Proposition 4.9. *Let D be a quantum system and $\{|x\rangle\}_{x\in\mathcal{X}}$ an orthonormal basis of $\mathcal{H}_{\mathsf{D}}$ where $\mathcal{X} = \{0,\dots,d-1\}$. Consider the unitary matrix $U_{X\mathsf{D}} \in \mathcal{U}(\mathcal{H}_X \otimes \mathcal{H}_{\mathsf{D}})$, where $\mathcal{H}_X = \mathcal{H}_{\mathsf{D}}$, such that $U_{X\mathsf{D}}|x,y\rangle = |x+y\rangle|y\rangle$ (where the addition is modulo d). Then, for any density matrix $\sigma_{\mathsf{DE}} \in \mathcal{D}(\mathcal{H}_{\mathsf{D}} \otimes \mathcal{H}_{\mathsf{E}})$, the state*

$$\rho_{X\mathsf{E}} = \mathrm{tr}_{\mathsf{D}}\big((U_{X\mathsf{D}} \otimes \mathbb{I}_{\mathsf{E}})(|0\rangle\langle 0| \otimes \sigma_{\mathsf{DE}})(U_{X\mathsf{D}} \otimes \mathbb{I}_{\mathsf{E}})^\dagger\big)$$

has classical X, and P_X equals the distribution of the outcome of measuring subsystem D of σ_{DE} in basis $\{|x\rangle\}_{x\in\mathcal{X}}$ and $\rho_{\mathsf{E}|X=x}$ equals the state to which E collapses when x is observed.

Proof. To simplify notation, we skip some of the indices. First, we consider the case of an "empty" E, i.e. $\sigma \in \mathcal{D}(\mathcal{H}_{\mathsf{D}})$. Due to linearity, it suffices to show the claim for a *pure* state $\sigma = |\varphi\rangle\langle\varphi|$, where we can write $|\varphi\rangle = \sum_x \alpha_x|x\rangle$. Then,

$$\begin{aligned}\mathrm{tr}_{\mathsf{D}}\big(U(|0\rangle\langle 0| \otimes \sigma)U^\dagger\big) &= \sum_{x,y} \alpha_x\bar{\alpha}_y \mathrm{tr}_{\mathsf{D}}\big(U(|0\rangle\langle 0| \otimes |x\rangle\langle y|)U^\dagger\big) = \sum_{x,y} \alpha_x\bar{\alpha}_y \mathrm{tr}_{\mathsf{D}}\big(U|0,x\rangle\langle 0,x|U^\dagger\big)\\ &= \sum_{x,y} \alpha_x\bar{\alpha}_y \mathrm{tr}_{\mathsf{D}}\big(|x\rangle|x\rangle\langle y|\langle y|\big) = \sum_{x,y} \alpha_x\bar{\alpha}_y \langle y|x\rangle|x\rangle\langle y| = \sum_x |\alpha_x|^2|x\rangle\langle x|\,.\end{aligned}$$

This is a classical state (with respect to $\{|x\rangle\}_x$), and $P_X(x) = |\alpha_x|^2$ is the probability of observing x when measuring σ in basis $\{|x\rangle\}_x$.

In case of a composite state $\sigma \in P(\mathcal{H}_{\mathsf{D}} \otimes \mathcal{H}_{\mathsf{E}})$, it still suffices to show the claim for a *pure* state $\sigma = |\varphi\rangle\langle\varphi|$, where now $|\varphi\rangle = \sum_x \alpha_x|x\rangle|\psi_x\rangle$. Similar to above, except that $|\psi_x\rangle\langle\psi_y|$ is carried along, it follows that

$$\rho_{X\mathsf{E}} = \sum_x |\alpha_x|^2|x\rangle\langle x| \otimes |\psi_x\rangle\langle\psi_x|$$

where $p_x = \sum_{i,j} \alpha_{x,i}\bar{\alpha}_{x,j} = \big|\sum_i \alpha_{x,i}\big|^2$ equals the probability to observe x when measuring the first part of σ, and $\rho_{\mathsf{E}|X=x} = |\psi_x\rangle\langle\psi_x|$ equals the state to which the second part of σ collapses when x is observed. □

5 Privacy Amplification

In this section we show how a *weak* (or *raw*) key X, about which an adversary has some but limited quantum information, can be transformed into a *strong* key K which looks essentially random to the adversary. The transformation is fully public, no secrecy is involved there. Such a process of "amplifying privacy" is called **privacy amplification**. The assumption that the adversary has only *limited* information on the weak key X will be expressed here by essentially requiring that the size of the adversary's quantum state, measured in qubits, is smaller than the a priory collision entropy $\mathrm{H}_2(X)$ or the min-entropy $\mathrm{H}_\infty(X)$ of X; we refer to Appendix B.5 for the definitions of $\mathrm{H}_2(X)$ and $\mathrm{H}_\infty(X)$ (as well as of the conditional versions).

Definition 5.1. *A function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{W}$ is called* **universal** *if*

$$P\big[f(S,x) = f(S,x')\big] \leq \frac{1}{|\mathcal{W}|}$$

for any $x \neq x' \in \mathcal{X}$ and for S uniformly distributed over $\mathcal{S}$.

Although a universal function does not have to be hashing, i.e. compressing, they are usually referred to as universal *hash* functions. Examples of universal (hash) functions are

$$f : \{0,1\}^{\ell\times n} \times \{0,1\}^n \to \{0,1\}^\ell,\ (A,x) \mapsto Ax$$

with $\ell \leq n$ and with all operations modulo 2, and

$$f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^\ell,\ (a,x) \mapsto [a \cdot x]_\ell\,,$$

with $\ell \leq n$ and where the multiplication is to be understood as multiplication in $\mathbb{F}_{2^\ell}$ and $[a \cdot x]_\ell$ stands for the first ℓ bits of $a \cdot x$ with respect to some fixed basis of the $\mathbb{F}_2$-vector-space $\mathbb{F}_{2^\ell}$.

Definition 5.2. *The* **max-entropy** *of a state $\rho \in \mathcal{D}(\mathcal{H})$ is defined as* $\mathrm{H}_0(\rho) = \log(\mathrm{rank}(\rho))$.

The max-entropy of a state ρ_{E} can be used to measure (or at least upper-bound) the entanglement between E and some, say classical, state X. In particular, if $\mathrm{H}_0(\rho_E) = 0$, and thus ρ_{E} is pure, $\rho_{X\mathsf{E}}$ must be of the form $\rho_X \otimes \rho_{\mathsf{E}}$ and hence E gives no information on X at all. Clearly, the max-entropy of an n-qubit state ρ is bounded by $\mathrm{H}_0(\rho) \leq n$.

Theorem 5.3 (Privacy Amplification). *Let X be a random variable over $\mathcal{X}$, and let E be a quantum system whose state may depend on X. Let $f : \mathcal{S} \times \mathcal{X} \to \{0,1\}^\ell$ be a universal hash function and let S be uniformly distributed over $\mathcal{S}$ (independent of X and E). Then*

$$\delta\big(\rho_{f(S,X)S\mathsf{E}}, \rho_U \otimes \rho_{S\mathsf{E}}\big) \leq \frac{1}{2}\, 2^{-\frac{1}{2}(\mathrm{H}_2(X) - \mathrm{H}_0(\rho_{\mathsf{E}}) - \ell)} \leq \frac{1}{2}\, 2^{-\frac{1}{2}(\mathrm{H}_\infty(X) - \mathrm{H}_0(\rho_{\mathsf{E}}) - \ell)}\,.$$

where U is uniformly distributed over $\{0,1\}^\ell$ (independent of X, S and E).

Informally, this means that if $\mathrm{H}_2(X)$ is significantly larger than $\mathrm{H}_0(\rho_{\mathsf{E}}) + \ell$, the max-entropy of the quantum state plus the size of the hash value, then $K = f(S,X)$ is essentially random when given S and the quantum state E; see Figure 9.

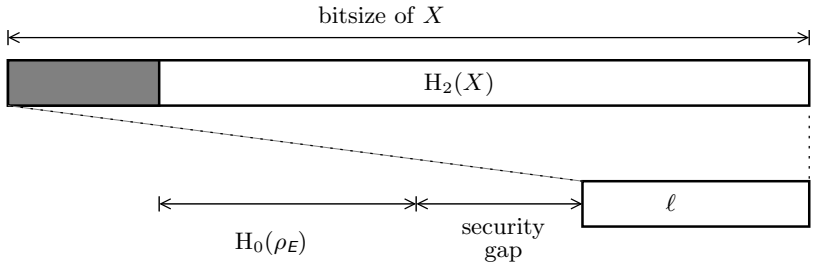

Fig. 9. Size of extractable key.

Remark 5.4. With an appropriate definition of the min-entropy $\mathrm{H}_\infty(X|\mathsf{E})$ of a classical X conditioned on a quantum state E, one can actually show that $\delta\big(\rho_{f(S,X)S\mathsf{E}}, \rho_U \otimes \rho_{S\mathsf{E}}\big)$ is bounded by $\frac{1}{2}\, 2^{-\frac{1}{2}(\mathrm{H}_\infty(\rho_{X\mathsf{E}}|\mathsf{E}) - \ell)}$. Theorem 5.3 is then a simple corollary, since $\mathrm{H}_\infty(X|\mathsf{E}) \geq \mathrm{H}_\infty(X) - \mathrm{H}_0(\rho_{\mathsf{E}})$. However, the proof is significantly more involved.

In the proof, we will actually consider the (square of the) **Hilbert-Schmidt distance**: for two density matrices ρ and σ in $\mathcal{D}(\mathcal{H})$ define $\Delta(\rho,\sigma) := \mathrm{tr}\big((\rho-\sigma)^2\big)$, i.e., the sum of the squares of the eigenvalues of $\rho - \sigma$. The following lemma shows that an upper bound on the (square of the) Hilbert-Schmidt distance translates to an upper bound on the trace norm distance.

Lemma 5.5. *For any density matrices $\rho, \sigma \in \mathcal{D}(\mathcal{H})$: $\delta(\rho,\sigma) \leq \frac{1}{2}\sqrt{\mathrm{rank}(\rho-\sigma)\Delta(\rho,\sigma)}$.*

Proof. Let $r = \operatorname{rank}(\rho - \sigma)$. We may assume that $\rho - \sigma$ is diagonal with the eigenvalues $\lambda_1, \dots, \lambda_r, 0, \dots, 0$ on its diagonal. It now follows that

$$2\delta(\rho,\sigma) = \sum_i |\lambda_i| = r \sum_i \frac{1}{r}\sqrt{\lambda_i^2} \leq r\sqrt{\sum_i \frac{1}{r}\lambda_i^2} = \sqrt{r\sum_i \lambda_i^2} = \sqrt{r\Delta(\rho,\sigma)}$$

where the inequality is Jensen's inequality (Proposition B.4). □

Proof (of Theorem 5.3). By Lemma 4.8 and using the fact that S, U and E are all independent,

$$\begin{aligned}&\delta\big(\rho_{f(S,X)S\mathsf{E}}, \rho_U \otimes \rho_{S\mathsf{E}}\big) = E_{s\leftarrow S}\big[\delta\big(\rho_{f(S,X)\mathsf{E}|S=s}, \rho_U \otimes \rho_{\mathsf{E}|S=s}\big)\big] = E_{s\leftarrow S}\big[\delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big)\big] \\ &\leq \frac{1}{2}2^{\frac{1}{2}(\ell+\mathrm{H}_0(\rho_{\mathsf{E}}))} E_{s\leftarrow S}\left[\sqrt{\Delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big)}\right] \leq \frac{1}{2}2^{\frac{1}{2}(\ell+\mathrm{H}_0(\rho_{\mathsf{E}}))}\sqrt{E_{s\leftarrow S}\big[\Delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big)\big]},\end{aligned}$$

where the inequalities follow from Lemma 5.5 and Jensen's inequality (Proposition B.4). It thus suffices to show that $E_{s\leftarrow S}\big[\Delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big)\big] \leq 2^{-\mathrm{H}_2(X)}$.

Let s be a fixed value in $\mathcal{S}$. Writing $W = f(s,X)$, we can re-write the square of the Hilbert-Schmidt distance as follows.

$$\begin{aligned}\Delta\big(\rho_{W\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big) &= \operatorname{tr}\Big(\big(\sum_w P_W(w)|w\rangle\langle w| \otimes \rho_{\mathsf{E}}^w - 2^{-\ell}\sum_w |w\rangle\langle w| \otimes \rho_{\mathsf{E}}\big)^2\Big) \\ &= \operatorname{tr}\Big(\big(\sum_w |w\rangle\langle w| \otimes (P_W(w)\rho_{\mathsf{E}}^w - 2^{-\ell}\rho_{\mathsf{E}})\big)^2\Big) \\ &= \operatorname{tr}\Big(\sum_w \big(P_W(w)\rho_{\mathsf{E}}^w - 2^{-\ell}\rho_{\mathsf{E}}\big)^2\Big) \\ &= \operatorname{tr}\Big(\sum_w \big(P_W(w)^2(\rho_{\mathsf{E}}^w)^2 - 2P_W(w)2^{-\ell}\rho_{\mathsf{E}}^w\rho_{\mathsf{E}} + 2^{-2\ell}\rho_{\mathsf{E}}^2\big)\Big) \\ &= \operatorname{tr}\Big(\sum_w P_W(w)^2(\rho_{\mathsf{E}}^w)^2 - 2^{-\ell}\rho_{\mathsf{E}}^2\Big)\end{aligned}$$

where the third equality follows from the fact that the matrix at hand is in block-diagonal form. Plugging in $f(s,X)$ for W and using linearity of the trace, we thus get

$$\Delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big) = \operatorname{tr}\Big(\sum_w P_{f(s,X)}(w)^2(\rho_{\mathsf{E}}^w)^2\Big) - \operatorname{tr}\Big(2^{-\ell}\rho_{\mathsf{E}}^2\Big).$$

Now note that

$$\begin{aligned}\operatorname{tr}\Big(\sum_w P_{f(s,X)}(w)^2(\rho_{\mathsf{E}}^w)^2\Big) &= \operatorname{tr}\Big(\sum_w \big(\sum_{x:f(s,x)=w} P_X(x)\rho_{\mathsf{E}}^x\big)\big(\sum_{x':f(s,x')=w} P_X(x')\rho_{\mathsf{E}}^{x'}\big)\Big) \\ &= \operatorname{tr}\Big(\sum_{x,x'} P_X(x)P_X(x')\delta_{f(s,x),f(s,x')}\rho_x\rho_{x'}\Big) = \sum_{x,x'} P_X(x)P_X(x')\delta_{f(s,x),f(s,x')}\operatorname{tr}\big(\rho_{\mathsf{E}}^x\rho_{\mathsf{E}}^{x'}\big)\end{aligned}$$

where $\delta_{w,w'}$ is the Kronecker Delta with $\delta_{w,w'} = 1$ if $w = w'$ and 0 otherwise, and

$$\operatorname{tr}\big(2^{-\ell}\rho_{\mathsf{E}}^2\big) = \operatorname{tr}\Big(2^{-\ell}\big(\sum_x P_X(x)\rho_{\mathsf{E}}^x\big)^2\Big) = \sum_{x,x'} P_X(x)P_X(x')2^{-\ell}\operatorname{tr}\big(\rho_{\mathsf{E}}^x\rho_{\mathsf{E}}^{x'}\big)$$

Using that due to the universal property $E_{s\leftarrow S}[\delta_{f(s,x),f(s,x')}] = P[f(S,x) = f(S,x')] \leq 2^{-\ell}$ for all $x \neq x'$, and using that $0 \leq \operatorname{tr}\big((\rho_{\mathsf{E}}^x)^2\big) \leq 1$ and $0 \leq \operatorname{tr}\big(\rho_{\mathsf{E}}^x\rho_{\mathsf{E}}^{x'}\big)$, it follows that

$$E_{s\leftarrow S}\big[\Delta\big(\rho_{f(s,X)\mathsf{E}}, \rho_U \otimes \rho_{\mathsf{E}}\big)\big] = \sum_{x,x'} P_X(x)P_X(x')\big(E_{s\leftarrow S}[\delta_{f(s,x),f(s,x')}] - 2^{-\ell}\big)\operatorname{tr}\big(\rho_{\mathsf{E}}^x\rho_{\mathsf{E}}^{x'}\big)$$

$$\leq \sum_x P_X(x)^2 \mathrm{tr}\big(\rho_x^2\big) \leq \sum_x P_X(x)^2 = 2^{-\mathrm{H}_2(X)}\,,$$

which was to be shown. □

The proof of the following is given as an exercise.

Corollary 5.6. *Let X be a random variable over $\mathcal{X}$, let Y be a random variable over $\mathcal{Y}$, and let E be a quantum system whose state may depend on X and Y. Let $f : \mathcal{S} \times \mathcal{X} \to \{0,1\}^\ell$ be a universal hash function and let S be uniformly distributed over $\mathcal{S}$ (independent of X, Y and E). Then*

$$\delta\big(\rho_{f(S,X)S\mathsf{E}}, \rho_U \otimes \rho_{S\mathsf{E}}\big) \leq \frac{1}{2}\, 2^{-\frac{1}{2}(\mathrm{H}_2(X|Y)-\mathrm{H}_0(\rho_\mathsf{E})-\ell)} \leq \frac{1}{2}\, 2^{-\frac{1}{2}(\mathrm{H}_\infty(X|Y)-\mathrm{H}_0(\rho_\mathsf{E})-\ell)}\,,$$

where U is uniformly distributed over $\{0,1\}^\ell$ (independent of X, S and E).

6 Subset Sampling – Classical and Quantum

The **Hamming weight** $W(x)$ of a bit-string $x = (x_1, \ldots, x_m) \in \{0,1\}^m$ is defined to be the number of 1's occurring within x. Similarly, the **relative** Hamming weight $\omega(x)$ of x is given by its Hamming weight divided by its bit-length m: $\omega(x) = W(x)/m$. We say that the relative Hamming weight of x is *ε-close* to $\beta \in \mathbb{R}$, denoted as $\omega(x) \approx_\varepsilon \beta$, if $|\omega(x) - \beta| \leq \varepsilon$. For any subset $T \subseteq \{1, \ldots, m\}$ of size k, we write X_T for the restriction of x to the positions in T: $x_T = (x_i)_{i \in T} \in \{0,1\}^k$.

Consider the following problem: we want to *estimate* the (relative) Hamming weight of an unknown but fixed string $x \in \{0,1\}^m$ (of known bit-length m) by only looking at a small number of positions in x. A canonical way to do so is as follows: choose at random a sample subset $T \subset \{1, \ldots, m\}$ of linear size (i.e. size αm for some constant $0 < \alpha < 1$), and take $\omega(x_T)$ as estimate for $\omega(x)$. Very generally, we allow the following kinds of sampling strategies: choose a sample subset $T \subset \{1, \ldots, m\}$ according to *some* fixed probability distribution P_T, and compute the estimate for $\omega(x)$ as *some* (possibly randomized) function of x_T.

Definition 6.1. *A* **sampling strategy** *consists of a triple $\Sigma = (P_T, P_S, estim)$, where P_T is a distribution over the subsets of $\{1, \ldots, m\}$, P_S is a (independent) distribution over a finite set $\mathcal{S}$, and estim is a function $estim : \mathcal{S} \times \{(t, v) : t \subseteq [n], v \in \{0,1\}^{|t|}\} \to \mathbb{R}$.*

We want to measure the reliability of such a general estimation strategy, i.e., how well it predicts the (relative) Hamming weight of the string x. Actually, for technical reasons (and because the positions within the sample subset T are anyway revealed), we want to measure how well such a general strategy predicts the (relative) Hamming weight of $x_{\bar{T}} \in \{0,1\}^n$ (where $n = m - |T|$), i.e., of x restricted to the positions $\bar{T} = \{1, \ldots, m\} \setminus T$ outside of the sample T (see Figure 10, top).

Definition 6.2. *For any $\varepsilon > 0$, we define the* **error probability**

$$err_\varepsilon(\Sigma) := \max_{x \in \{0,1\}^m} P\big[\omega(x_{\bar{T}}) \not\approx_\varepsilon estim(S, T, x_T)\big]$$

where the probability is over the choices of T and S according to $P_{TS} = P_T P_S$.

By definition, for any choice of $x \in \{0,1\}^m$: $\omega(x_{\bar{T}}) \approx_\varepsilon estim(S, T, x_T)$ except with probability at most $err_\varepsilon(\Sigma)$. Using classical sampling theory (see e.g. [**?**]), one can e.g. show that for