

Sampling in a Quantum Population, and Applications

Niek J. Bouman and Serge Fehr

Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
{n.j.bouman,s.fehr}@cwi.nl

Abstract. We propose a framework for analyzing classical sampling strategies for estimating the Hamming weight of a large string from a few sample positions, when applied to a multi-qubit quantum system instead. The framework shows how to interpret the result of such a strategy and how to define its accuracy when applied to a quantum system. Furthermore, we show how the accuracy of any strategy relates to its accuracy in its classical usage, which is well understood for the important examples. We show the usefulness of our framework by using it to obtain new and simple security proofs for the following quantum-cryptographic schemes: BB84 quantum-key-distribution, and quantum oblivious-transfer from bit-commitment.

1 Introduction

Sampling allows to learn some information on a large population by merely looking at a comparably small number of individuals. For instance it is possible to predict the outcome of an election with very good accuracy by analyzing a relatively small subset of all the votes. In this work, we study sampling in a *quantum* population: we want to learn information about a large quantum state by measuring only a small part. Specifically, we investigate the quantum version of the following classical sampling problem (and of variants thereof). Given a bit-string $\mathbf{q} = (q_1, \dots, q_n) \in \{0, 1\}^n$ of length n , the task is to estimate the Hamming weight of \mathbf{q} by sampling and looking at only a few positions within \mathbf{q} . This classical sampling problem is well understood. For instance, the following particular *sampling strategy* works well: sample (with or without replacement) a linear number of positions uniformly at random, and compute an estimate for the Hamming weight of \mathbf{q} by scaling the Hamming weight of the sample accordingly; Hoeffding's bounds guarantee that the estimate is close to the real Hamming weight except with small probability. In particular, one can use a sampling strategy to *test* whether \mathbf{q} is close to the all-zero string $(0, \dots, 0)$ by looking only at a relatively small number of positions, where the test is accepted if and only if all the sample positions are zero, i.e., the estimated Hamming weight vanishes.

In the quantum version of the sampling problem from above, the string \mathbf{q} is replaced by a n -qubit quantum system A . Obviously, a sampling strategy from the classical can be *applied* to the quantum setting as well: pick a sample of

qubit positions within A , measure (in the computational basis) these sample positions, and compute the estimate as dictated by the sampling strategy from the observed values (i.e., typically, scale the Hamming weight of the measured sample appropriately). However, due to the special nature of quantum states, it is not clear and to the best of our knowledge so far not well understood, how to formally *interpret* the computed estimate. Simply extending the classical results in a straightforward way to the quantum setting does not work due to several reasons (e.g., one reason being that it is not clear what the Hamming weight of a quantum state should be).

In this work, we present a framework that addresses the above and fully characterizes the behavior of a classical sampling strategy when applied to a quantum population, i.e., to a n -qubit system or, more general, to n copies of an arbitrary “atomic” system. Our framework incorporates the following. First, we specify an abstract property on the state of A (after the measurements done by the sampling strategy), with the intended meaning that this is the property one should conclude from the outcome of the sampling strategy when applied to A . We also demonstrate that this property has useful consequences: specifically, that a suitable measurement will lead to a high-entropy outcome; this is handy in particular for quantum-cryptographic purposes. Then, we define a meaningful measure, sort of a “quantum error probability” (although technically speaking it is not a probability), that tells how reliable it is to conclude the specified property from the outcome of the sampling strategy. Finally, we show that for *any* sampling strategy, the quantum error probability of the strategy, as we define it, is bounded by the square-root of its classical error probability. This means that in order to understand how well a sampling strategy performs in the quantum setting, it suffices to analyze it in the classical setting, which is typically much simpler. Furthermore, for typical sampling strategies, like when picking the sample uniformly at random, there are well-known good bounds on the classical error probability.

We demonstrate the usefulness of our framework by means of two applications. Our applications do not constitute actual new results, but they provide new and simple(r) proofs for known results, both in the area of quantum cryptography. We take this as strong indication for the usefulness of the framework, and that the framework is likely to prove valuable in other applications as well.

The first application is to quantum key-distribution (QKD). We show how our framework for analyzing sampling strategies in the quantum setting leads to a conceptually very simple and easy-to-understand security proof for the BB84 QKD scheme.¹ The main idea behind the proof is that the checking phase of the BB84 scheme can be viewed as executing a specific sampling strategy. From the framework, it then follows that the raw key has high min-entropy from the adversary’s point of view, and the proof is concluded by applying the privacy amplification theorem.

¹ Actually, we prove security for an entanglement-based version of BB84 that implies security for the original BB84 scheme.

QKD schemes initially came without security proofs, and proving QKD schemes rigorously secure turned out to be an extremely challenging and subtle task. Nowadays, though, the security of QKD schemes is better understood, and we know of different ways of proving, say, BB84 secure, ranging from Shor and Preskill’s proof based on quantum error-correcting codes [9] to Renner’s approach using a quantum De Finetti theorem which allows to reduce security against general attacks to security against the much weaker class of so-called collective attacks [7]. Nonetheless, we think that our proof is interesting because of the following reasons. It provides an *explicit* expression for the security of the scheme, given in terms of an easy-to-compute function of the observed error-rate, the parameters of the code used to do error correction, and the number of extracted key-bits (and the parameters of the scheme). This is in contrast to most proofs in the literature which merely provide an asymptotic analysis. Furthermore, the proof is technically very accessible (e.g. compared to quantum-De-Finetti-based proofs) and as such for instance particularly well-suited for teaching. Finally, it does not require any “symmetrization of the qubits” (e.g. by applying a random permutation) from the protocol, and it gives a *direct* security proof, rather than a reduction to the security against collective attacks.

The second application is to quantum oblivious transfer (QOT). It is well known that QOT is not possible from scratch, but one can build a secure QOT scheme when given a bit-commitment (BC) primitive “for free”. Also for this cryptographic primitive, our framework allows for a simple and easy-to-understand security proof. Due to space restriction, this second application is only given in the full version [2] of this paper. The security of QOT (when given bit commitments) has also recently been rigorously proven in [4]. Although at the technical level similar ideas are used, our work distinguishes from [4] in that we introduce and rigorously study the concept of a general sampling strategy. This not only gives a nice framework and makes the security of QOT easier to understand, but it also opens the door for other applications (as we demonstrate).

We find it particularly interesting that with our framework, the protocols for QKD and QOT can be proven secure by means of very similar techniques, even though they implement fundamentally different cryptographic primitives, and are intuitively secure due to different reasons.

2 Notation, Terminology, and Some Tools

Strings and Hamming Weight. Throughout the paper, \mathcal{A} denotes some fixed finite alphabet with $0 \in \mathcal{A}$. It is safe to think of \mathcal{A} as $\{0, 1\}$, but our claims also hold for larger alphabets. For a string $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$ of arbitrary length $n \geq 0$, the *Hamming weight* of \mathbf{q} is defined as: $\text{wt}(\mathbf{q}) := |\{i \in [n] : q_i \neq 0\}|$, where $[n]$ is a short hand for $\{1, \dots, n\}$. The *relative Hamming weight* of \mathbf{q} is defined as $\omega(\mathbf{q}) := \text{wt}(\mathbf{q})/n$. By convention, the relative Hamming weight of the empty string \perp is set to $\omega(\perp) := 0$. For a subset $J \subset [n]$, we write $\mathbf{q}_J := (q_i)_{i \in J}$ for the restriction of \mathbf{q} to the positions $i \in J$.

Quantum Systems and States. We assume the reader to be familiar with the basic concepts of quantum information theory; we merely fix some specific terminology and notation here.

By default, we write \mathcal{H}_A for the state space of system A , and ρ_A for the density matrix and $|\varphi_A\rangle$ for the state vector (in case of a pure state) describing the state of A . To simplify language we are sometimes a bit sloppy in distinguishing between a quantum system, its state, and the state vector or density matrix describing the state. A *qubit* is a quantum system A with state space $\mathcal{H}_A = \mathbb{C}^2$. The *computational basis* $\{|0\rangle, |1\rangle\}$ (for a qubit) is given by $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, and the *Hadamard basis* by $H\{|0\rangle, |1\rangle\} = \{H|0\rangle, H|1\rangle\}$, where H denotes the 2-dimensional *Hadamard matrix* $H = 2^{-1/2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. The state space of an n -qubit system $A = A_1 \cdots A_n$ is given by $\mathcal{H}_A = (\mathbb{C}^2)^{\otimes n} = \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2$. For $\mathbf{x} = (x_1, \dots, x_n)$ and $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ in $\{0, 1\}^n$, we write $|\mathbf{x}\rangle$ for $|\mathbf{x}\rangle = |x_1\rangle \cdots |x_n\rangle$ and $H^\boldsymbol{\theta}$ for $H^\boldsymbol{\theta} = H^{\theta_1} \otimes \cdots \otimes H^{\theta_n}$, and thus $H^\boldsymbol{\theta}|\mathbf{x}\rangle$ for $H^\boldsymbol{\theta}|\mathbf{x}\rangle = H^{\theta_1}|x_1\rangle \cdots H^{\theta_n}|x_n\rangle$. Finally, we write $\{|0\rangle, |1\rangle\}^{\otimes n} = \{|\mathbf{x}\rangle : \mathbf{x} \in \{0, 1\}^n\}$ for the computational basis on an n -qubit system, and $H^\boldsymbol{\theta}\{|0\rangle, |1\rangle\}^{\otimes n} = \{H^\boldsymbol{\theta}|\mathbf{x}\rangle : \mathbf{x} \in \{0, 1\}^n\} = H^{\theta_1}\{|0\rangle, |1\rangle\} \otimes \cdots \otimes H^{\theta_n}\{|0\rangle, |1\rangle\}$ for the basis that is made up of the computational basis on the subsystems A_i with $\theta_i = 0$ and of the Hadamard basis on the subsystems A_i with $\theta_i = 1$. To simplify notation, we will sometimes abuse terminology and speak of the basis $\boldsymbol{\theta}$ when we actually mean $H^\boldsymbol{\theta}\{|0\rangle, |1\rangle\}^{\otimes n}$.

Measuring a system A in basis $\{|i\rangle\}_{i \in I}$, where $\{|i\rangle\}_{i \in I}$ is an orthonormal basis of \mathcal{H}_A , means applying the measurement described by the projectors $\{|i\rangle\langle i|\}_{i \in I}$, such that outcome $i \in I$ is observed with probability $p_i = \text{tr}(|i\rangle\langle i|\rho_A)$ (respectively $p_i = |\langle i|\varphi_A\rangle|^2$ in case of a pure state). If A is a subsystem of a bipartite system AB , then it means applying the measurement described by the projectors $\{|i\rangle\langle i| \otimes \mathbb{I}_B\}_{i \in I}$, where \mathbb{I}_B is the identity operator on \mathcal{H}_B .

We measure closeness of two states ρ and σ by their *trace distance*: $\Delta(\rho, \sigma) := \frac{1}{2} \text{tr}|\rho - \sigma|$, where for any square matrix M , $|M|$ denotes the positive-semi-definite square-root of $M^\dagger M$. For *pure* states $|\varphi\rangle$ and $|\psi\rangle$, the trace distance of the corresponding density matrices coincides with $\Delta(|\varphi\rangle\langle\varphi|, |\psi\rangle\langle\psi|) = \sqrt{1 - |\langle\varphi|\psi\rangle|^2}$. If the states of two systems A and B are ϵ -close, i.e. $\Delta(\rho_A, \rho_B) \leq \epsilon$, then A and B cannot be distinguished with advantage greater than ϵ ; in other words, A behaves exactly like B , except with probability ϵ .

Classical and Hybrid Systems (and States). Subsystem X of a bipartite quantum system XE is called *classical*, if the state of XE is given by a density matrix of the form $\rho_{XE} = \sum_{x \in \mathcal{X}} P_X(x) |x\rangle\langle x| \otimes \rho_E^x$, where \mathcal{X} is a finite set of cardinality $|\mathcal{X}| = \dim(\mathcal{H}_X)$, $P_X : \mathcal{X} \rightarrow [0, 1]$ is a probability distribution, $\{|x\rangle\}_{x \in \mathcal{X}}$ is some fixed orthonormal basis of \mathcal{H}_X , and ρ_E^x is a density matrix on \mathcal{H}_E for every $x \in \mathcal{X}$. Such a state, called *hybrid* or *cq* (for classical-quantum) state, can equivalently be understood as consisting of a *random variable* X with distribution P_X , taking on values in \mathcal{X} , and a system E that is in state ρ_E^x exactly when X takes on the value x . This formalism naturally extends to two (or more) classical systems X, Y etc.

If the state of XE satisfies $\rho_{XE} = \rho_X \otimes \rho_E$, where $\rho_X = \text{tr}_E(\rho_{XE}) = \sum_x P_X(x) |x\rangle\langle x|$ and $\rho_E = \text{tr}_X(\rho_{XE}) = \sum_x P_X(x) \rho_E^x$, then X is *independent*

of E , and thus no information on X can be obtained from system E . Moreover, if $\rho_{XE} = \frac{1}{|\mathcal{X}|} \mathbb{I}_X \otimes \rho_E$, where \mathbb{I}_X denotes the identity on \mathcal{H}_X , then X is *random-and-independent* of E . This is what is aimed for in quantum cryptography, when X represents a classical cryptographic key and E the adversary's potential quantum information on X .

It is not too hard to see that for two hybrid states ρ_{XE} and $\rho_{XE'}$ with the same (distribution of) X , the trace distance between ρ_{XE} and $\rho_{XE'}$ can be computed as $\Delta(\rho_{XE}, \rho_{XE'}) = \sum_x P_X(x) \Delta(\rho_E^x, \rho_{E'}^x)$.

Min-Entropy and Privacy Amplification. We make use of Renner's notion of the *conditional min-entropy* $H_{\min}(\rho_{XE}|E)$ of a system X conditioned on another system E [7]. Although the notion makes sense for arbitrary states, we restrict to hybrid states ρ_{XE} with classical X . If the hybrid state ρ_{XE} is clear from the context, we may write $H_{\min}(X|E)$ instead of $H_{\min}(\rho_{XE}|E)$. The formal definition is not very relevant to us, we merely rely on some elementary properties. For instance, the *chain rule* guarantees that $H_{\min}(X|YE) \geq H_{\min}(XY|E) - \log(|\mathcal{Y}|) \geq H_{\min}(X|E) - \log(|\mathcal{Y}|)$ for classical X and Y with respective ranges \mathcal{X} and \mathcal{Y} . Note that throughout this paper, \log denotes the *binary* logarithm (we write \ln for the *natural* logarithm). Furthermore, it holds that if E' is obtained from E by measuring (part of) E , then $H_{\min}(X|E') \geq H_{\min}(X|E)$.

Finally, we make use of Renner's privacy amplification theorem [8,7], as given below. Recall that a function $g : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ is called a *universal* (hash) function, if for the random variable R , uniformly distributed over \mathcal{R} , and for any distinct $x, y \in \mathcal{X}$: $\Pr[g(R, x) = g(R, y)] \leq 2^{-\ell}$.

Theorem 1 (Privacy amplification). *Let ρ_{XE} be a hybrid state with classical X . Let $g : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be a universal hash function, and let R be uniformly distributed over \mathcal{R} , independent of X and E . Then $K = g(R, X)$ satisfies*

$$\Delta(\rho_{KRE}, \frac{1}{|\mathcal{K}|} \mathbb{I}_K \otimes \rho_{RE}) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}(H_{\min}(X|E) - \ell)}.$$

Informally, Theorem 1 states that if X contains sufficiently more than ℓ bits of entropy when given E , then ℓ nearly random-and-independent bits can be extracted from X .

3 Sampling in a Classical Population

As a warm-up, and in order to study some useful examples and introduce some convenient notation, we start with the classical sampling problem, which is rather well-understood.

3.1 Sampling Strategies

Let $\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{A}^n$ be a string of given length n . We consider the problem of estimating the relative Hamming weight $\omega(\mathbf{q})$ by only looking at a substring

\mathbf{q}_t of \mathbf{q} , for a small subset $t \subset [n]$.² Actually, we are interested in the equivalent problem of estimating the relative Hamming weight $\omega(\mathbf{q}_{\bar{t}})$ of the *remaining* string $\mathbf{q}_{\bar{t}}$, where \bar{t} is the complement $\bar{t} = [n] \setminus t$ of t .³ A canonical way to do so would be to sample a uniformly random subset (say, of a certain small size) of positions, and compute the relative Hamming weight of the sample as estimate. Very generally, we allow any strategy that picks a subset $t \subset [n]$ according to some probability distribution and computes the estimate for $\omega(\mathbf{q}_{\bar{t}})$ as some (possibly randomized) function of t and \mathbf{q}_t , i.e., as $f(t, \mathbf{q}_t, s)$ for a *seed* s that is sampled according to some probability distribution from a finite set \mathcal{S} . This motivates the following formal definition.

Definition 1 (Sampling strategy). A sampling strategy Ψ consists of a triple (P_T, P_S, f) , where P_T is a distribution over the subsets of $[n]$, P_S is a (independent) distribution over a finite set \mathcal{S} , and f is a function

$$f : \{(t, v) : t \subset [n], v \in \mathcal{A}^{|t|}\} \times \mathcal{S} \rightarrow \mathbb{R}.$$

We stress that a sampling strategy Ψ , as defined here, specifies how to choose the sample subset as well as how to compute the estimate from the sample (thus a more appropriate but lengthy name would be a “sample-and-estimate strategy”).

Remark 1. By definition, the choice of the seed s is specified to be independent of t , i.e., $P_{TS} = P_T P_S$. Sometimes, however, it is convenient to allow s to depend on t . We can actually do so without contradicting Definition 1. Namely, to comply with the independence requirement, we would simply choose a (typically huge) “container” seed that contains a seed for every possible choice of t , each one chosen with the corresponding distribution, and it is then part of f ’s task, when given t , to select the seed that is actually needed out of the container seed.⁴

A sampling strategy Ψ can obviously also be used to *test* if \mathbf{q} (or actually $\mathbf{q}_{\bar{t}}$) is close to the all-zero string $0 \cdots 0$: compute the estimate for $\omega(\mathbf{q}_{\bar{t}})$ as dictated by Ψ , and *accept* if the estimate vanishes and else *reject*.

We briefly discuss a few example sampling strategies (two more examples, including random sampling *with* replacement, can be found in the full version [2]).

Example 1 (Random sampling without replacement). In random sampling without replacement, k *distinct* indices i_1, \dots, i_k within $[n]$ are chosen uniformly at random, where k is some parameter, and the relative Hamming weight of $\mathbf{q}_{\{i_1, \dots, i_k\}}$ is used as estimate for $\omega(\mathbf{q}_{\bar{t}})$. Formally, this sampling strategy is given

² More generally, we may consider the problem of estimating the Hamming *distance* of \mathbf{q} to some arbitrary *reference string* \mathbf{q}_o ; but this can obviously be done simply by estimating the Hamming weight of $\mathbf{q}' = \mathbf{q} - \mathbf{q}_o$.

³ In our applications, the sampled positions within \mathbf{q} will be *discarded*, and thus we will be interested merely in the remaining positions.

⁴ Alternatively, we could simply drop the independence requirement in Definition 1; however, we feel it is conceptually easier to think of the seed as being independently chosen.

by $\Psi = (P_T, P_S, f)$ where $P_T(t) = 1/\binom{n}{k}$ if $|t| = k$ and else $P_T(t) = 0$, $\mathcal{S} = \{\perp\}$ and thus $P_S(\perp) = 1$, and $f(t, \mathbf{q}_t, \perp) = f(t, \mathbf{q}_t) = \omega(\mathbf{q}_t)$. \diamond

Example 2 (Uniformly random subset sampling). The sample set t is chosen as a uniformly random subset of $[n]$, and the estimate is computed as the relative Hamming weight of the sample \mathbf{q}_t : $P_T(t) = 1/2^n$ for any $t \subseteq [n]$, and $\mathcal{S} = \{\perp\}$ and $f(t, \mathbf{q}_t, \perp) = f(t, \mathbf{q}_t) = \omega(\mathbf{q}_t)$. \diamond

The next example is a somewhat unnatural and in some sense non-optimal sampling strategy, but it will be of use for the QKD proof in Section 5.

Example 3 (Pairwise one-out-of-two sampling, using only part of the sample). For this example, it is convenient to consider the index set from which the subset t is chosen, to be of the form $[n] \times \{0, 1\}$. Namely, we consider the string $\mathbf{q} \in \mathcal{A}^{2n}$ to be indexed by *pairs* of indices, $\mathbf{q} = (q_{ij})$, where $i \in [n]$ and $j \in \{0, 1\}$; in other words, we consider \mathbf{q} to consist of n pairs (q_{i0}, q_{i1}) . The subset $t \subseteq [n] \times \{0, 1\}$ is chosen as $t = \{(1, j_1), \dots, (n, j_n)\}$ where every j_k is picked independently at random in $\{0, 1\}$. In other words, t selects one element from each pair (q_{i0}, q_{i1}) . Furthermore, the estimate for $\omega(\mathbf{q}_{\bar{t}})$ is computed from \mathbf{q}_t as $f(t, \mathbf{q}_t, s) = \omega(\mathbf{q}_s)$ where the seed s is a random subset $s \subseteq t$ of size k . \diamond

3.2 The Error Probability

We formally define a measure that captures for a given sampling strategy how well it performs, i.e., with what probability the estimate, $f(t, \mathbf{q}_t, s)$, is how close to the real value, $\omega(\mathbf{q}_{\bar{t}})$. For the definition and for later purposes, it will be convenient to introduce the following notation. For a given sampling strategy $\Psi = (P_T, P_S, f)$, consider arbitrary but fixed choices for the subset $t \subseteq [n]$ and the seed $s \in \mathcal{S}$ with $P_T(t) > 0$ and $P_S(s) > 0$. Furthermore, fix an arbitrary $\delta > 0$. Define $B_{t,s}^\delta(\Psi) \subseteq \mathcal{A}^n$ as

$$B_{t,s}^\delta(\Psi) := \{\mathbf{b} \in \mathcal{A}^n : |\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta\},$$

i.e., as the set of all strings \mathbf{q} for which the estimate is δ -close to the real value, assuming that subset t and seed s have been used. To simplify notation, if Ψ is clear from the context, we simply write $B_{t,s}^\delta$ instead of $B_{t,s}^\delta(\Psi)$. By replacing the specific values t and s by the corresponding (independent) random variables T and S , with distributions P_T and P_S , respectively, we obtain the *random variable* $B_{T,S}^\delta$, whose range consists of subsets of \mathcal{A}^n . By means of this random variable, we now define the *error probability* of a sampling strategy as follows.

Definition 2 (Error probability). *The (classical) error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is defined as the following value, parametrized by $0 < \delta < 1$:*

$$\varepsilon_{\text{class}}^\delta(\Psi) := \max_{\mathbf{q} \in \mathcal{A}^n} \Pr\left[\mathbf{q} \notin B_{T,S}^\delta(\Psi)\right].$$

By definition of the error probability, it is guaranteed that for any string $\mathbf{q} \in \mathcal{A}^n$, the estimated value is δ -close to the real value except with probability at most $\varepsilon_{\text{class}}^\delta(\Psi)$. When used as a sampling strategy to test closeness to the all-zero string, $\varepsilon_{\text{class}}^\delta(\Psi)$ determines the probability of accepting even though $\mathbf{q}_{\bar{T}}$ is “not close” to the all-zero string, in the sense that its relative Hamming weight exceeds δ . Whenever Ψ is clear from the context, we will write $\varepsilon_{\text{class}}^\delta$ instead of $\varepsilon_{\text{class}}^\delta(\Psi)$.

Below, we analyze the error probability of the sampling strategy discussed in Example 3, because this sampling strategy is used in our QKD security proof (Section 5). The error probabilities of the other examples can be found in the full version of this paper [2]. To bound the error probability, we use Hoeffding’s inequality [5]. The following theorem summarizes this inequality, tailored to our needs.

Theorem 2 (Hoeffding). *Let $\mathbf{b} \in \{0, 1\}^n$ be a bit string with relative Hamming weight $\mu = \omega(\mathbf{b})$. Let the random variables X_1, X_2, \dots, X_k be obtained by sampling k random entries from \mathbf{b} with replacement, i.e., the X_i ’s are independent and $\Pr_{X_i}(1) = \mu$. Furthermore, let the random variables Y_1, Y_2, \dots, Y_k be obtained by sampling k random entries from \mathbf{b} without replacement. Then, for any $\delta > 0$, the random variables $\bar{X} := \frac{1}{k} \sum_i X_i$ and $\bar{Y} := \frac{1}{k} \sum_i Y_i$ satisfy*

$$\Pr[|\bar{Y} - \mu| \geq \delta] \leq \Pr[|\bar{X} - \mu| \geq \delta] \leq 2 \exp(-2\delta^2 k).$$

Error Probability of Example 3. For $\mathcal{A} = \{0, 1\}$, a bound on the error probability $\varepsilon_{\text{class}}^\delta$ is obtained as follows. Let \mathbf{q} be arbitrary, indexed as discussed earlier. First, we show that $\omega(\mathbf{q}_{\bar{T}})$ is likely to be close to $\omega(\mathbf{q}_T)$. For this, consider the pairs (q_{i0}, q_{i1}) for which $q_{i0} \neq q_{i1}$. Let there be ℓ such pairs (where obviously $\ell \leq n$). We denote the restrictions of \mathbf{q}_T and $\mathbf{q}_{\bar{T}}$ to these indices i with $q_{i0} \neq q_{i1}$ by $\tilde{\mathbf{q}}_T$ and $\tilde{\mathbf{q}}_{\bar{T}}$, respectively. It is easy to see that $\text{wt}(\tilde{\mathbf{q}}_T) + \text{wt}(\tilde{\mathbf{q}}_{\bar{T}}) = \ell$. It follows that for any $\epsilon > 0$ we have

$$\begin{aligned} \Pr[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_T)| \geq \epsilon] &= \Pr[|\text{wt}(\mathbf{q}_T) - \text{wt}(\mathbf{q}_{\bar{T}})| \geq n\epsilon] \\ &= \Pr[|\text{wt}(\tilde{\mathbf{q}}_T) - \text{wt}(\tilde{\mathbf{q}}_{\bar{T}})| \geq n\epsilon] = \Pr[|2\text{wt}(\tilde{\mathbf{q}}_T) - \ell| \geq n\epsilon] \\ &\leq 2 \exp\left(-2\left(\frac{n\epsilon}{2\ell}\right)^2 \ell\right) = 2 \exp\left(-\frac{n\epsilon^2}{2} \cdot \frac{n}{\ell}\right) \leq 2 \exp\left(-\frac{1}{2}\epsilon^2 n\right), \end{aligned}$$

where the third equality follows from replacing $\text{wt}(\tilde{\mathbf{q}}_{\bar{T}}$) by $\ell - \text{wt}(\tilde{\mathbf{q}}_T)$, and the first inequality follows from Hoeffding’s inequality (as each entry of $\text{wt}(\tilde{\mathbf{q}}_T)$ is 0 with independent probability $\frac{1}{2}$).

Furthermore, for any $\gamma > 0$ we have the following relation involving \mathbf{q}_S :

$$\Pr[|\omega(\mathbf{q}_T) - \omega(\mathbf{q}_S)| \geq \gamma] \leq 2 \exp(-2k\gamma^2),$$

which follows from directly applying Hoeffding’s inequality. Applying the union bound and letting $\delta = \epsilon + \gamma$, we obtain

$$\begin{aligned} \varepsilon_{\text{class}}^\delta &= \Pr[|\omega(\mathbf{q}_{\bar{T}}) - \omega(\mathbf{q}_S)| \geq \delta] < 2 \min_{\epsilon \in (0, \delta)} [\exp(-\frac{1}{2}\epsilon^2 n) + \exp(-2k(\delta - \epsilon)^2)] \\ &\leq 4 \exp\left(-\frac{2kn\delta^2}{(2\sqrt{k} + \sqrt{n})^2}\right) \leq 4 \exp\left(-\frac{1}{3}\delta^2 k\right), \end{aligned}$$

where the last line follows from choosing ϵ such that the two exponents coincide, and from doing some simplifications while assuming $k \leq n/2$.

4 Sampling in a Quantum Population

We now want to study the behavior of a sampling strategy when applied to a quantum population. More specifically, let $A = A_1 \cdots A_n$ be an n -partite quantum system, where the state space of each system A_i equals $\mathcal{H}_{A_i} = \mathbb{C}^d$ with $d = |\mathcal{A}|$, and let $\{|a\rangle\}_{a \in \mathcal{A}}$ be a fixed orthonormal basis of \mathbb{C}^d . We allow A to be entangled with some additional system E with arbitrary finite-dimensional state-space \mathcal{H}_E . We may assume the joint state of AE to be pure, and as such be given by a state vector $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$; if not, then it can be purified by increasing the dimension of \mathcal{H}_E .

Similar to the classical sampling problem of testing closeness to the all-zero string, we can consider here the problem of testing if the state of A is close to the all-zero *reference state* $|\varphi_A^\circ\rangle = |0\rangle \cdots |0\rangle$ by looking at, which here means *measuring*, only a few of the subsystems of A . More generally, we will be interested in the sampling problem of estimating the ‘‘Hamming weight of the state of A ’’, although it is not clear at the moment what this should mean. Actually, like in the classical case, we are interested in testing closeness to the all-zero state, respectively estimating the Hamming weight, of the *remaining subsystems* of A .

It is obvious that a sampling strategy $\Psi = (P_T, P_S, f)$ can be applied in a straightforward way to the setting at hand: sample t according to P_T , measure the subsystems A_i with $i \in t$ in basis $\{|a\rangle\}_{a \in \mathcal{A}}$ to observe $\mathbf{q}_t \in \mathcal{A}^{|t|}$, and compute the estimate as $f(t, \mathbf{q}_t, s)$ for s chosen according to P_S (respectively, for testing closeness to the all-zero state, accept or reject depending on the value of the estimate). However, it is a-priori *not* clear, how to interpret the outcome. Measuring a random subset of the subsystems of A and observing 0 all the time indeed seems to suggest that the original state of A , and thus the remaining subsystems, must be in some sense close to the all-zero state; but in what formal sense is this true? And what can we conclude about the remaining state in case of a general sampling strategy for estimating the (relative) Hamming weight?

We give in this section a rigorous analysis of sampling strategies when applied to a n -partite quantum system A . Our analysis completely answers above questions. Later in the paper, we demonstrate the usefulness of our analysis of sampling strategies for studying and analyzing quantum-cryptographic schemes.

4.1 Analyzing Sampling Strategies in the Quantum Setting

We start by suggesting the property on the remaining subsystems of A that one should expect to be able to conclude from the outcome of a sampling strategy. A somewhat natural approach is as follows.

Definition 3. *For system AE , and similarly for any subsystem of A , we say that the state $|\varphi_{AE}\rangle$ of AE has relative Hamming weight β within A if it is of the form $|\varphi_{AE}\rangle = |\mathbf{b}\rangle|\varphi_E\rangle$ with $\mathbf{b} \in \mathcal{A}^n$ and $\omega(\mathbf{b}) = \beta$.*

Now, given the outcome $f(t, \mathbf{q}_t, s)$ of a sampling strategy when applied to A , we want to be able to conclude that, up to a small error, the state of the remaining subsystem $A_{\bar{t}}E$ is a *superposition* of states with relative Hamming weight close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$. To analyze this, we extend some of the notions introduced in the classical setting. Recall the definition of $B_{t,s}^\delta$, consisting of all strings $\mathbf{b} \in \mathcal{A}^n$ with $|\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta$. By slightly abusing notation, we extend this notion to the quantum setting and write

$$\text{span}(B_{t,s}^\delta) := \text{span}(\{|\mathbf{b}\rangle : \mathbf{b} \in B_{t,s}^\delta\}) = \text{span}(\{|\mathbf{b}\rangle : |\omega(\mathbf{b}_{\bar{t}}) - f(t, \mathbf{b}_t, s)| < \delta\}).$$

Note that if the state $|\varphi_{AE}\rangle$ of AE happens to be in $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$ for some t and s , and if exactly these t and s are chosen when applying the sampling strategy to A , then *with certainty* the state of $A_{\bar{t}}E$ (after the measurement) is in a superposition of states with relative Hamming weight δ -close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$, regardless of the measurement outcome \mathbf{q}_t .

Next, we want to extend the notion of error probability (Definition 2) to the quantum setting. For this, we consider the *hybrid* system $TSAE$, consisting of the classical random variables T and S with distribution $P_{TS} = P_T P_S$, describing the choices of t and s , respectively, and of the actual quantum systems A and E . The state of $TSAE$ is given by

$$\rho_{TSAE} = \sum_{t,s} P_{TS}(t, s) |t, s\rangle\langle t, s| \otimes |\varphi_{AE}\rangle\langle\varphi_{AE}|.$$

Note that TS is independent of AE : $\rho_{TSAE} = \rho_{TS} \otimes \rho_{AE}$; indeed, in a sampling strategy t and s are chosen independently of the state of AE . We compare this *real* state of $TSAE$ with an *ideal* state which is of the form

$$\begin{aligned} \tilde{\rho}_{TSAE} &= \sum_{t,s} P_{TS}(t, s) |t, s\rangle\langle t, s| \otimes |\tilde{\varphi}_{AE}^{ts}\rangle\langle\tilde{\varphi}_{AE}^{ts}| \quad \text{with} \\ &|\tilde{\varphi}_{AE}^{ts}\rangle \in \text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E \quad \forall t, s \end{aligned} \tag{1}$$

for some given $\delta > 0$. Thus, T and S have the same distribution as in the real state, but here we allow AE to depend on T and S , and for each particular choice t and s for T and S , respectively, we require the state of AE to be in $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$. Thus, in an “ideal world” where the state of the hybrid system $TSAE$ is given by $\tilde{\rho}_{TSAE}$, it holds *with certainty* that the state $|\psi_{A_{\bar{t}}E}\rangle$ of $A_{\bar{t}}E$, after having measured A_t and having observed \mathbf{q}_t , is in a superposition of states with relative Hamming weight δ -close to $\beta := f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$. We now define the quantum error probability of a sampling strategy by looking at how far away the closest ideal state $\tilde{\rho}_{TSAE}$ is from the real state ρ_{TSAE} .

Definition 4 (Quantum error probability). *The quantum error probability of a sampling strategy $\Psi = (P_T, P_S, f)$ is defined as the following value, parametrized by $0 < \delta < 1$:*

$$\varepsilon_{\text{quant}}^\delta(\Psi) = \max_{\mathcal{H}_E} \max_{|\varphi_{AE}\rangle} \min_{\tilde{\rho}_{TSAE}} \Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}),$$

where the first max is over all finite-dimensional state spaces \mathcal{H}_E , the second max is over all state vectors $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, and the min is over all ideal states $\tilde{\rho}_{TSAE}$ as in (1).⁵

As with $B_{t,s}^\delta$ and $\varepsilon_{\text{class}}^\delta$, we simply write $\varepsilon_{\text{quant}}^\delta$ when Ψ is clear from the context. We stress the meaningfulness of the definition: it guarantees that on average over the choice of t and s , the state of $A_{\bar{t}}E$ is $\varepsilon_{\text{quant}}^\delta$ -close to a superposition of states with Hamming weight δ -close to $f(t, \mathbf{q}_t, s)$ within $A_{\bar{t}}$, and as such it behaves like a superposition of such states, except with probability $\varepsilon_{\text{quant}}^\delta$. We will argue below and demonstrate in the subsequent sections that being (close to) a superposition of states with given approximate (relative) Hamming weight has some useful consequences.

Remark 2. Similarly to footnote 2, also here the results of the section immediately generalize from the all-zero reference state $|0\rangle \cdots |0\rangle$ to an arbitrary reference state $|\varphi_A^\circ\rangle$ of the form $|\varphi_A^\circ\rangle = U_1|0\rangle \otimes \cdots \otimes U_n|0\rangle$ for unitary operators U_i acting on \mathbb{C}^d . Indeed, the generalization follows simply by a suitable change of basis, defined by the U_i 's. Or, in the special case where $\mathcal{A} = \{0, 1\}$ and

$$|\varphi_A^\circ\rangle = H^{\hat{\theta}}|\hat{\mathbf{x}}\rangle = H^{\hat{\theta}_1}|\hat{x}_1\rangle \otimes \cdots \otimes H^{\hat{\theta}_n}|\hat{x}_n\rangle$$

for a fixed reference basis $\hat{\theta} \in \{0, 1\}^n$ and a fixed reference string $\hat{\mathbf{x}} \in \{0, 1\}^n$, we can, alternatively, replace in the definitions and results the computational by the Hadamard basis whenever $\hat{\theta}_i = 1$, and speak of the (relative) Hamming distance to $\hat{\mathbf{x}}$ rather than of the (relative) Hamming weight.

4.2 The Quantum vs. the Classical Error Probability

It remains to discuss how difficult it is to actually *compute* the quantum error probability for given sampling strategies, and how the *quantum* error probability $\varepsilon_{\text{quant}}^\delta$ relates to the corresponding *classical* error probability $\varepsilon_{\text{class}}^\delta$. To this end, we show the following simple relationship between $\varepsilon_{\text{quant}}^\delta$ and $\varepsilon_{\text{class}}^\delta$.

Theorem 3. *For any sampling strategy Ψ and for any $\delta > 0$:*

$$\varepsilon_{\text{quant}}^\delta(\Psi) \leq \sqrt{\varepsilon_{\text{class}}^\delta(\Psi)}.$$

As a consequence of this theorem, it suffices to analyze a sampling strategy in the classical setting, which is much easier, in order to understand how it behaves in the quantum setting. In particular, sampling strategies that are known to behave well in the classical setting, like Example 1 to 3, are also automatically guaranteed to behave well in the quantum setting. We will use this for our applications.

Our bound on $\varepsilon_{\text{quant}}^\delta$ is in general tight. Indeed, in [2] we show tightness for an explicit class of sampling strategies, which e.g. includes Example 1 and Example 3. Here, we just mention the tightness result.

⁵ It is not too hard to see, in particular after having gained some more insight via the proof of Theorem 3 below, that these min and max exist.

Proposition 1. *There exist natural sampling strategies for which the inequality in Theorem 3 is an equality.*

Proof (of Theorem 3). We need to show that for any $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$, with arbitrary \mathcal{H}_E , there exists a suitable ideal state $\tilde{\rho}_{TSAE}$ with $\Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) \leq (\varepsilon_{\text{class}}^\delta)^{1/2}$. We construct $\tilde{\rho}_{TSAE}$ as in (1), where the $|\tilde{\varphi}_{AE}^{ts}\rangle$'s are defined by the following decomposition.

$$|\varphi_{AE}\rangle = \langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle |\tilde{\varphi}_{AE}^{ts}\rangle + \langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle |\tilde{\varphi}_{AE}^{ts\perp}\rangle,$$

with $|\tilde{\varphi}_{AE}^{ts}\rangle \in \text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$, $|\tilde{\varphi}_{AE}^{ts\perp}\rangle \in \text{span}(B_{t,s}^\delta)^\perp \otimes \mathcal{H}_E$ and $|\langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle|^2 + |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2 = 1$. In other words, $|\tilde{\varphi}_{AE}^{ts}\rangle$ is obtained as the re-normalized projection of $|\varphi_{AE}\rangle$ into $\text{span}(B_{t,s}^\delta) \otimes \mathcal{H}_E$. Note that $|\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2$ equals the probability $\Pr[\mathbf{Q} \notin B_{t,s}^\delta]$, where the random variable \mathbf{Q} is obtained by measuring subsystem A of $|\varphi_{AE}\rangle$ in basis $\{|a\rangle\}_{a \in A}$. Furthermore,

$$\begin{aligned} \sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2 &= \sum_{t,s} P_{TS}(t,s) \Pr[\mathbf{Q} \notin B_{t,s}^\delta] = \Pr[\mathbf{Q} \notin B_{T,S}^\delta] \\ &= \sum_{\mathbf{q}} P_{\mathbf{Q}}(\mathbf{q}) \Pr[\mathbf{q} \notin B_{T,S}^\delta], \end{aligned}$$

where by definition of $\varepsilon_{\text{class}}^\delta$, the latter is upper bounded by $\varepsilon_{\text{class}}^\delta$. From elementary properties of the trace distance, and using Jensen's inequality, we can now conclude that

$$\begin{aligned} \Delta(\rho_{TSAE}, \tilde{\rho}_{TSAE}) &= \sum_{t,s} P_{TS}(t,s) \Delta(|\varphi_{AE}\rangle\langle\varphi_{AE}|, |\tilde{\varphi}_{AE}^{ts}\rangle\langle\tilde{\varphi}_{AE}^{ts}|) \\ &= \sum_{t,s} P_{TS}(t,s) \sqrt{1 - |\langle \tilde{\varphi}_{AE}^{ts} | \varphi_{AE} \rangle|^2} = \sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle| \\ &\leq \sqrt{\sum_{t,s} P_{TS}(t,s) |\langle \tilde{\varphi}_{AE}^{ts\perp} | \varphi_{AE} \rangle|^2} \leq \sqrt{\varepsilon_{\text{class}}^\delta}, \end{aligned}$$

which was to be shown. □

As a side remark, we point out that the particular ideal state $\tilde{\rho}_{TSAE}$ constructed in the proof minimizes the distance to ρ_{TSAE} ; this follows from the so-called Hilbert projection theorem.

4.3 Superpositions with a Small Number of Terms

We give here some argument why being (close to) a superposition of states with a given approximate Hamming weight may be a useful property in the analyses of quantum-cryptographic schemes. For simplicity, and since this will be the case in our applications, we now restrict to the binary case where $\mathcal{A} = \{0, 1\}$. Our argument is based on the following lemma, which follows immediately from Lemma 3.1.13 in [7]; we give a direct proof of Lemma 1 in the full version [2].

Informally, it states that measuring (part of) a *superposition* of a small number of orthogonal states produces a similar amount of uncertainty as when measuring the *mixture* of these orthogonal states.

Lemma 1. *Let A and E be arbitrary quantum systems, let $\{|i\rangle\}_{i \in I}$ and $\{|w\rangle\}_{w \in \mathcal{W}}$ be orthonormal bases of \mathcal{H}_A , and let $|\varphi_{AE}\rangle$ and ρ_{AE}^{mix} be of the form*

$$|\varphi_{AE}\rangle = \sum_{i \in J} \alpha_i |i\rangle |\varphi_E^i\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E \quad \text{and} \quad \rho_{AE}^{\text{mix}} = \sum_{i \in J} |\alpha_i|^2 |i\rangle\langle i| \otimes |\varphi_E^i\rangle\langle \varphi_E^i|$$

for some subset $J \subseteq I$. Let ρ_{WE} and ρ_{WE}^{mix} describe the hybrid systems obtained by measuring subsystem A of $|\varphi_{AE}\rangle$ and ρ_{AE}^{mix} , respectively, in basis $\{|w\rangle\}_{w \in \mathcal{W}}$ to observe outcome W . Then, $H_{\min}(\rho_{WE}|E) \geq H_{\min}(\rho_{WE}^{\text{mix}}|E) - \log |J|$.

We apply Lemma 1 to an n -qubit system A where $|\varphi_{AE}\rangle$ is a superposition of states with relative Hamming weight δ -close to β within A .⁶

$$|\varphi_{AE}\rangle = \sum_{\substack{\mathbf{b} \in \{0,1\}^n \\ |\omega(\mathbf{b}) - \beta| \leq \delta}} |\mathbf{b}\rangle |\varphi_E^{\mathbf{b}}\rangle.$$

It is well known that $|\{\mathbf{b} \in \{0,1\}^n : |\omega(\mathbf{b}) - \beta| \leq \delta\}| \leq 2^{h(\beta+\delta)n}$ for $\beta + \delta \leq \frac{1}{2}$, where $h(p) := -(p \log(p) + (1-p) \log(1-p))$ denotes the binary entropy function.

Since measuring qubits within a state $|\mathbf{b}\rangle$ in the *Hadamard* basis produces uniformly random bits, we can conclude the following.

Corollary 1. *Let A be an n -qubit system, let the state $|\varphi_{AE}\rangle$ of AE be a superposition of states with relative Hamming weight δ -close to β within A , where $\delta + \beta \leq \frac{1}{2}$, and let the random variable \mathbf{X} be obtained by measuring A in basis $H^\theta\{|0\rangle, |1\rangle\}^{\otimes n}$ for $\theta \in \{0,1\}^n$. Then*

$$H_{\min}(\mathbf{X}|E) \geq \text{wt}(\theta) - h(\beta + \delta)n.$$

Consider now the following quantum-cryptographic setting. Bob prepares and hands over to Alice an n -qubit quantum system A , which ought to be in state $|\varphi_A^\circ\rangle = |0\rangle \cdots |0\rangle$. However, since Bob might be dishonest, the state of A could be anything, even entangled with some system E controlled by Bob. Our results now imply the following: Alice can apply a suitable sampling strategy to convince herself that the joint state of the remaining subsystem of A and of E is (close to) a superposition of states with bounded relative Hamming weight. From Corollary 1, we can then conclude that with respect to the min-entropy of the measurement outcome, the state of A behaves similarly to the case where Bob honestly prepares A to be in state $|\varphi_A^\circ\rangle$. By Remark 2, i.e., by doing a suitable change of basis, the same holds if $|\varphi_A^\circ\rangle = H^{\hat{\theta}}|\hat{\mathbf{x}}\rangle$ for arbitrary fixed $\hat{\theta}, \hat{\mathbf{x}} \in \{0,1\}^n$, where $\text{wt}(\theta)$ is replaced by the Hamming distance between θ and $\hat{\theta}$. We will make use of this in the application in the upcoming section.

⁶ System A considered here corresponds to the subsystem $A_{\bar{t}}$ in the previous section, after having measured A_t of the ideal state.

5 Application: Quantum Key Distribution (QKD)

In quantum key distribution (QKD), Alice and Bob want to agree on a secret key in the presence of an adversary Eve. Alice and Bob are assumed to be able to communicate over a quantum channel and over an authenticated classical channel. Eve may eavesdrop the classical channel (but not insert or modify messages), and she has full control over the quantum channel. The first and still most prominent QKD scheme is the famous BB84 QKD scheme due to Bennett and Brassard [1].

In this section, we show how our sampling-strategy framework leads to a simple security proof for the BB84 QKD scheme. Proving QKD schemes rigorously secure is a highly non-trivial task, and as such our new proof nicely demonstrates the power of the sampling-strategy framework. Furthermore, our new proof has some nice features. For instance, it allows us to explicitly state (a bound on) the error probability of the QKD scheme for any given choices of the parameters. Additionally, our proof does not seem to take unnecessary detours or to make use of “loose bounds”, and therefore we feel that the bound on the error probability we obtain is rather tight (although we have no formal argument to support this). Our proof strategy can also be applied to other QKD schemes that are based on the BB84 encoding. For example, Lo *et al.*'s QKD scheme⁷ [6] can be proven secure by following exactly our proof, except that one needs to analyze a slightly different sampling strategy. On the other hand, it is yet unknown whether our framework can be used to prove e.g. the six-state QKD protocol [3] secure.

As a matter of fact, the QKD scheme we analyze is an entanglement-based version of the BB84 scheme. However, it is very well known and not too hard to show that security of the entanglement-based version implies security of the original BB84 QKD scheme.

The entanglement-based QKD scheme, QKD, is parametrized by the total number n of qubits sent in the protocol and the number k of qubits used to estimate the error rate of the quantum channel (where we require $k \leq n/2$). Additional parameters, which are determined during the course of the protocol, are the observed error rate β and the number $\ell \in \mathbb{N} \cup \{0\}$ of extracted key bits. QKD makes use of a universal hash function $g : \mathcal{R} \times \{0, 1\}^{n-k} \rightarrow \{0, 1\}^\ell$ and a linear binary error correcting code of length $n - k$ that allows to correct up to a β' -fraction of errors (except maybe with negligible probability) for some $\beta' > \beta$. The choice of how much β' exceeds β is a trade-off between keeping the probability that Alice and Bob end up with different keys small and increasing the size of the extractable key. We will write m for the bit size of the syndrome of this error-correcting code. Protocol QKD can be found below.

It is not hard to see that $\mathbf{k} = \hat{\mathbf{k}}$ except with negligible probability (in n). Furthermore, if no Eve interacts with the quantum communication in the qubit distribution phase then $\mathbf{x} = \mathbf{y}$ in case of a noise-free quantum channel, or more generally, $\omega(\mathbf{x} - \mathbf{y}) \approx \phi$ in case the quantum channel is noisy and introduces

⁷ In this scheme, Alice and Bob bias the choice of the bases so that they measure a bigger fraction of the qubits in the same basis.

Protocol QKD

1. (*Qubit distribution*) Alice prepares n EPR pairs of the form $(|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2}$, and sends one qubit of each pair to Bob, who confirms the receipt of the qubits. Then, Alice picks random $\theta \in \{0, 1\}^n$ and sends it to Bob, and Alice and Bob measure their respective qubits in basis θ to obtain \mathbf{x} on Alice’s side and \mathbf{y} on Bob’s side.
 2. (*Error estimation*) Alice chooses a random subset $s \subset [n]$ of size k and sends it to Bob. Then, Alice and Bob exchange \mathbf{x}_s and \mathbf{y}_s and compute $\beta := \omega(\mathbf{x}_s \oplus \mathbf{y}_s)$.
 3. (*Error correction*) Alice sends the syndrome syn of \mathbf{x}_s to Bob with respect to a suitable linear error correcting code (as described above). Bob uses syn to correct the errors in \mathbf{y}_s and obtains $\hat{\mathbf{x}}_s$. Let m be the bit-size of syn .
 4. (*Key distillation*) Alice chooses a random seed r for a universal hash function g with range $\{0, 1\}^\ell$, where ℓ satisfies $\ell < (1 - h(\beta))n - k - m$ (or $\ell = 0$ if the right-hand side is not positive), and sends it to Bob. Then, Alice and Bob compute $\mathbf{k} := g(r, \mathbf{x}_s)$ and $\hat{\mathbf{k}} := g(r, \hat{\mathbf{x}}_s)$, respectively.
-

an error probability $0 \leq \phi < \frac{1}{2}$. It follows that $\beta \approx \phi$, so that using an error correcting code that approaches the Shannon bound, Alice and Bob can extract close to $(1 - 2h(\phi))(n - k)$ bits of secret key, which is positive for ϕ smaller than approximately 11%. The difficult part is to prove security against an active adversary Eve. We first state the formal security claim.

Note that we cannot expect that Eve has (nearly) no information on \mathbf{K} , i.e. that $\Delta(\rho_{\mathbf{K}E}, \frac{1}{|\mathbf{K}|} \mathbb{I}_{\mathbf{K}} \otimes \rho_E)$ is small, since the bit-length ℓ of \mathbf{K} is not fixed but depends on the course of the protocol, and Eve can influence and thus obtain information on ℓ (and thus on \mathbf{K}). Theorem 4 though guarantees that the bit-length ℓ is the *only* information Eve learns on \mathbf{K} , in other words, \mathbf{K} is essentially random-and-independent of E when given ℓ .

Theorem 4 (Security of QKD). *Consider an execution of QKD in the presence of an adversary Eve. Let \mathbf{K} be the key obtained by Alice, and let E be Eve’s quantum system at the end of the protocol. Let $\tilde{\mathbf{K}}$ be chosen uniformly at random of the same bit-length as \mathbf{K} . Then, for any δ with $\beta + \delta \leq \frac{1}{2}$:*

$$\Delta(\rho_{\mathbf{K}E}, \rho_{\tilde{\mathbf{K}}E}) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}((1-h(\beta+\delta))n-k-m-\ell)} + 2 \exp(-\frac{1}{6}\delta^2k).$$

From an application point of view, the following question is of interest. Given the parameters n and k , and given a course of the protocol with observed error rate β and where an error-correcting code with syndrome length m was used, what is the maximal size ℓ of the extractable key \mathbf{K} if we want $\Delta(\rho_{\mathbf{K}E}, \rho_{\tilde{\mathbf{K}}E}) \leq \epsilon$ for a given ϵ ? From the bound in Theorem 4, it follows that for every choice of δ (with $\beta + \delta \leq \frac{1}{2}$), one can easily compute a possible value for ℓ simply by solving for ℓ . In order to compute the optimal value, one needs to maximize ℓ over the choice of δ .

The formal proof of Theorem 4 is given below. Informally, the argument goes as follows. The error estimation phase can be understood as applying a sampling strategy. From this, we can conclude that the state from which the raw key,

$\mathbf{x}_{\bar{s}}$, is obtained, is a superposition of states with bounded Hamming weight, so that Corollary 1 guarantees a certain amount of min-entropy within $\mathbf{x}_{\bar{s}}$. Privacy amplification then finishes the proof.

To indeed be able to model the error estimation procedure as a sampling strategy, we will need to consider a modified but *equivalent* way for Alice and Bob to jointly obtain \mathbf{x}_s and \mathbf{y}_s from the initial joint state, which will allow them to obtain the XOR-sum $\mathbf{x}_s \oplus \mathbf{y}_s$, and thus to compute β , *before* they measure the remaining part of the state, whose outcome then determines $\mathbf{x}_{\bar{s}}$. This modification is based on the so-called CNOT operation, U_{CNOT} , acting on $\mathbb{C}^2 \otimes \mathbb{C}^2$, and its properties that

$$U_{\text{CNOT}}(|b\rangle|c\rangle) = |b\rangle|b \oplus c\rangle \quad \text{and} \quad U_{\text{CNOT}}(H|b\rangle H|c\rangle) = H|b \oplus c\rangle H|c\rangle, \quad (2)$$

where the first holds by definition of U_{CNOT} , and the second is trivial to verify.

Proof. Throughout the proof, we use capital letters, Θ, \mathbf{X} etc. for the *random variables* representing the corresponding choices of θ, \mathbf{x} etc. in protocol QKD. Let the state, shared by Alice, Bob and Eve right after the quantum communication in the qubit distribution phase, be denoted by $|\psi_{ABE_o}\rangle$; ⁸ without loss of generality, we may indeed assume the shared state to be pure. For every $i \in [n]$, Alice and Bob then measure the respective qubits A_i and B_i from $|\psi_{ABE_o}\rangle$ in basis Θ_i , obtaining X_i and Y_i . This results in the hybrid state $\rho_{\Theta \mathbf{X} \mathbf{Y} E_o}$. For the proof, it will be convenient to introduce the additional random variables $\mathbf{W} = (W_1, \dots, W_n)$ and $\mathbf{Z} = (Z_1, \dots, Z_n)$, defined by

$$W_i := \begin{cases} X_i & \text{if } \Theta_i = 0 \\ Y_i & \text{if } \Theta_i = 1 \end{cases} \quad \text{and} \quad Z_i := X_i \oplus Y_i. \quad (3)$$

Note that, when given Θ , the random variables \mathbf{W} and \mathbf{Z} are uniquely determined by \mathbf{X} and \mathbf{Y} *and vice versa*, and thus we may equivalently analyze the hybrid state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$.

For the analysis, we will consider a slightly *different* experiment for Alice and Bob to obtain the very *same* state $\rho_{\Theta \mathbf{W} \mathbf{Z} E_o}$; the advantage of the modified experiment is that it can be understood as a sampling strategy. The modified experiment is as follows. First, the CNOT transformation is applied to every qubit pair $A_i B_i$ within $|\psi_{ABE_o}\rangle$ for $i \in [n]$, such that the state $|\varphi_{ABE_o}\rangle = (U_{\text{CNOT}}^{\otimes n} \otimes \mathbb{I}_{E_o})|\psi_{ABE_o}\rangle$ is obtained. Next, Θ is chosen at random as in the original scheme, and for every $i \in [n]$ the qubit pair $A_i B_i$ of the transformed state is measured as in the original scheme depending on Θ_i ; however, if $\Theta_i = 0$ then the resulting bits are denoted by W_i and Z_i , respectively, and if $\Theta_i = 1$ then they are denoted by Z_i and W_i , respectively, such that which bit is assigned to which variable depends on Θ_i . This is illustrated in Figure 1 (left and middle), where light and dark colored ovals represent measurements in the computational and Hadamard basis, respectively. It now follows immediately from the properties (2)

⁸ E_o represents Eve’s quantum state just after the quantum communication stage, whereas E represents Eve’s entire state at the end of the protocol (i.e., her quantum information and all classical information gathered during execution of QKD).




Fig. 1. Original and modified experiments for obtaining the same state $\rho_{\theta W Z E_0}$.

of the CNOT transformation and from the relation (3) between \mathbf{X}, \mathbf{Y} and \mathbf{W}, \mathbf{Z} that the state $\rho_{\theta W Z E_0}$ (or, equivalently, $\rho_{\theta X Y E_0}$) obtained in this modified experiment is exactly the same as in the original.

An additional modification we may do without influencing the final state is to *delay* some of the measurements: we assume that first the qubits are measured that lead to the Z_i 's, and only at some later point, namely after the *error estimation* phase, the qubits leading to the W_i 's are measured (as illustrated in Figure 1, right). This can be done since the relative Hamming weight of $X_S \oplus Y_S$ for a random subset $S \subset [n]$ (of size k) can be computed given \mathbf{Z} alone.

The crucial observation is now that this modified experiment can be viewed as a particular sampling strategy Ψ , as a matter of fact as the sampling strategy discussed in Example 3, being applied to systems A and B of the state $|\varphi_{ABE_0}\rangle$. Indeed: first, a subset of the $2n$ qubit positions is selected according to some probability distribution, namely of each pair $A_i B_i$ one qubit is selected at random (determined by θ_i). Then, the selected qubits are measured to obtain the bit string $\mathbf{Z} = (Z_1, \dots, Z_n)$. And, finally, a value β is computed as a (randomized) function of \mathbf{Z} : $\beta = \omega(\mathbf{Z}_S)$ for a random $S \subset [n]$ of size k . We point out that here the reference basis (as explained in Remark 2) is not the computational basis for all qubits, but is the Hadamard basis on the qubits in system A and the computational basis in system B ; however, as discussed in Remark 2, we may still apply the results from Section 4 (appropriately adapted).

It thus follows that for any fixed $\delta > 0$, the remaining state, from which \mathbf{W} is then obtained, is (on average over Θ and S) $\varepsilon_{\text{quant}}^\delta$ -close to a state which is (for any possible values for Θ, \mathbf{Z} and S) a superposition of states with relative Hamming weight in a δ -neighborhood of β . Note that the latter has to be understood with respect to the fixed reference basis (i.e., the Hadamard basis on A and the computational basis on B). In the following, we assume that the remaining state *equals* such a superposition, but we remember the error

$$\varepsilon_{\text{quant}}^\delta \leq \sqrt{\varepsilon_{\text{class}}^\delta} \leq 2 \exp\left(-\frac{1}{6}\delta^2 k\right).$$

where the bound on $\varepsilon_{\text{class}}^\delta$ was derived in Section 3.2.

Recall that \mathbf{W} is now obtained by measuring the remaining qubits; however, the basis used is opposite to the reference basis, namely the computational basis on the qubits A_i and the Hadamard basis on the qubits B_i . Hence, by Corollary 1 (and the subsequent discussion) we get a lower bound on the min-entropy of \mathbf{W} :

$$H_{\min}(\mathbf{W}|\boldsymbol{\Theta}\mathbf{Z}SE_o) \geq (1 - h(\beta + \delta))n.$$

Since \mathbf{W} is uniquely determined by \mathbf{X} (and vice versa) when given $\boldsymbol{\Theta}$ and \mathbf{Z} , the same lower bound also holds for $H_{\min}(\mathbf{X}|\boldsymbol{\Theta}\mathbf{Z}SE_o)$. Note that in QKD, the k qubit-pairs that are used for estimating β are not used anymore in the key distillation phase, so we are actually interested in the min-entropy of $\mathbf{X}_{\bar{S}}$. Additionally, we should take into account that Alice sends an m -bit syndrome SYN during the error correction phase. Hence, by using the chain rule, we obtain

$$H_{\min}(\mathbf{X}_{\bar{S}}|\boldsymbol{\Theta}\mathbf{Z}\mathbf{X}_{\bar{S}}SYNE_o) \geq (1 - h(\beta + \delta))n - k - m.^9$$

Finally, we apply privacy amplification (Theorem 1) to conclude the proof. \square

References

1. Bennett, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. In: Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing, pp. 175–179 (1984)
2. Bouman, N.J., Fehr, S.: Sampling in a quantum population, and applications (2009), <http://arxiv.org/abs/0907.4246>
3. Brass, D.: Optimal eavesdropping in quantum cryptography with six states. *Physical Review Letters* 81, 3018 (1998), <http://arxiv.org/abs/quant-ph/9805019>
4. Damgård, I., Fehr, S., Lunemann, C., Salvail, L., Schaffner, C.: Improving the security of quantum protocols via commit-and-open. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 408–427. Springer, Heidelberg (2009)
5. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301), 13–30 (1963)
6. Lo, H.K., Chau, H.F., Ardehali, M.: Efficient quantum key distribution scheme and a proof of its unconditional security. *J. Cryptol.* 18(2), 133–165 (2005)
7. Renner, R.: Security of Quantum Key Distribution. Ph.D. thesis, ETH Zürich (Switzerland) (September 2005), <http://arxiv.org/abs/quant-ph/0512258>
8. Renner, R.S., König, R.: Universally composable privacy amplification against quantum adversaries. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 407–425. Springer, Heidelberg (2005)
9. Shor, P.W., Preskill, J.: Simple Proof of Security of the BB84 Quantum Key Distribution Protocol. *Phys. Rev. Lett.* 85, 441–444 (2000)

⁹ Probably, it is possible to prove the lower bound: $(1 - h(\beta + \delta))(n - k) - m$ using a different sampling strategy. However, for that case the error probability of the related classical sampling strategy becomes harder to analyze. We have chosen for the current proof strategy and bound for the sake of simplicity.