

Hybrid Quantum–Computational Key Distribution with Reduced Assumptions (Short Paper)

Jean Snyman

Surrey Centre for Cyber Security
University of Surrey
Guildford, United Kingdom
research@snyman.ca

Ioana Boureanu

Surrey Centre for Cyber Security
University of Surrey
Guildford, United Kingdom
i.boureanu@surrey.ac.uk

Daniel Gardham

Surrey Centre for Cyber Security
University of Surrey
Guildford, United Kingdom
daniel.gardham@surrey.ac.uk

Abstract

The *HM-QCT* protocol, presented by Mazzoncini et al. in 2025, uses a unique feature of quantum information to achieve *quantum key distribution (QKD)* with secret-key rates beyond standard constructions. The protocol has been designed in the hybrid *quantum–computational time-lock (QCT)* security model which imposes two independent (but related) restrictions on the attacker: (1) a computationally secure encryption scheme which a successful attacker must break within a narrow timeframe and (2) an assumption that the attacker can only store quantum information with low fidelity.

We first review the protocol and its two-assumption model by presenting, in explicit form, the two natural “harvest now, decrypt later” strategies against it, one using quantum memory and one using classical memory. We show how each is excluded by exactly one of the two assumptions, so that harvesting resistance rests on their conjunction and on the presumed existence of a technological gap: the adversary’s quantum memory must decohere much faster than the time it takes to break the short-term encryption.

Finally, we introduce **HM-PRF**, a variant that replaces the short-term encryption with a *pseudorandom function (PRF)*, evaluated locally. By eliminating the transmission of enciphered classical information, the adversary does not need to race an impending expiry event. We compare the security of our construction against the original and argue that our PRF-based assumption yields a more realistic adversary. It simply requires the existence of *one-way functions* rather than a predicted lasting limitation of quantum computation and memory technology.

CCS Concepts

• **Security and privacy** → **Security protocols**; *Cryptography*; • **Hardware** → *Quantum communication and cryptography*.

Keywords

Quantum Key Distribution, Noisy-Storage Model

ACM Reference Format:

Jean Snyman, Ioana Boureanu, and Daniel Gardham. 2026. Hybrid Quantum–Computational Key Distribution with Reduced Assumptions (Short Paper). In *Post-Quantum and Quantum-based Security (PQQS ’26)*, November 02–04, 2026, Milpitas, CA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3839338.3857249>



This work is licensed under a Creative Commons Attribution 4.0 International License. PQQS ’26, Milpitas, CA, USA

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2920-1/2026/11
<https://doi.org/10.1145/3839338.3857249>

1 Introduction

Standard quantum key distribution (QKD) is a form of multi-party computation whereby (usually) two communication partners arrive at a common key used to secure subsequent communications. QKD works differently from classical key agreement, preferring to harness the behaviour of subatomic particles over security derived from computationally hard problems. This has certain benefits: for one, QKD schemes typically have *everlasting security* [10], meaning the negotiated key cannot be recovered by an attacker once the protocol completes, even if computational security is later broken.

The utility of QKD, however, is constrained by its inefficiency. Known as the *PLOB bound*, the key rate achievable by a typical protocol is limited to $-\log_2(1 - \eta)$ bits per channel use [8]. This rate decreases as the fractional loss of the channel (i.e., $1 - \eta$) increases with transmission distance. It is therefore not uncommon to see transmissions at a rate of less than one bit per channel use even at distances of only a few kilometres. A 56K dial-up modem transmits around seven bits with each channel use, for context.

The thesis of Mazzoncini [6] presents a novel construction capable of exceeding the PLOB bound. This surprising result is made possible by the *hidden matching problem* [1, 3] whereby information about a particular graph is partitioned among two parties, one of whom aims to find an edge belonging to the graph with only partial knowledge thereof. A known quantum protocol has been shown to be exponentially more efficient than any possible classical variant. This is repurposed for QKD by the HM-QCT hybrid protocol [7] which combines the usual quantum channel with an encrypted classical one to partition the information about the chosen graph. The increased efficiency allows the transmitted data to be repeated which lowers the error rate and increases the key rate by a factor $O(\sqrt{n}/\log n)$ over BB84 [2].

Thanks to the information-theoretic properties of the hidden matching problem, an attacker can only obtain the negotiated key if they have both the vertex information transmitted over the quantum channel and the edge information transmitted over the classical channel. Security is achieved if either or both is inaccessible during and after protocol execution. The secrecy of the classical edge information of course hinges on the computational security of the encryption scheme used. HM-QCT attempts to reduce this reliance by allowing the scheme to be broken after a certain time t_{comp} greater than the protocol’s execution time. The quantum vertex information cannot therefore be available to the attacker after t_{comp} , or security would be compromised. Any attacker is thus presumed to only have access to imperfect quantum memory which degrades to random noise by t_{comp} , an assumption imported from the *noisy-storage model (NSM)* [13].

This set of threat model assumptions together make up the *quantum-computational time-lock (QCT) model*¹ of Vyas and Al-léaume [12]. Within these assumptions, however, there exists a tension which may be uncomfortable to cryptographers. Conventionally, as far as an attacker's capabilities are concerned, restricted computational power is traded for unlimited memory capacity, or vice versa. This bargain is illustrated by Maurer's original *bounded-storage model* [5], the classical ancestor of the NSM and QCT models. Conversely, the QCT model restricts the attacker in terms of both computation *and* memory.

Contributions. The QCT model and HM-QCT protocol aim to achieve efficient QKD by incorporating a computational assumption and an assumption from the noisy-storage literature. We show that these assumptions are not fundamental. By deriving each session's edge information from a shared long-term secret, the protocol no longer requires delayed disclosure of the edge information and therefore no longer relies on assumptions about the adversary's quantum storage capabilities or its ability to break encryption.

2 Background

2.1 The Hidden Matching Problem and the Communication Complexity Gap

The *hidden matching problem* described by Bar-Yossef et al. [1] has a curious property: no classical one-way protocol can solve it without transmitting at least $\Omega(\sqrt{n})$ bits, but $O(\log n)$ qubits are all that is needed to do so quantumly. Gavinsky et al. [3] also showed this to be the case for a boolean variant which [6] calls the β -*partial matching* (β PM) *problem*.²

The problem. Alice has a string $x \in \{0, 1\}^n$ representing the n vertices of a graph. Bob has a sequence of βn edges represented by pairs of unique vertex indices. This is called a *matching*. Taken together, Bob's matching M and Alice's x form a graph where each vertex is joined to at most one other. When all vertices are joined ($\beta = 0.5$), the matching is said to be *perfect*. Otherwise, if one or more vertices are left unconnected ($\beta < 0.5$), it is known as a *partial matching*. Suppose that for each edge $(i, j) \in M$ we computed $x_i \oplus x_j$ and concatenated the resulting values to form a bit string v of length βn . In addition to M , Bob is also given $w = v \oplus a^{\beta n}$ where $a \in \{0, 1\}$. Bob's goal is to determine whether he received w equal to v or $\bar{v}$, i.e., v with all bits flipped. If it were possible for Bob to determine one bit of v , then he could check the corresponding bit in w to arrive at the answer $a = 0$ or $a = 1$.

Classical protocol. Bob cannot know anything about v without learning something about x . Alice sends $\sqrt{n/\beta}$ bits of x to Bob which is likely to contain at least one of the edges in M due to the birthday paradox. If so, this allows Bob to determine at least one bit in v and thereby a .

Quantum protocol. If Alice can communicate with Bob over a quantum network, she can simply encode x as the linear combination of n basis states such that $|\psi_x\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n (-1)^{x_i} |i\rangle$. For example, $|\psi_{1001}\rangle = \frac{1}{\sqrt{4}} (|1\rangle + |2\rangle + |3\rangle - |4\rangle)$. These $\log(n)/2\beta$

qubits are sent to Bob who can then use his knowledge of M and w to compute a . He does this by applying a projector $|i\rangle\langle i| + |j\rangle\langle j|$ to $|\psi_x\rangle$ for each $(i, j) \in M$. With probability 2β (which equals 1 when M is a perfect matching), this measurement reveals one of the edges (i, j) in the graph described by M and x and collapses the state to a simple combination of the two corresponding basis states. Performing a further projective measurement in the Hadamard basis has the effect of XORing the two components together to obtain a single bit of v . Since the edge (i, j) which produced this bit $x_i \oplus x_j$ is known, Bob can check the corresponding bit in w to determine a .

2.2 The HM-QCT Protocol




Figure 1: One round of the HM-QCT protocol.

Adapted from Fig. 7.1 in [6] and Fig. 1 in [7].

The *hidden matching QCT (HM-QCT) protocol* proceeds according to a number of rounds where each round solves an instance of the β PM problem to arrive at a single key bit (see Fig. 1). The protocol is performed as follows:

Setup. The preconditions are:

- Alice and Bob each hold a pre-shared secret key k of length ℓ_k (established out of band or by prior key establishment run).
- A quantum channel and an authenticated classical channel exists between them.
- The parties have agreed on a set of possible matchings $\mathcal{M}_{\beta n}$ where each element is a β -matching of n vertices.
- The parties have also agreed on a number of copies m .

Protocol execution. In each round i :

- (1) Alice selects a bit string $x_i \leftarrow \{0, 1\}^n$ and a matching $M_i \leftarrow \mathcal{M}_{\beta n}$ uniformly at random. She computes v_i from M_i and x_i and sets $w_i := v_i$ or $w_i := \bar{v}_i$ by coin flip.
- (2) Alice forms the message $y_i := (M_i, w_i)$ and sends $\text{Enc}_k(y_i)$ to Bob over the classical channel.
- (3) Alice encodes x_i as a quantum state $|\psi_{x_i}\rangle$ and transmits m copies of this state to Bob over the quantum channel. She also evaluates a function $\beta\text{PM}(x_i, y_i)$ which computes $a_i^{\beta n} = v_i \oplus w_i$ and returns a_i , the key bit for the round.
- (4) Bob computes his own copy of the round bit b_i using the information contained in y_i to appropriately measure $|\psi_{x_i}\rangle$. When M is an imperfect matching, Bob obtains an edge $(i, j) \notin M$ with probability $1 - 2\beta$ and outputs $b_i = \perp$. Otherwise, $b_i = a_i$.

Post-processing. Alice and Bob perform reconciliation and privacy amplification to arrive at a final key.

¹The NSM and QCT models are not adversarial models of the type typical in classical cryptography but serve to enumerate the assumptions which apply to a given protocol.

²Originally referred to as the α -partial matching (α PM) problem by Gavinsky et al.

2.3 Standard QKD Assumptions

As with all QKD schemes, the HM-QCT protocol (and our derivative in Section 4) relies on the following basic assumptions [9]:

- (1) The devices used by the parties and all communication channels between them are correctly and completely explained by quantum theory.
- (2) The parties have access to an authenticated channel on which no message modified or injected by an attacker is accepted by the receiving party.
- (3) The devices used locally by the parties faithfully perform all instructions given by the parties.

For a specific protocol, these base-level assumptions are combined with certain additional threat model assumptions such as those given by the NSM or QCT models.

2.4 QCT Protocol Assumptions

A *quantum-computational time-lock (QCT) protocol*, a term coined by Vyas et al. [12] for their MUB-QCT construction, achieves confidential key agreement by combining two separate security mechanisms: (1) a computational mechanism to protect the negotiated key up to a certain time and (2) the presumption of memory noise meant to protect the key beyond that.

The set of assumptions applicable to QCT protocols such as MUB-QCT and HM-QCT are described by what is alternately called the *quantum-computational hybrid (QCH) model* [12] or the *QCT model* [7]. These are twofold:

ASSUMPTION 1 (TIME-LOCK ENCRYPTION). *The chosen encryption scheme has indistinguishable encryptions under a chosen ciphertext attack (IND-CCA security) for a duration $\geq t_{\text{comp}}$.*

ASSUMPTION 2 (NOISY QUANTUM STORAGE). *The adversary's quantum memory is a noisy channel Φ_t such that a state stored for time t_{comp} is δ -indistinguishable from the output of the completely depolarising channel $\mathcal{F}$. The choice of t_{comp} is such that the coherence time of available quantum memories t_{coh} is much less (i.e., $t_{\text{coh}} \ll t_{\text{comp}}$), resulting in $\delta \ll 1$.*

The scheme described by the first assumption could be seen as a form of time-release encryption [11]. The second assumption is lifted from the *noisy-storage model (NSM)* [4, 13], originally described in the context of quantum oblivious transfer.

Notice that no assumption is made as to the “hardness” of the β PM problem. This is because, in a given round i , the protocol carries the information-theoretic guarantee that a_i is unobtainable without y_i . Recall that $a_i^{\beta n} = v_i \oplus w_i$. Since $a_i^{\beta n}$ is either $0^{\beta n}$ or $1^{\beta n}$, only w_i and a single bit from v_i is required to derive a_i . Still, computing a single bit of v_i requires measuring $|\psi_{x_i}\rangle$ with projectors corresponding to M_i . Therefore, as $y_i = (M_i, w_i)$, knowledge of y_i is required to derive a single bit of $v_i \oplus w_i$ and thereby a_i .

To see how the two assumptions of QCT combine, we now explain how HM-QCT achieves its security.

2.5 HM-QCT Security

Let A stand for the classical key bit Alice holds in a given round (the value the protocol aims to keep secret from Eve) and Q for the

information Alice transmits on the quantum channel (i.e., $|\psi_{x_i}\rangle$) before it is subjected to the noise of the channel. Upon intercepting Q , Eve can choose to measure or process the information in whole or in part resulting in a new quantum state Q_{in} and classical measurement outcome Z , both of which Eve will store in memory. Y is the information obtained when Eve breaks the IND-CCA security of the encryption scheme after time t_{comp} . This is shown in Fig. 2.




Figure 2: Eve's general attack strategy.

Adapted from Fig. 1 in [13], Fig. 6.2 in [6] and Fig. 4 in [7].

The quantity $P_{\text{guess}}(A | \cdot)$ denotes the *guessing probability*: the largest probability with which an optimal strategy recovers A , given the information listed after the bar. Thus, upon the release of Y , $P_{\text{guess}}(A | YZ)$ is Eve's success probability having measured everything immediately and kept only classical data, while $P_{\text{guess}}(A | YZ \Phi_t(Q_{\text{in}}))$ is her success probability if she does not measure the entire state and instead stores all or part of it in her (imperfect) quantum memory Φ_t . Here $\Phi_t(\cdot)$ is the completely positive map which takes the state Eve stored and returns the degraded state that survives after the storage time t .

Consider now the map $\mathcal{F}(\cdot)$ which models worst-case decoherence. It simply erases its input and outputs pure noise in the form of the maximally mixed state. The diamond norm $\|\Phi_t - \mathcal{F}\|_\diamond$ measures how distinguishable the two maps are over all possible inputs. The condition $\|\Phi_t - \mathcal{F}\|_\diamond \leq \delta$ therefore says that any state Eve stores in memory is at most δ away from being useless after time t .

From theorem 6.4.1 in [6], the following result is obtained: if $\|\Phi_{t_{\text{comp}}} - \mathcal{F}\|_\diamond \leq \delta$, then

$$P_{\text{guess}}(A | YZ \Phi_{t_{\text{comp}}}(Q_{\text{in}})) \leq P_{\text{guess}}(A | YZ) + \delta. \quad (1)$$

In simple terms, whatever Eve gains by keeping Q_{in} in her noisy quantum memory is at most δ better than what she could achieve from measuring the state immediately, regardless of her choice of measurement. If we choose our encryption scheme such that t_{comp} occurs long after t_{coh} , which is the point at which the data stored in Eve's memory is no longer useable (i.e., $t_{\text{coh}} \ll t_{\text{comp}}$), then $\delta \ll 1$. With δ near zero, the two sides are essentially equal, so storing the quantum information buys her nothing: she may as well have measured it the moment it arrived.

This is what lets the analysis reduce any attack to an immediate-measurement strategy: instead of reasoning about the vast space of actions Eve might take with a quantum memory over time, it suffices to bound what she learns by measuring at once and keeping classical records. Security then follows by lower-bounding the min-entropy, $H_{\text{min}}(A | YZ \Phi_{t_{\text{comp}}}(Q_{\text{in}}))$, which quantifies how much genuine uncertainty about A remains from Eve's standpoint after t_{comp} .

3 Two Harvesting Attacks

Due to its time-lock encryption, a successful attack on HM-QCT necessitates capturing and storing transmissions until the encryption can be broken. Many permutations of this tactic are possible, varying the information stored. The general strategy (Section 2.5) is called an *encoding attack* in [6], but to simplify, it can be understood by its two natural forms: one storing quantum states and one using only classical memory (see Fig. 3). Only *i.i.d.* attacks are considered, as in the original work, meaning that the attacker uses an identical strategy in each round, i.e., for every negotiated key bit.




Figure 3: The two harvesting attacks on a single round of the HM-QCT protocol.

Attack 1. Harvest now, decrypt later (using quantum memory). For each round i , Eve intercepts at least one copy of $|\psi_{x_i}\rangle$ which encodes the classical bit string x_i Bob uses to derive the negotiated key bit a_i at the end of the round. Simultaneously, she works on breaking $\text{Enc}_k(y_i)$ and succeeds at a time after $t_0 + t_{\text{comp}}$. Doing so, she obtains y_i and thereby the measurement U_i which reveals a_i .

To succeed, Eve must preserve coherence of the intercepted states for a time $t \geq t_{\text{comp}} \gg t_{\text{coh}}$. By Eq. (1), after that time, her register is within δ of fully depolarised noise, so measuring is no better than guessing. Coherence of $|\psi_{x_i}\rangle$ for duration t is exactly what noisy storage prevents and so the strategy is defeated by Assumption 2.

Attack 2. Harvest now, decrypt later (using classical memory). In this variant, Eve does not store quantum information. Instead, she chooses an arbitrary measurement U and applies this to the intercepted $|\psi_{x_i}\rangle$ immediately, storing only the classical outcome b_i which suffers no decoherence. As a result, this attack is also referred to as the *immediate measurement strategy*. Once y_i is obtained after time $t_0 + t_{\text{comp}}$, Eve knows whether her guess of U was correct and therefore whether b_i corresponds to the negotiated key bit or a random value.

Because y_i is time-locked, the measurement basis must be committed to *before* it is known. As a wrong measurement irreversibly destroys the information needed to recover the bit, it is necessary to guess correctly with a non-negligible probability. The number of measurement possibilities increases with n and the information which can be extracted from m copies of $|\psi_{x_i}\rangle$ is limited by the Holevo bound. To improve on this would require breaking the βPM problem. As a result, Eve's success probability reduces to the

classical information complexity of βPM : with y_i hidden, Eve's stored classical record leaves her per-round probability of guessing a_i bounded away from 1 whenever $m \leq O(\sqrt{n}/\log n)$ [6, 7]. Therefore, the strategy is defeated by the unconditional hardness of the βPM problem and the existence of time-limited encryption (Assumption 1).

Comparing the two attacks. The attacks are two sides of the same dilemma: to use the intercepted quantum states, Eve must either *wait* or *measure blind*. Neither threat model assumption (Section 2.4) eliminates both attacks. Drop the noisy storage assumption (Assumption 2) and Attack 1 succeeds. Or drop the short-term encryption (Assumption 1) and Attack 2 succeeds. HM-QCT's resistance to eavesdropping is therefore a property of the conjunction of these, hinging on the non-overlap of the two windows, $t_{\text{coh}} \ll t_{\text{comp}}$. This isolates the cost: the guarantee is contingent on a predicted, lasting technological gap between memory and computational capabilities. Specifically, it assumes a ceiling exists on the coherence of future quantum memory technology. Similar bets on the impracticability of quantum computing have made “harvest now, decrypt later” attacks on classical public-key cryptography an ever-present worry.

4 HM-PRF: From QCT to Pseudorandom Function




Figure 4: One round of the HM-PRF protocol.

4.1 The Construction

The analysis above shows that the two QCT assumptions play very different roles. The presence of time-limited encryption (Assumption 1) is the operative assumption as it is what hides y_i during transmission. The noisy storage (Assumption 2) is purely reactive. It exists only to plug the gap opened by the time-lock: namely Attack 1 in which Eve waits for the lock to expire before measuring stored states. This suggests that a simplified protocol should target the first assumption.

Of course, the reason the time-lock exists at all is that y_i is *transmitted*. One remedy therefore is to not transmit y_i over any channel. If y_i is instead generated locally and identically by each party, there is no encrypted transmission to protect. Thus, the short-term encryption assumption is no longer needed, eliminating Attack 2. And because there is no time-lock, the requirement for noisy storage also disappears, eliminating Attack 1. The QCT assumptions therefore become unnecessary for the security of the construction.

We call the resulting protocol **HM-PRF** (Hidden-Matching with Pseudorandom Function). It replaces the short-term encryption of Step 2 by a *pseudorandom function (PRF)* shared by Alice and Bob. A run of the protocol proceeds as described below (see also Fig. 4):

Setup. All preconditions are the same as in the HM-QCT protocol described in Section 2.2. Notably, Alice and Bob have a preshared key k and have agreed to a set of possible matchings $\mathcal{M}_{\beta n}$.

Protocol execution. In each round i :

- (1) Alice selects a bit string r_i of length ℓ_r uniformly at random. This is sent to Bob over a classical channel.
- (2) Alice and Bob both apply a publicly-known keyed function $F : \{0, 1\}^{\ell_k} \times \{0, 1\}^{\ell_r} \rightarrow \{0, 1\}^{\ell_y}$ to derive pseudorandom material $y_i := F(k, r_i)$. Evaluating the same pseudorandom function on the same k and r_i , the two devices agree on y_i without communicating.
- (3) Alice samples x_i uniformly from random, encodes it m times as $|\psi_{x_i}\rangle$, and transmits these states over the quantum channel. She evaluates $a_i := \beta PM(x_i, y_i)$ to obtain the key bit for the round.
- (4) Bob, having derived his own copy of y_i , computes the round bit. Reconciliation and privacy amplification are unchanged.

The quantum channel and the βPM -based key extraction are exactly as in HM-QCT. The only change is the *provenance* of y_i : locally derived rather than transmitted using short-term encryption.

4.2 Harvesting Resistance of HM-PRF

Attack 1 and HM-PRF. In HM-PRF, y_i is never transmitted over any channel accessible to Eve and therefore does not require protection in transit. Without time-limited encryption that later lapses, there is no expiry event after which Eve learns y_i . Since **Attack 1** has no target y_i to wait for, the attack disappears in HM-PRF.

Indeed, in the original HM-QCT variant, **Attack 1** pays the cost of quantum storage in exchange for eventually learning y_i once the time-lock expires at t_{comp} . In HM-PRF, there is no benefit to storing $|\psi_{x_i}\rangle$ as y_i only exists as the result of computation done at the two endpoints and is never disclosed. Neither is k which gives y_i its pseudorandomness. The noisy storage assumption (**Assumption 2**), whose sole role was to defeat the wait-for- y_i strategy, becomes superfluous and can be dropped.

Attack 2 and HM-PRF. An adversary performing **Attack 2** chooses not to store $|\psi_{x_i}\rangle$ upon interception and must therefore measure blind, an ineffectual strategy. In the original HM-QCT protocol, this is because y_i is not known for the duration of the protocol run. A PRF achieves this at least as strongly as time-limited encryption (in fact, it ensures y_i is unavailable indefinitely rather than for t_{comp}). **Attack 2** therefore remains impractical in HM-PRF, even without the existence of short-term encryption (**Assumption 1**).

Assumptions applicable in HM-PRF. Besides the universal QKD assumptions given in Section 2.3, the security of HM-PRF rests on the PRF and on the hardness of the βPM problem, i.e., the inability to derive a_i without y_i . The later is unconditional, as it is in HM-QCT, since a_i is the result of an XOR with w_i (contained in y_i). Thus, the βPM problem operates similarly to a one-time pad. Because the existence of one-way functions implies the existence of PRFs, one must only assume that one-way functions exist to trust that the HM-PRF protocol is secure against harvesting attacks. As discussed above, neither **Assumption 1** or **2** of the QCT model need to hold, thanks to the PRF.

5 Security and Efficiency Evaluation of HM-PRF

Use of the βPM problem for QKD requires y_i to be made known to the legitimate endpoints without falling into the hands of any other party. Whether this is done by transmitting an enciphered y_i across a communication link or by deriving y_i locally, achieving this necessitates a *trusted setup*. This means there is a trusted entity³ at each end pre-provisioned with key material.

One's choice of trusted setup can affect the probability space in which the parties, and the adversary, operate. In HM-QCT, the round inputs $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$ are drawn from $\mu_{\beta PM}$ which is the uniform distribution over $\mathcal{X} \times \mathcal{Y}$. We can say that $\mu_{\beta PM}$ is a joint distribution of distributions $\mu_{\mathcal{X}}$ and $\mu_{\mathcal{Y}}$, over $\mathcal{X}$ and $\mathcal{Y}$ respectively, each independently uniform. The pre-shared key $k \in \mathcal{K}$, sampled from uniform distribution $\mu_{\mathcal{K}}$, is used as input to probabilistic algorithm $\text{Enc}(k, y_i)$ which performs short-term encryption of y_i . The output distribution of this algorithm is taken over the uniform choices of k and y_i as well as the randomness of Enc .

The main change introduced by HM-PRF is that the round input y_i is no longer drawn from true uniform distribution $\mu_{\mathcal{Y}}$. Instead, it is drawn from pseudorandom distribution ρ_F over $\mathcal{K} \times \mathcal{R}$ where $\mathcal{K}$ is the key space of the pseudorandom function F , and $\mathcal{R}$ is the space of all possible values of r_i . The distribution of inputs given to function βPM is thus no longer perfectly uniform. Rather, it is the joint distribution of $\mu_{\mathcal{X}}$ and ρ_F . We call this distribution $\rho_{\beta PM}$.

As an unbounded Eve could distinguish $\text{Enc}(k, y_i)$ or $F(k, r_i)$ from uniform, both protocols rest on a computational assumption. In the case of HM-PRF, F is assumed to be a pseudorandom function (PRF) secure against any efficient polynomial-time adversary. HM-QCT conversely assumes IND-CCA security of the encryption scheme but only against an adversary completing in time $< t_{\text{comp}}$.

Overall, our HM-PRF protocol (1) is no longer dependant on the technological assumption of the NSM and (2) replaces the computational assumption of time-limited encryption with a PRF assumption. The security of the overall construction is therefore no longer parameterised on t_{comp} , as is required in the QCT model.⁴

5.1 Security Result

PROPOSITION 1 (EAVESDROPPER SECURITY OF HM-PRF). *If we accept that*

- (i) *the function F is a PRF, i.e., the derived $y_i = F(k, r_i)$ is distinguishable from a truly random y'_i sampled from a uniform distribution with a negligible probability of ϵ_{PRF} ; and*
- (ii) *the standard QKD assumptions hold (see Section 2.3),*

then, due to the hardness of the βPM problem (i.e., the one-way information complexity of βPM obeys the hidden-matching lower bound $\text{IC}_{\mu_{\beta PM}}^{\text{Lext}}(\beta PM, \epsilon) = \Omega(\sqrt{n})$ where ϵ is the average error and n is the length of each x_i), Eve's per-round guessing probability for any i.i.d. encoding attack and any $n \geq 4$ is bounded by

$$P_{\text{guess}}^{\text{Real}}(A | Z) \leq \frac{1}{2} + 2 \left(\sqrt[3]{\frac{1}{3}} + \sqrt{\frac{p}{3}} \right) + \epsilon_{\text{PRF}}, \quad (2)$$

³This entity is entrusted not to reveal the long-term secret to an outside party even if the result of computation using the secret is made known.

⁴It is best practice for the security of a scheme to be dependent on a security parameter adjustable by the user, not a parameter controlled by the adversary.

with p and q functions of n , the number of copies m and the matching density β , as given in Theorem 7.2.1 of [6].

PROOF. The argument is a two-step reduction:

- (1) we replace the PRF-derived inputs by truly random ones at the cost of ϵ_{PRF} ; and
- (2) we update the parameterised bounds given for HM-QCT to reflect the decrease in available information.

Step 1: Randomisation of the trusted setup. Consider two experiments, shown in Fig. 5, producing Eve's view of a protocol round. In the *real* experiment, value r_i is drawn from uniform distribution $\mu_{\mathcal{R}}$. The pre-shared key k is also chosen uniformly at random. Value y_i is generated by evaluating pseudorandom function $F(k, r_i)$, so y_i is chosen according to pseudorandom distribution $\mu_{\mathcal{Y}}$. In the *ideal* experiment, r_i is chosen in the same way as before, but y_i is instead drawn from the true distribution $\mu_{\mathcal{Y}}$ of possible y_i values, exactly as it was in HM-QCT. No k is generated.

In both experiments, a fixed procedure $\mathcal{D}$ is performed, receiving r_i and y_i as inputs. First, x_i is drawn from the uniform distribution $\mu_{\mathcal{X}}$. This is encoded quantumly m times as $|\psi_{x_i}\rangle^{\otimes m}$. Then, a_i is obtained by evaluating $a_i := \beta\text{PM}(x_i, y_i)$. Finally, upon receiving r_i and $|\psi_{x_i}\rangle^{\otimes m}$, Eve applies an arbitrary encoding attack and produces b_i as her guess of the negotiated key. $\mathcal{D}$ compares Eve's guess against its knowledge of the actual key bit, outputting 1 if $b_i = a_i$ and 0 otherwise.

As Eve never receives y_i or k , and r_i is always produced in the same manner, the only way Eve's success probability could differ between the two experiments is if there were a difference in the distribution of x_i , which there is not (the method for encoding x_i as quantum information is also identical in both experiments). Representing a_i as random variable A and the outcome of an optimal measurement performed on $|\psi_{x_i}\rangle^{\otimes m}$ as random variable Z , we can therefore say that

$$|p_{\text{guess}}^{\text{Real}}(A | Z) - p_{\text{guess}}^{\text{Ideal}}(A | Z)| \leq \epsilon_{\text{PRF}}. \quad (3)$$

Any gap larger than ϵ_{PRF} would make $\mathcal{D}$ able to distinguish between $F(k, r_i)$ and $\mu_{\mathcal{Y}}$ violating (i) of Proposition 1.

Step 2: The HM-QCT bounds with reduced information. Theorem 7.2.1 of [6] bounds Eve's success probability as follows:

$$P_{\text{guess}}(A | Y Z \Phi_{t_{\text{comp}}}(Q_{\text{in}})) \leq \frac{1}{2} + 2 \left(\sqrt[3]{-q} + \sqrt{\frac{p}{3}} \right) + \delta. \quad (4)$$

However, in the Ideal experiment, there is no additional information Y which is unlocked after time t_{comp} . Nor is any function of Q (i.e., Q_{in}) ever stored in quantum storage for t_{comp} , the result of which is denoted by $\Phi_{t_{\text{comp}}}(Q_{\text{in}})$. As such, the knowledge of Eve in the experiment is strictly less than if she were interacting with the original HM-QCT protocol. We can therefore state that

$$p_{\text{guess}}^{\text{Ideal}}(A | Z) \leq \frac{1}{2} + 2 \left(\sqrt[3]{-q} + \sqrt{\frac{p}{3}} \right). \quad (5)$$

Since we drop $\Phi_{t_{\text{comp}}}(Q_{\text{in}})$ on the LHS, we can also drop δ on the RHS per Theorem 6.4.1 of [6].

Conclusion. By combining (3) and (5), we obtain bound (2). $\square$




Figure 5: The HM-PRF security experiments.

5.2 Efficiency Result

PROPOSITION 2 (SECRET KEY RATE OF HM-PRF). *If Proposition 1 holds and we take its assumptions (i) and (ii) as true, then the achievable key rate is*

$$R_{\infty}^{\text{HM-PRF}} \geq \gamma_{\text{sift}} \left[-\log \left(p_{\text{guess}}^{\text{Real}}(A | Z) \right) - H_2(\text{QBER}) \right], \quad (6)$$

where γ_{sift} is the sifting factor as described in Section 7.2.2 of [6], $H_2(x)$ is the binary Shannon entropy of x , and QBER is the observed quantum bit-error rate.

The key-rate bound (6) follows from (2) as per Section 7.2.2 of [6]: the extractable secret fraction per sifted round is the min-entropy of A given Eve's view, less the cost of error correction.

6 Everlasting Security of HM-PRF

Protocols with *everlasting security* [10] remain secure even if the attacker becomes computationally unbounded upon completion. The HM-QCT protocol achieves a form of this but only after t_{comp} which, given $t_{\text{comp}} \gg t_{\text{coh}}$, may occur long after execution is finished.

By removing the encryption scheme and associated expiry event, HM-PRF eliminates fragile reliance on a gap between quantum memory and quantum computational capabilities. Doing away with the time-lock encryption and noisy-storage assumptions allows the security of the protocol to rest instead on a single static assumption: the existence of one-way functions plus a trusted setup. Guarding a long-term secret in certified hardware is auditable, and potentially more realistic, whereas the QCT model must predict a lasting ceiling on quantum technology.

The trade-off with eliminating both assumptions is that the secrecy of the negotiated key (shown in Section 5.1) is not everlasting. Because y_i is only pseudorandomly hidden and $|\psi_{x_i}\rangle$ can be stored indefinitely, past key bits could be revealed if the adversary were to break F after protocol execution. However, if one is willing to accept the noisy storage assumption, HM-PRF does achieve everlasting security after decoherence time t_{coh} .

7 Conclusion

We recast HM-QCT's guarantee as resistance to two *harvest now, decrypt later* strategies, showing its assumptions play connected roles: short-term encryption hides the matching in transit, while noisy storage closes the window this opens. In response, we proposed HM-PRF which derives the matching locally via a pseudorandom function, eliminating both the time-lock and (optionally) the storage assumption, while retaining the beyond-BB84 key rate.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council.

References

- [1] Ziv Bar-Yossef, T. S. Jayram, and Iordanis Kerenidis. 2004. Exponential separation of quantum and classical one-way communication complexity. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing* (Chicago, IL, USA) (STOC '04). Association for Computing Machinery, New York, NY, USA, 128–137. doi:10.1145/1007352.1007379
- [2] Charles H. Bennett and Gilles Brassard. 2014. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science* 560 (2014), 7–11. Theoretical Aspects of Quantum Cryptography – celebrating 30 years of BB84. doi:10.1016/j.tcs.2014.05.025
- [3] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. 2007. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing* (San Diego, California, USA) (STOC '07). Association for Computing Machinery, New York, NY, USA, 516–525. doi:10.1145/1250790.1250866
- [4] Robert König, Stephanie Wehner, and Jürg Wullschleger. 2012. Unconditional Security From Noisy Quantum Storage. *IEEE Transactions on Information Theory* 58, 3 (2012), 1962–1984. doi:10.1109/TIT.2011.2177772
- [5] Ueli M. Maurer. 1992. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology* 5, 1 (01 Jan 1992), 53–66. doi:10.1007/BF00191321
- [6] Francesco Mazzoncini. 2024. *Multimode Quantum Communications and Hybrid Cryptography*. Theses. Institut Polytechnique de Paris. HAL: tel-04656723. doi:10.70675/ac2f8f17z80faz4276za40fzfe7c124fb9fc
- [7] Francesco Mazzoncini, Balthazar Bauer, Peter Brown, and Romain Alléaume. 2025. Hybrid Quantum Cryptography from Communication Complexity. *Quantum* 9 (Sept. 2025), 1862. doi:10.22331/q-2025-09-24-1862
- [8] Stefano Pirandola, Riccardo Laurenza, Carlo Ottaviani, and Leonardo Banchi. 2017. Fundamental limits of repeaterless quantum communications. *Nature Communications* 8, 1 (26 Apr 2017), 15043. doi:10.1038/ncomms15043
- [9] Christopher Portmann and Renato Renner. 2022. Security in quantum cryptography. *Rev. Mod. Phys.* 94 (Jun 2022), 025008. Issue 2. doi:10.1103/RevModPhys.94.025008
- [10] Dominique Unruh. 2013. Everlasting Multi-party Computation. In *Advances in Cryptology – CRYPTO 2013*, Ran Canetti and Juan A. Garay (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 380–397. doi:10.1007/978-3-642-40084-1_22
- [11] Dominique Unruh. 2014. Revocable Quantum Timed-Release Encryption. In *Advances in Cryptology – EUROCRYPT 2014*, Phong Q. Nguyen and Elisabeth Oswald (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 129–146. doi:10.1007/978-3-642-55220-5_8
- [12] Nilesh Vyas and Romain Alléaume. 2020. Everlasting Secure Key Agreement with performance beyond QKD in a Quantum Computational Hybrid security model. arXiv:2004.10173 [quant-ph] doi:10.48550/arXiv.2004.10173
- [13] Stephanie Wehner, Christian Schaffner, and Barbara M. Terhal. 2008. Cryptography from Noisy Storage. *Phys. Rev. Lett.* 100 (Jun 2008), 220502. Issue 22. doi:10.1103/PhysRevLett.100.220502