting Software". In: *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS '06)*. Pisa, Italy: ACM, 2006. DOI: 10.1145/1190095.1190133.