Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Active Exploitation of TID Environmental Routing for Key Recovery in Microwave Quantum Channels

victoria mellor

emmann242@gmail.com

University of Portsmouth

Research Article

Keywords: quantum key distribution, Temporal Information Dynamics, environmental decoherence routing, information recording capacity, microwave quantum networks, Helstrom bound, Devetak-Winter key rate, privacy amplification, superconducting quantum communication

Posted Date: March 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9076817/v1

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

Additional Declarations: No competing interests reported.

Active Exploitation of TID Environmental Routing for Key Recovery in Microwave Quantum Channels

Victoria Mellor[1]

[1]University of Portsmouth, Portsmouth, United Kingdom
Victoria.Mellor@myport.ac.uk

March 2026

Abstract

This research derives a complete information-theoretic attack chain in which an adversary actively engineers environmental decoherence modes to exploit the capacity-weighted routing proposed by the Temporal Information Dynamics (TID) framework for key recovery from a quantum channel. The adversary places a cryogenic high-Q resonator (thermal occupation $\bar{n} \approx 0$) adjacent to a warm microwave quantum channel ($\bar{n} \gg 1$), creating a capacity gradient that, under TID, redirects information from lost photons preferentially toward the adversary's mode. I compute the per-bit environmental states, verify optimality of the adversary's measurement via the Helstrom bound, and derive the mutual information between key bits and adversarial measurements, yielding a Devetak-Winter key rate analysis.

A critical finding emerges: the attack is effective only when the quantum bit error rate (QBER) is low. At QBER = 0.5% with 10 dB channel loss, the adversary extracts 0.328 bits of mutual information per sifted key bit while standard BB84 privacy amplification removes only 0.045 bits, yielding 0.283 uncompensated bits per sifted bit. Paradoxically, higher-quality channels with lower QBER are more vulnerable because reduced privacy amplification leaves less margin against capacity-routed leakage. The crossover occurs at QBER $\approx$ 6%: below this threshold, the attack breaches security; above it, standard privacy amplification is sufficient.

A second critical finding is that this attack surface exists only for microwave quantum communication (superconducting quantum networks), not for standard optical fibre QKD, because at optical frequencies all environmental modes have $\bar{n} \approx 0$ regardless of temperature, eliminating the capacity gradient.

All results are conditional on two unproven assumptions: (A1) TID is physically valid and (A2) the decoherence routing basis is the coherent-state basis. The physical experiment required to test these assumptions is identified.

Keywords: quantum key distribution, Temporal Information Dynamics, environmental decoherence routing, information recording capacity, microwave quantum networks, Helstrom bound, Devetak-Winter key rate, privacy amplification, superconducting quantum communication

1 Introduction

The Temporal Information Dynamics (TID) framework [1] proposes that the information recording capacity of each environmental decoherence mode enters the routing probability as an independent physical variable. If correct, an adversary could exploit this by engineering high-capacity environmental modes adjacent to a quantum channel, redirecting information from channel losses preferentially toward modes they control without disturbing the transmitted signal or increasing the quantum bit error rate (QBER).

This paper derives the complete attack chain from first principles: the adversary's environmental engineering protocol, the conditional environmental states produced by each key bit, the optimal measurement for bit recovery (Helstrom bound), and the resulting secret key rate under the Devetak-Winter framework [4]. I outline the parameter regimes where the attack breaches security and where it does not. The attack is distinct from known QKD side-channel vulnerabilities surveyed in [5]: rather than exploiting implementation imperfections in sources or detectors, it targets the environmental decoherence pathway itself.

Two results are unexpected. First, the attack is most effective against high-quality channels with low QBER, because reduced privacy amplification leaves less margin to absorb capacity-routed leakage. Second, the attack surface is restricted to microwave quantum communication; standard optical fibre QKD is immune because the thermal occupation of optical-frequency modes is negligible at all relevant temperatures.

All results are conditional on TID being physically valid (Assumption A1) and the decoherence routing basis being the coherent-state basis rather than the Fock basis (Assumption A2). I also discuss the physical experiment that would test both assumptions.

1.1 Methodological Approach

This work derives the adversary's information gain analytically under the assumption that TID is correct, computes the resulting key rate, and identifies the conditions under which the attack succeeds. Rather than simulating standard dynamics (which produce zero TID signal by construction, as the standard Lindblad master equation does not contain the capacity-weighted routing term), the analysis proceeds from TID's modified routing formula to its information-theoretic consequences.

2 Channel and Adversary Model

2.1 Microwave Quantum Channel

I consider a microwave quantum communication channel operating at frequency ν connecting two superconducting quantum nodes. The channel traverses a segment at temperature T_{env} with total loss L dB. The channel transmittance is $\eta = 10^{-L/10}$ and the probability that each transmitted photon is lost to the environment is $p_{\text{lost}} = 1 - \eta$.

The thermal occupation of environmental modes at frequency ν and temperature T is:

$$\bar{n}(T) = \frac{1}{\exp(h\nu/k_B T) - 1}. \tag{1}$$

At microwave frequencies ($\nu = 5$ GHz), $\bar{n}(300\,\text{K}) \approx 1249$, while $\bar{n}(20\,\text{mK}) \approx 0$. At optical frequencies ($\nu \sim 200$ THz), $\bar{n}(300\,\text{K}) \approx 0$ because $h\nu \gg k_B T$. This frequency

dependence is the origin of the attack surface restriction to microwave channels.

2.2 Adversary Model

The adversary places a cryogenic high-Q resonator mode at a temperature $T_{\text{adv}} = 20$ mK adjacent to the warm channel segment. The resonator couples to the channel with a fraction f_{int} of the total environmental coupling. The adversary's mode has a thermal occupation of $\bar{n}_{\text{adv}} \approx 0$; the natural environmental modes have $\bar{n}_{\text{env}} \approx 1249$.

Definition 1 (TID Information Recording Capacity)**.** *For a thermal environmental mode at occupation $\bar{n}$, the TID recording capacity is [1]:*

$$C(\bar{n}) = \frac{1}{2\pi(2\bar{n}+2)}. \tag{2}$$

This formula has a direct physical interpretation grounded in bosonic channel theory. For a single bosonic mode in a thermal state with occupation $\bar{n}$, the Husimi Q-function is a Gaussian with variance proportional to $\bar{n}+1$. The capacity $C(\bar{n})$ is the inverse participation ratio of this Q-function integrated over phase space [1], which quantifies how concentrated the mode's phase-space distribution is. A cold mode ($\bar{n} \approx 0$) has a tightly concentrated Q-function, meaning that the absorption of a single photon produces a large, distinguishable displacement in phase space—a high-fidelity record. A hot mode ($\bar{n} \gg 1$) has a broadly spread Q-function, so absorbing one additional photon produces a negligible change relative to the existing thermal fluctuations—a poor record. The $1/(2\bar{n}+2)$ scaling captures this: recording quality degrades inversely with the thermal noise floor. This is the thermal-state special case of the general Gaussian capacity formula $C = 1/[2\pi\sqrt{\det(\mathbf{V}+\mathbf{V}^{\mathcal{T}}+\mathbf{I})}]$, where $\mathbf{V}$ is the covariance matrix and $\mathbf{V}^{\mathcal{T}}$ its time-reversal (see [1], Proposition 8.4).

The capacity ratio between the adversary's cryogenic mode and the warm environment is:

$$\frac{C_{\text{adv}}}{C_{\text{env}}} = \frac{2\bar{n}_{\text{env}}+2}{2\bar{n}_{\text{adv}}+2} \approx \bar{n}_{\text{env}}+1 \approx 1250 \tag{3}$$

for the microwave parameters above.

2.3 TID-Modified Routing

Under standard theory, the fraction of lost photon information routed to the adversary's mode is simply proportional to the coupling fraction:

$$f_{\text{route}}^{\text{std}} = f_{\text{int}}. \tag{4}$$

Under TID, routing is weighted by information recording capacity:

$$f_{\text{route}}^{\text{tid}} = \frac{C_{\text{adv}} \cdot f_{\text{int}}}{C_{\text{adv}} \cdot f_{\text{int}} + C_{\text{env}} \cdot (1-f_{\text{int}})}. \tag{5}$$

For $C_{\text{adv}}/C_{\text{env}} \approx 1250$ and $f_{\text{int}} = 0.01$ (1% of environmental coupling), Eq. (5) gives $f_{\text{route}}^{\text{tid}} = 0.927$: the adversary captures 92.7% of the environmental information despite controlling only 1% of the coupling. Under standard theory, they would capture only 1%.

This is the core of the attack: TID's capacity weighting acts as a massive amplifier, redirecting information toward the adversary's high-capacity mode. Fig. 1 illustrates the

dramatic contrast: under TID, even minimal coupling ($f_{\text{int}} \sim 0.01$) routes nearly all environmental information to the adversary, whereas standard routing remains proportional to the coupling fraction.

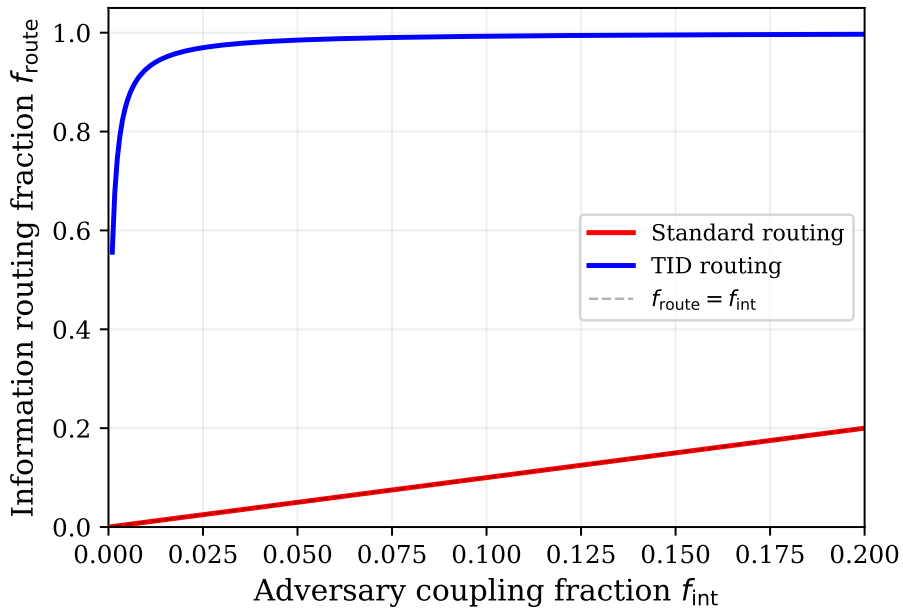


Figure 1: Information routing fraction f_{route} as a function of adversary coupling fraction f_{int} for standard (red) and TID (blue) routing at $\nu = 5$ GHz, $C_{\text{adv}}/C_{\text{env}} \approx 1250$. Under TID, even 1% coupling captures over 90% of the environmental information due to the 1250× capacity ratio between the cryogenic adversary mode and the warm environment.

2.4 Why Optical QKD Is Immune

At optical frequencies ($\nu \sim 200$ THz, corresponding to a telecom wavelength of 1550 nm), $h\nu/k_B T \approx 31$ at room temperature, giving $\bar{n}(300\,\text{K}) \approx e^{-31} \approx 0$. Both the adversary's mode and the natural environment have $\bar{n} \approx 0$, so $C_{\text{adv}} \approx C_{\text{env}}$ and no capacity gradient exists. The TID routing enhancement factor is unity: Eq. (5) reduces to Eq. (4).

Standard optical fibre QKD is therefore immune to this class of attack regardless of whether TID is physically valid. The attack surface is restricted to microwave quantum communication, including superconducting quantum networks, quantum interconnects between cryostats, and any channel where the natural environmental modes are thermally populated.

3 Per-Bit Environmental State Analysis

3.1 Z-Basis Key Bits

Consider a BB84 [2] key bit encoded in the Z basis. Alice sends either $|0\rangle$ (no photon emitted into the environment) or $|1\rangle$ (photon emitted and may be captured by the environment).

Let $p_{\text{capt}} = p_{\text{lost}} \cdot f_{\text{route}}$ be the probability that a photon emitted by a $|1\rangle$-state qubit is captured by the adversary's mode. The conditional environmental states in the adver-

sary's resonator are:

$$\rho_E^{(0)} = |0\rangle\langle 0|, \tag{6}$$

$$\rho_E^{(1)} = (1 - p_{\text{capt}})|0\rangle\langle 0| + p_{\text{capt}}|1\rangle\langle 1|. \tag{7}$$

For a vacuum resonator ($\bar{n}_{\text{adv}} = 0$), detection of a photon in the resonator unambiguously identifies the key bit as 1. The absence of a photon is ambiguous: it could mean the key bit was 0 (no emission), or it was 1 but the photon went elsewhere.

Encoding assumption. Equations 6–7 assume ideal single-photon sources encoding key bits as Fock states $|0\rangle$ and $|1\rangle$. This is appropriate for transmon-based microwave QKD, where the qubit naturally encodes in the $\{|0\rangle, |1\rangle\}$ Fock basis. For architectures employing coherent-state pulses $|\alpha\rangle$, the environmental states produced by loss are themselves coherent (with reduced amplitude) rather than the Fock-diagonal mixtures above, which would modify both the adversary's optimal measurement and the resulting mutual information. The single-photon model is used throughout; the coherent-state generalisation is left for future work.

3.2 Adversary's Optimal Measurement

The adversary performs photon number measurement on their resonator after each transmission. First, I verify that this is optimal via the Helstrom bound. With a uniform prior $P(k{=}0) = P(k{=}1) = 1/2$, the minimum error probability for distinguishing $\rho_E^{(0)}$ from $\rho_E^{(1)}$ is [6]:

$$P_{\text{err}}^{\text{Hel}} = \frac{1}{2}\left(1 - \frac{1}{2}\left\|\rho_E^{(0)} - \rho_E^{(1)}\right\|_1\right). \tag{8}$$

Since both states are diagonal in the Fock basis, the trace norm is $\left\|\rho_E^{(0)} - \rho_E^{(1)}\right\|_1 = p_{\text{capt}}$, giving $P_{\text{err}}^{\text{Hel}} = (1 - p_{\text{capt}}/2)/2$. Photon number measurement achieves this bound exactly; it is the optimal measurement for this state pair.

The conditional probabilities are:

$$P(\text{click} \mid k = 1) = p_{\text{capt}}, \qquad P(\text{click} \mid k = 0) = 0, \tag{9}$$

$$P(\text{no click} \mid k = 1) = 1 - p_{\text{capt}}, \qquad P(\text{no click} \mid k = 0) = 1. \tag{10}$$

The adversary's optimal strategy is: if click, declare $k = 1$ (always correct); if no click, declare $k = 0$ (correct with posterior probability $P(k{=}0 \mid \text{no click})$).

Using Bayes' theorem with a uniform prior $P(k{=}0) = P(k{=}1) = 1/2$:

$$P(k{=}0 \mid \text{no click}) = \frac{1}{2 - p_{\text{capt}}}. \tag{11}$$

3.3 Mutual Information

The mutual information between the Z-basis key bit K_Z and the adversary's measurement outcome E is:

$$I(K_Z; E) = H(K_Z) - H(K_Z \mid E) = 1 - H(K_Z \mid E), \tag{12}$$

where $H(K_Z) = 1$ for uniform key bits. The conditional entropy decomposes as:

$$H(K_Z \mid E) = P(\text{click}) \cdot \underbrace{H(K_Z \mid \text{click})}_{=0} + P(\text{no click}) \cdot H(K_Z \mid \text{no click}), \tag{13}$$

with $P(\text{click}) = p_{\text{capt}}/2$ and:

$$H(K_Z \mid \text{no click}) = h\left(\frac{1}{2 - p_{\text{capt}}}\right). \tag{14}$$

3.4 X-Basis Key Bits

For X-basis key bits ($|+\rangle$ and $|-\rangle$), both states have equal probability 1/2 of emitting a photon (both are symmetric superpositions of $|0\rangle$ and $|1\rangle$). The adversary's photon number measurement provides no information about which X-basis state was sent:

$$I(K_X; E) = 0. \tag{15}$$

This result depends on the adversary's choice of measurement. Homodyne detection, which measures a field quadrature rather than photon number, could in principle distinguish $|+\rangle$ from $|-\rangle$ via their relative phase signatures, potentially yielding $I(K_X; E) > 0$. However, Assumption A2 restricts TID routing to the coherent-state basis, and the environmental states (Eqs. 6–7) are diagonal in the Fock basis with no off-diagonal coherence. Homodyne detection of a state $\rho_E^{(1)} = (1 - p_{\text{capt}})|0\rangle\langle 0| + p_{\text{capt}}|1\rangle\langle 1|$ yields a symmetric quadrature distribution regardless of the phase of the original qubit state because the mixture has destroyed the phase information. Consequently, homodyne detection provides no advantage over photon counting for this attack model. If the routing mechanism were to preserve coherence in the environmental state producing displaced or phase-dependent states rather than Fock-diagonal mixtures, then the analysis would change substantially, and this remains an open question tied to the detailed physics of TID routing.

3.5 Sifted Key Information

After basis reconciliation (publicly announced), the adversary knows which sifted key bits were Z-basis and which were X-basis. The average mutual information per sifted key bit is:

$$I_{\text{sift}} = \frac{1}{2} I(K_Z; E) + \frac{1}{2} I(K_X; E) = \frac{1}{2} I(K_Z; E). \tag{16}$$

4 Key Rate Analysis

4.1 Devetak-Winter Bound

The BB84 [2] secret key rate under one-way error correction, following the Devetak-Winter bound [4, 3], is:

$$R = I(A; B) - I(A; E) = 1 - h(\text{QBER}) - I_{\text{sift}}. \tag{17}$$

Alice and Bob, unaware of the TID-routed leakage, estimate the key rate using the standard Shor-Preskill formula:

$$R_{\text{std}} = 1 - 2h(\text{QBER}). \tag{18}$$

A critical subtlety arises here. The Shor-Preskill proof [3] bounds Eve's information via the phase-error rate, which in standard BB84 equals the bit-error rate by symmetry. The TID attack does not violate this symmetry. It does not introduce bit errors or phase errors into the transmitted signal, and the standard security proof remains valid

against any eavesdropper whose information is captured by the channel's error statistics. Instead, TID routing creates an additional information pathway that is entirely invisible to QBER monitoring. The adversary gains information about Z-basis key bits through the environmental side channel (Section 3.1), but this leakage does not disturb the transmitted qubits and therefore does not appear in any error rate that Alice and Bob can measure. The security breach occurs not because the Shor-Preskill bound is wrong, but because it does not account for information that bypasses the quantum channel entirely.

The security gap is:

$$\Delta R = R_{\text{std}} - R = I_{\text{sift}} - h(\text{QBER}). \tag{19}$$

Proposition 1 (Security Breach Condition)**.** *The TID routing attack breaches BB84 security if and only if:*

$$I_{\text{sift}} > h(\text{QBER}), \tag{20}$$

i.e., the adversary's actual information per sifted bit exceeds the information removed by standard privacy amplification. When this condition holds, the distilled key contains ΔR bits per sifted bit of uncompensated adversarial information. When it does not hold, standard privacy amplification is sufficient, and the attack fails.

4.2 The Low-QBER Vulnerability

The breach condition (Eq. 20) reveals quite a counterintuitive vulnerability: lower QBER makes the channel more vulnerable, not less.

The adversary's information I_{sift} depends on the channel loss (which determines p_{capt}) and the TID routing fraction, but is independent of QBER. The privacy amplification removal $h(\text{QBER})$ depends only on QBER. As QBER decreases, $h(\text{QBER})$ decreases faster than I_{sift}, eventually falling below I_{sift} and opening the security gap.

For a 10 dB microwave channel with $f_{\text{int}} = 0.01$, the crossover occurs at QBER $\approx 6\%$:

Table 1: Security breach analysis at 10 dB channel loss, $f_{\text{int}} = 0.01$, $\nu = 5$ GHz.

QBER	$h(\text{QBER})$	I_{sift}	ΔR	Status
0.001	0.011	0.328	0.317	**Breach**
0.002	0.021	0.328	0.307	**Breach**
0.005	0.045	0.328	0.283	**Breach**
0.010	0.081	0.328	0.247	**Breach**
0.020	0.141	0.328	0.187	**Breach**
0.050	0.286	0.328	0.042	**Breach**
0.060	0.327	0.328	0.001	**Breach**
0.065	0.347	0.328	−0.019	Secure
0.080	0.402	0.328	−0.074	Secure
0.110	0.500	0.328	−0.172	Secure

Remark 1. *This inversion of the usual security intuition arises because TID routing operates on lost photons, not transmitted ones. High-quality channels (low QBER) apply less privacy amplification, leaving more room for environmental leakage to be uncompensated. The conventional wisdom that lower QBER implies higher security relies on the assumption that environmental leakage is bounded by the QBER, an assumption that TID routing would violate.*

Fig. 2 visualises this vulnerability: the adversary's information I_{sift} (blue, constant with respect to QBER) exceeds the privacy amplification removal $h(\text{QBER})$ (green, increasing) for all QBER values below the crossover at $\approx 6\%$. The shaded region represents the parameter space where the attack breaches security.

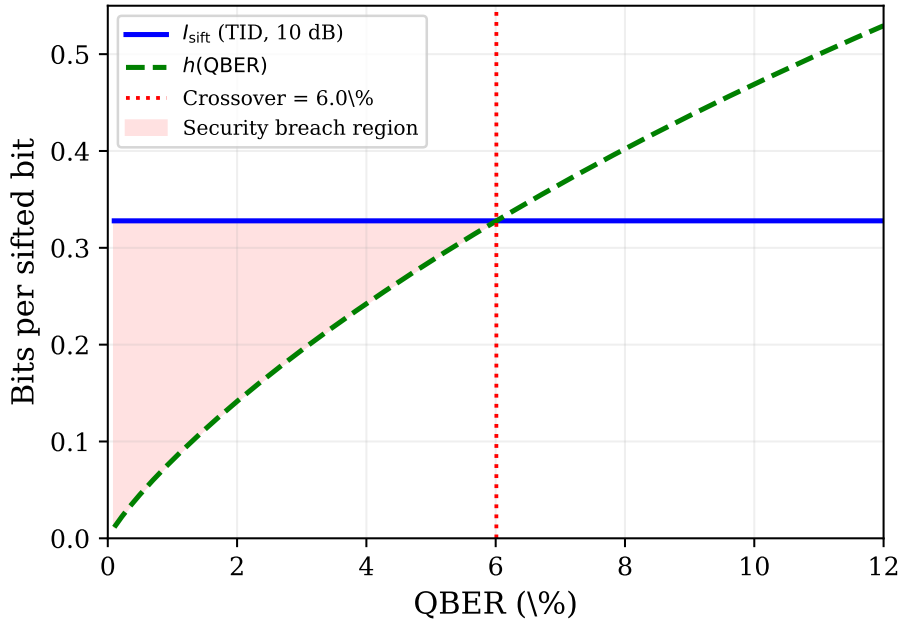


Figure 2: The low-QBER vulnerability. The adversary's mutual information $I_{\text{sift}} = 0.328$ bits (blue, independent of QBER) is compared with the privacy amplification removal $h(\text{QBER})$ (green dashed). The shaded region marks the security breach: for QBER $\lesssim 6\%$, the TID-routed leakage exceeds the information removed by standard privacy amplification. Parameters: 10 dB loss, $f_{\text{int}} = 0.01$, $\nu = 5$ GHz.

4.3 Channel Loss Dependence

The adversary's information I_{sift} increases with channel loss because more photons are available for environmental capture:

Table 2: Security gap vs channel loss at QBER $= 0.5\%$, $f_{\text{int}} = 0.01$, $\nu = 5$ GHz.

Loss (dB)	p_{lost}	p_{capt}	I_{sift}	ΔR	Bits/256-bit key
1	0.206	0.191	0.051	0.006	1.5
2	0.369	0.342	0.098	0.053	13.5
5	0.684	0.634	0.213	0.168	43.0
10	0.900	0.834	0.328	0.283	72.4
15	0.968	0.898	0.377	0.331	84.8
20	0.990	0.918	0.395	0.349	89.4

At 10 dB loss (typical for a 5-metre microwave waveguide at room temperature) with 0.5% QBER, the adversary obtains 72 uncompensated bits of information per 256-bit key. Alice and Bob believe their key is secure when it is not. Fig. 3 shows the continuous dependence: the TID adversary's leakage rises steeply with loss (blue), while a stan-

dard adversary with the same 1% coupling obtains negligible information (red) because standard routing does not benefit from the capacity gradient.

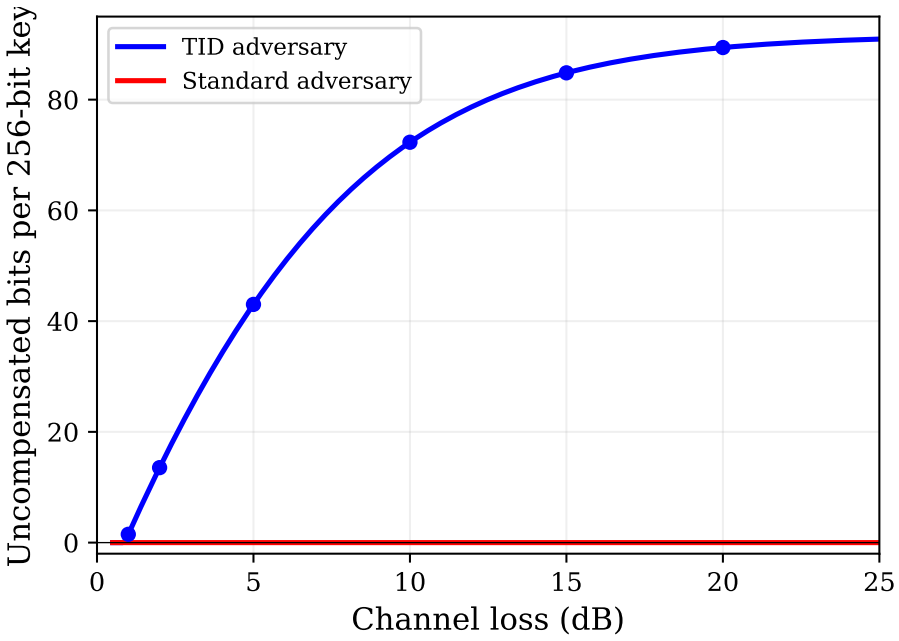


Figure 3: Uncompensated adversarial information per 256-bit key as a function of channel loss, for TID (blue) and standard (red) adversaries at $f_{\text{int}} = 0.01$, QBER = 0.5%. Dots mark the values from Table 2. Under standard routing, 1% coupling yields negligible leakage at all losses; under TID, the same coupling yields up to 89 uncompensated bits per key.

4.4 The 1% Coupling Sufficiency

A remarkable feature of the TID routing enhancement is that the adversary needs only minimal physical coupling to the channel. At $f_{\text{int}} = 0.01$ (1% coupling), TID routes 92.7% of environmental information to the adversary (vs. 1% under standard theory). Increasing f_{int} beyond 0.01 provides diminishing returns because the routing is already near-saturated:

Table 3: TID routing enhancement vs adversary coupling fraction. 10 dB loss, QBER = 0.5%.

f_{int}	$f_{\text{route}}^{\text{std}}$	$f_{\text{route}}^{\text{tid}}$	Enhancement	I_{sift}	ΔR
0.001	0.001	0.556	556×	0.156	0.110
0.005	0.005	0.863	173×	0.292	0.247
0.010	0.010	0.927	93×	0.328	0.283
0.020	0.020	0.962	48×	0.351	0.306
0.050	0.050	0.985	20×	0.368	0.323
0.100	0.100	0.993	10×	0.374	0.328

This has important practical implications. The adversary's physical footprint can be minimal. A single cryogenic resonator with weak coupling to the channel waveguide is sufficient.

5 Attack Protocol

The complete attack protocol, conditional on assumptions A1 and A2, is:

Phase 1: Environmental Preparation. The adversary positions a cryogenic high-Q microwave resonator (temperature $T_{\text{adv}} \leq 20$ mK, $\bar{n} \approx 0$) with weak coupling ($f_{\text{int}} \sim 1\%$) to the warm segment of the microwave quantum channel. No modification of the channel itself is required.

Phase 2: Passive Information Capture. During QKD transmission, each lost photon's information is routed by TID capacity weighting. With $C_{\text{adv}}/C_{\text{env}} \approx 1250$, even 1% coupling captures ∼93% of the environmental information. The adversary's resonator accumulates a photon-number record correlated with each Z-basis key bit.

Phase 3: Mode Readout. After each transmission window, the adversary performs dispersive readout of the resonator's photon number state. Detection of a photon ($|1\rangle$) identifies the key bit as 1 with certainty. No detection is consistent with key bit 0 (high posterior probability).

Phase 4: Key Reconstruction. After public basis reconciliation, the adversary identifies the Z-basis key bits and applies their measurement outcomes. For the X-basis bits, the adversary has no information. The adversary's mutual information per sifted bit is $I_{\text{sift}} = I(K_Z; E)/2$.

Phase 5: Exploitation. Alice and Bob apply privacy amplification calibrated to QBER, removing $h(\text{QBER})$ bits per sifted bit. If $I_{\text{sift}} > h(\text{QBER})$, the distilled key contains $\Delta R = I_{\text{sift}} - h(\text{QBER})$ bits per sifted bit of uncompensated adversarial information.

6 Conditional Assumptions and Falsifiability

All results in this paper are conditional on two assumptions:

Assumption A1 (TID Validity). Environmental information recording capacity modifies decoherence routing probabilities beyond what standard coupling strengths determine. This is the core physical claim of the TID framework [1].

Assumption A2 (Coherent-State Routing Basis). The decoherence routing occurs in the coherent-state basis, not the Fock basis. TID predicts zero deviation in the Fock basis [1]. If the routing basis is Fock, all results in this paper produce null values and no attack exists.

Falsifiability. Both assumptions can be tested by the decisive experiment proposed in [1]: a transmon qubit coupled symmetrically to thermal and squeezed vacuum microwave modes at matched mean photon number, with branching ratio readout. Standard QM predicts 50:50 branching; TID predicts 60.7:39.3 at squeezing parameter $r = 1$. This experiment is feasible with current superconducting circuit technology [7].

The routing basis question can be resolved by performing the branching measurement with both homodyne detection (coherent-state basis) and photon number detection (Fock basis). If the branching shift appears with homodyne but not with photon counting, A2 is confirmed. If neither shows a shift, A1 is falsified.

Additionally, for practical microwave QKD, the detector type determines the routing basis. If the receiving node uses dispersive readout (which projects onto the Fock basis), then TID deviations may vanish even if A1 is true. This question requires dedicated theoretical analysis connecting the detector POVM to the routing basis in a microwave quantum network architecture.

7 Countermeasures

If assumptions A1 and A2 are experimentally confirmed, the following countermeasures become necessary for microwave quantum communication:

Capacity-corrected privacy amplification. Replace $h(\text{QBER})$ with $\max(h(\text{QBER}), I_{\text{sift}}^{\text{bound}})$ in the privacy amplification step, where $I_{\text{sift}}^{\text{bound}}$ is a conservative upper bound on the adversary's mutual information computed from the monitored environmental capacity gradient. This requires no hardware modification and can be implemented in software, making it the most immediately deployable countermeasure. The cost is a reduced key rate, since more bits are sacrificed to privacy amplification.

Environmental capacity monitoring. Real-time characterisation of the capacity structure of environmental modes adjacent to the channel, with sensitivity to capacity ratios ≥ 10. This would require microwave noise thermometry or spectroscopy of the environmental mode spectrum along the channel path. Current cryogenic noise measurement techniques [13] can resolve thermal occupation numbers at the level required, but integrating such monitoring into a deployed quantum link adds significant infrastructure complexity.

Defensive thermal flooding. Actively heat the channel environment to ensure all modes are in high-$\bar{n}$ thermal equilibrium, where TID deviations vanish by the KMS condition (Lemma 1). This is straightforward in principle—the channel segments are already at room temperature—but requires ensuring that no anomalously cold modes exist in the vicinity, which may be difficult to guarantee in complex multi-cryostat architectures.

Cryogenic shielding. Prevent adversarial placement of cold modes by shielding the warm channel segment against external cryogenic coupling. Physical shielding of microwave waveguides is well-established technology, although ensuring complete coverage along an extended channel path presents engineering challenges.

Lemma 1 (KMS Condition [1])**.** *For environmental modes in thermal equilibrium ρ_{th}, the KMS condition guarantees $\mathcal{T}[\rho_{\text{th}}] = \rho_{\text{th}}$, so TID reduces to standard Lindblad dynamics and no capacity-based routing bias exists.*

8 Discussion

8.1 The Paradox of Low-QBER Vulnerability

The finding that lower QBER increases vulnerability (Section 4.2) inverts the conventional security hierarchy. In standard QKD security analysis, a lower QBER always implies a higher secure key rate. Under TID routing, this remains true for the *standard estimate* $R_{\text{std}} = 1 - 2h(\text{QBER})$, but the actual key rate $R = 1 - h(\text{QBER}) - I_{\text{sift}}$ can decrease relative to R_{std} as QBER falls.

The physical reason is that TID routing exploits the lost photons, whose information content is independent of QBER (which affects the transmitted photons). Privacy amplification, calibrated to QBER, removes less information as QBER falls, while the TID-routed leakage remains constant. Fig. 4 shows this directly: the gap between the estimated key rate R_{std} (green dashed) and the actual key rate R_{actual} (blue solid) is largest at low QBER and closes at the crossover near 6%.

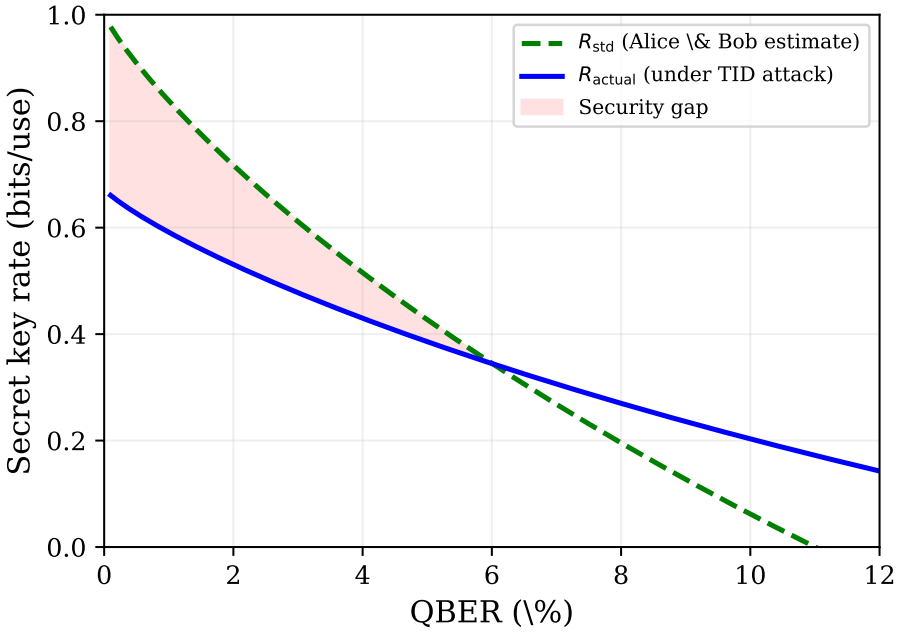


Figure 4: Secret key rate as a function of QBER for the standard estimate $R_{\text{std}} = 1 - 2h(\text{QBER})$ (green dashed, what Alice and Bob compute) and the actual rate $R_{\text{actual}} = 1 - h(\text{QBER}) - I_{\text{sift}}$ under TID attack (blue solid). The shaded region represents the security gap: uncompensated adversarial information in the distilled key. The gap is largest at low QBER and closes at the crossover ($\approx 6\%$). Parameters: 10 dB loss, $f_{\text{int}} = 0.01$, $\nu = 5$ GHz.

8.2 Attack Feasibility

The attack requires: (i) a cryogenic resonator at ≤ 20 mK adjacent to the warm channel segment; (ii) weak coupling to the channel ($\sim 1\%$); (iii) dispersive readout of the resonator after each transmission. All three are achieved routinely in superconducting circuit laboratories [7]. The adversary does not need to intercept, modify, or inject photons into the channel, making the attack fundamentally different from, and potentially less detectable than, standard intercept-resend or Trojan-horse attacks [8, 9].

8.3 Implications for Superconducting Quantum Networks

Superconducting quantum networks connecting multiple cryostats via microwave quantum links [10, 11, 12] are a rapidly developing technology. The warm channel segments between cryostats (typically at room temperature) provide exactly the thermal environment that maximises TID capacity gradients. If TID is valid, these warm segments are the most vulnerable points in the network. This is a finding with direct implications for the physical security architecture of quantum data centres and modular quantum computers.

9 Limitations

First, all results are conditional on two as yet unproven physical assumptions (A1 and A2).

Second, the channel model treats photon loss as a binary event (absorbed or transmitted) without modelling partial decoherence, phase noise, or non-Markovian dynamics.

Third, the analysis assumes ideal single-photon Fock-state encoding (see the encoding assumption noted in Section 3.1). Extensions to coherent-state encoding would modify the environmental states and the adversary's optimal measurement.

Fourth, the adversary's measurement is modelled as ideal photon number detection. Realistic detector inefficiency $\eta_{\text{det}} < 1$ reduces the effective capture probability to $p_{\text{capt}}^{\text{eff}} = \eta_{\text{det}} \cdot p_{\text{capt}}$, which lowers I_{sift} and raises the QBER threshold for breach. At $\eta_{\text{det}} = 0.9$ (typical for dispersive readout in current superconducting circuits), the crossover shifts from QBER $\approx 6\%$ to $\approx 5.2\%$; at $\eta_{\text{det}} = 0.7$, it shifts to $\approx 3.5\%$. Dark counts would further reduce effective I_{sift} by introducing false positives. The attack remains effective at realistic detector efficiencies for low-QBER channels, though the security margin narrows.

Fifth, the key rate analysis uses the Shor–Preskill one-way bound. Finite-key effects, two-way error correction, and more sophisticated privacy amplification protocols could modify the security margin.

Sixth, the connection between the detector POVM and routing basis is proposed but not derived from first principles. This is the most critical theoretical gap: without it, A2 remains a free parameter.

10 Conclusion

I have derived a complete information-theoretic attack chain exploiting TID capacity-weighted environmental routing for key recovery from microwave quantum channels. Three principal findings emerge.

First, the attack operates through a surprising amplification mechanism: even 1% adversarial coupling captures $\sim$93% of environmental information under TID routing, because the 1250$\times$ capacity ratio between cryogenic and room-temperature modes overwhelmingly biases information flow.

Second, the attack inverts the conventional security hierarchy: lower QBER means less privacy amplification, leaving more uncompensated leakage. At QBER = 0.5% with 10 dB channel loss, the adversary obtains 72 uncompensated bits per 256-bit key.

Third, the attack surface is restricted to microwave quantum communication. Standard optical fibre QKD is immune because all environmental modes at optical frequencies have negligible thermal occupation.

All results are conditional on TID being physically valid and the routing basis being the coherent-state basis. The decisive experiment, competitive branching between thermal and squeezed microwave modes, will determine whether this attack surface exists. If confirmed, the security architecture of superconducting quantum networks requires fundamental reassessment.

Declarations

Funding: This research received no external funding.

Informed Consent: Not applicable.

Author Contribution: All work was done by the sole author.

Data Availability: All analytical results are derived in the text. Computational scripts are available as supplementary materials.

Conflict of Interest: The author is a Quantum Encryption PhD candidate at the University of Portsmouth and serves as Head of Research at Argo DevOps Solutions.

The research was conducted independently. No proprietary data or resources from Argo DevOps Solutions were used.

References

[1] V. Mellor, "Information Recording Capacity as a Physical Constraint on Quantum Decoherence," University of Portsmouth preprint (2026). DOI: 10.5281/zenodo.18914394. Under review at *Nature Discovery Quantum*.

[2] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," in *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, Bangalore, India (IEEE, New York, 1984), pp. 175–179.

[3] P. W. Shor and J. Preskill, "Simple proof of security of the BB84 quantum key distribution protocol," *Phys. Rev. Lett.* **85**, 441 (2000).

[4] I. Devetak and A. Winter, "Distillation of secret key and entanglement from quantum states," *Proc. R. Soc. A* **461**, 207–235 (2005).

[5] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, "The security of practical quantum key distribution," *Rev. Mod. Phys.* **81**, 1301–1350 (2009).

[6] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, New York, 1976).

[7] J. Koch *et al.*, "Charge-insensitive qubit design derived from the Cooper pair box," *Phys. Rev. A* **76**, 042319 (2007).

[8] L. Lydersen *et al.*, "Hacking commercial quantum cryptography systems by tailored bright illumination," *Nature Photonics* **4**, 686–689 (2010).

[9] N. Jain *et al.*, "Trojan-horse attacks threaten the security of practical quantum cryptography," *New J. Phys.* **16**, 123030 (2014).

[10] P. Magnard *et al.*, "Microwave quantum link between superconducting circuits housed in spatially separated cryogenic systems," *Phys. Rev. Lett.* **125**, 260502 (2020).

[11] P. Kurpiers *et al.*, "Deterministic quantum state transfer and remote entanglement using microwave photons," *Nature* **558**, 264–267 (2018).

[12] N. Leung *et al.*, "Deterministic bidirectional communication and remote entanglement generation between superconducting qubits," *npj Quantum Inf.* **5**, 18 (2019).

[13] L. Spietz, K. W. Lehnert, I. Siddiqi, and R. J. Schoelkopf, "Primary electronic thermometry using the shot noise of a tunnel junction," *Science* **300**, 1929–1932 (2003).

[14] C. W. Gardiner and P. Zoller, *Quantum Noise*, 3rd ed. (Springer, Berlin, 2004).

[15] S. L. Braunstein and A. K. Pati, "Quantum information cannot be completely hidden in correlations," *Phys. Rev. Lett.* **98**, 080502 (2007).

[16] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz, "Time symmetry in the quantum process of measurement," *Phys. Rev.* **134**, B1410 (1964).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- analyticalframeworkPythonCode.py
- MicrowaveAttack1.zip
- MicrowaveAttack1.zip