Private Quantum Channels and the Cost of Randomizing Quantum Information

Private Quantum Channels and the Cost of Randomizing Quantum Information

1 Introduction

Secure transmission of classical information is a well studied topic. Suppose Alice wants to send an n-bit message M to Bob over an insecure (i.e. spied-on) channel, in such a way that the eavesdropper Eve cannot obtain any information about M from tapping the channel. If Alice and Bob share some secret n-bit key K, then here is a simple way for them to achieve their goal: Alice exclusive-ors M with K and sends the result $M' = M \oplus K$ over the channel, Bob then xors M' again with K and obtains the original message $M' \oplus K = M$. Eve may see the encoded message M', but if she does not know K then this will give her no information about the real message M, since for any M there is a key K' giving rise to the same encoding M'. This scheme is known as the *Vernam cipher* or *one-time pad* ("one-time" because K can be used only once if we want information-theoretic security). It shows that n bits of shared secret key are sufficient to securely transmit n bits of information. Shannon [Sha48, Sha49] has shown that this scheme is optimal: n bits of shared key are also *necessary* in order to transmit an n-bit message in an information-theoretically secure way.

Now let us consider the analogous situation in the quantum world. Alice and Bob are connected by a one-way quantum channel, to which an eavesdropper Eve has complete access. Alice wants to transmit to Bob some n-qubit state ρ taken from some set $\mathcal{S}$, without allowing Eve to obtain any information about ρ. Alice and Bob could easily achieve such security if they share n EPR-pairs (or if they were able to establish EPR-pairs over a secure quantum channel), for then they can apply teleportation [BBC$^+$93] and transmit every qubit via 2 random classical bits, which will give Eve no information whatsoever. But now suppose Alice and Bob do not share EPR-pairs, but instead they only have the resource of shared randomness, which is weaker but easier to maintain.

A first question is: is it at all possible to send quantum information fully securely using only a finite amount of randomness? At first sight this may seem hard: Alice and Bob have to "hide" the amplitudes of a quantum state, which are infinitely precise complex numbers. Nevertheless, the question has a positive answer. More precisely, to privately send n qubits, a $2n$-bit classical key is sufficient. The encryption technique is fairly natural. Alice applies to the state ρ she wants to transmit a reversible quantum operation specified by the shared key K (basically, she applies a random Pauli matrix to each qubit), and she sends the result ρ' to Bob. In the most general setting this reversible operation can be represented as doing a unitary operation on the state ρ augmented with a known fixed ancilla state ρ_a. Knowing the key K that Alice used, Bob knows which operation Alice applied and he can reverse this, remove the ancilla, and retrieve ρ. In order for this scheme to be information-theoretically secure against the eavesdropper, we have to require that Eve always "sees" the same density matrix ρ_0 on the channel, no matter what ρ was. Because Eve does not know K, this condition can indeed be satisfied. Accordingly, an insecure quantum channel can be made secure (private) by means of shared classical randomness.

A second question is, then, *how much* key Alice and Bob need to share in order to be able to privately transmit any n-qubit state. A good way to measure key size is by the amount of entropy required to create it. As one might imagine, showing that $2n$ bits of key are also necessary is the most challenging part of the article. We prove this in Section 5. Accordingly, in analogy with the classical one-time pad, we have an optimal quantum one-time pad which uses $2n$ classical bits to completely "hide" n qubits from Eve. In particular, hiding a qubit is only twice as hard as hiding a classical bit, despite the fact that in the qubit we now have to hide amplitudes coming from a continuous set.

Now imagine an alternative scenario. Alice has a state ρ from some specific set and she wants to randomize it completely. How much entropy does she need for this? That is, what is the thermodynamical cost of forgetting quantum information? A natural and general way to do that is for Alice to perform a unitary transformation to ρ augmented with an ancilla and then to forget which one. The thermodynamical price of this operation is now the entropy of the probability distribution over the set of unitary transformations. The parallel between these two scenarios should be clear. If one has a private quantum channel, one automatically has a related randomization procedure. Consequently, we obtain the result that $2n$ bits are necessary and sufficient to randomize an n-qubit quantum register. For the case $n = 1$, this result has also been obtained by Braunstein, Lo, and Spiller [BLS99, Lo99].

2 Preliminaries

2.1 States and operators

We use $||v||$ for the Euclidean norm of vector v. If A is a matrix, then we use $A^\dagger$ for its conjugate transpose and $\mathrm{Tr}(A)$ for its trace (the sum of its diagonal entries). A square matrix A is *Hermitian* if $A = A^\dagger$, and *unitary* if $A^{-1} = A^\dagger$. Important examples of unitary transformations are the 4 *Pauli matrices*:

$$\sigma_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},\ \sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},\ \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix},\ \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Let $|0\rangle, \ldots, |M-1\rangle$ denote the basis states of some M-dimensional Hilbert space $\mathcal{H}_M$. We use $\mathcal{H}_{2^n}$ for the Hilbert space whose basis states are the 2^n classical n-bit strings. A *pure quantum state* $|\phi\rangle$ is a norm-1 vector in $\mathcal{H}_M$. We treat $|\phi\rangle$ as an M-dimensional column vector and use $\langle\phi|$ for the row vector that is its conjugate transpose. The *inner product* between pure states $|\phi\rangle$ and $|\psi\rangle$ is $\langle\phi|\psi\rangle$. A *mixed quantum state* or *density matrix* ρ is a non-negative Hermitian matrix that has trace $\mathrm{Tr}(\rho) = 1$. The density matrix corresponding to a pure state $|\phi\rangle$ is $|\phi\rangle\langle\phi|$. Because a density matrix ρ is Hermitian, it has a diagonalization $\rho = \sum_{i=1}^{N} p_i |\phi_i\rangle\langle\phi_i|$, where the p_i are its eigenvalues, $p_i \geq 0$, $\sum_i p_i = 1$, and the $|\phi_i\rangle$ form an orthonormal set. Thus ρ can be viewed as describing a probability distribution over pure states. We use $\tilde{I}_M = \frac{1}{M} I_M = \frac{1}{M} \sum_{i=1}^{M} |i\rangle\langle i|$ to denote the totally mixed state, which represents the uniform distribution on all basis states. If two systems are in pure states $|\phi\rangle$ and $|\psi\rangle$, respectively, then their joint state is the tensor product pure state $|\phi\rangle \otimes |\psi\rangle = |\phi\rangle|\psi\rangle$. If two systems are in mixed states ρ_1 and ρ_2, respectively, then their joint state is the tensor product $\rho_1 \otimes \rho_2$. Note that $(|\phi\rangle \otimes |\psi\rangle)(\langle\phi| \otimes \langle\psi|)$ is the same as $|\phi\rangle\langle\phi| \otimes |\psi\rangle\langle\psi|$.

Applying a unitary transformation U to a pure state $|\phi\rangle$ gives pure state $U|\phi\rangle$, applying U to a mixed state ρ gives mixed state $U\rho U^\dagger$. We will use $\mathcal{E} = \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}$ to denote the *superoperator* which applies U_i with probability p_i to its argument (we assume $\sum_i p_i = 1$). Thus $\mathcal{E}(\rho) = \sum_i p_i U_i \rho U_i^\dagger$. Quantum mechanics allows for more general superoperators, but this type suffices for our purposes. If two superoperators $\mathcal{E} = \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}$ and $\mathcal{E}' = \{\sqrt{p'_i}U'_i \mid 1 \leq i \leq N'\}$ are identical ($\mathcal{E}(\rho) = \mathcal{E}'(\rho)$ for all ρ), then they are unitarily related in the following way [Nie98, Section 3.2] (where we assume $N \geq N'$ and if $N > N'$ we pad $\mathcal{E}'$ with zero operators to make $\mathcal{E}$ and $\mathcal{E}'$ of equal size): there exists a unitary $N \times N$ matrix A such that for all i

$$\sqrt{p_i}U_i = \sum_{j=1}^{N} A_{ij}\sqrt{p'_j}U'_j.$$

2.2 Von Neumann entropy

Let density matrix ρ have the diagonalization $\sum_{i=1}^{N} p_i |\phi_i\rangle \langle\phi_i|$. The *Von Neumann entropy* of ρ is $S(\rho) = H(p_1, \ldots, p_N) = -\sum_{i=1}^{N} p_i \log p_i$, where H is the classical entropy function. This $S(\rho)$ can be interpreted as the minimal Shannon entropy of the measurement outcome, minimized over all possible complete measurements. Note that $S(\rho)$ only depends on the eigenvalues of ρ. The following properties of Von Neumann entropy will be useful later (for proofs see for instance [Weh78]).

1. $S(|\phi\rangle \langle\phi|) = 0$, for every pure state $|\phi\rangle$.
2. $S(\rho_1 \otimes \rho_2) = S(\rho_1) + S(\rho_2)$.
3. $S(U\rho U^\dagger) = S(\rho)$.
4. $S(\lambda_1\rho_1 + \lambda_2\rho_2 + \cdots + \lambda_n\rho_n) \geq \lambda_1 S(\rho_1) + \lambda_2 S(\rho_2) + \cdots + \lambda_n S(\rho_n)$ if $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$.
5. If $\rho = \sum_{i=1}^{N} p_i |\phi_i\rangle \langle\phi_i|$ with the $|\phi_i\rangle$ not necessarily orthogonal, then $S(\rho) \leq H(p_1, \ldots, p_N)$.

3 Private Quantum Channel

Let us sketch the scenario for a private quantum channel. There are N possible keys, which we identify for convenience with the numbers $1, \ldots, N$. The ith key has probability p_i, so the key has entropy $H(p_1, \ldots, p_N)$ when viewed as a random variable. Each key i corresponds to a unitary transformation U_i. Suppose Alice wants to send a pure state $|\phi\rangle$ from some set $\mathcal{S}$ to Bob. She appends some fixed ancilla qubits in state ρ_a to $|\phi\rangle\langle\phi|$ and then applies U_i to $|\phi\rangle\langle\phi| \otimes \rho_a$, where i is her key. She sends the resulting state to Bob. Bob, who shares the key i with Alice, applies U_i^{-1} to obtain $|\phi\rangle\langle\phi| \otimes \rho_a$, removes the ancilla ρ_a, and is left with Alice's message $|\phi\rangle\langle\phi|$. Now in order for this to be secure against an eavesdropper Eve, we have to require that if Eve does not know i, then the density matrix ρ_0 that she gets from monitoring the channel is independent of $|\phi\rangle$. This implies that she gets no information at all about $|\phi\rangle$. Of course, Eve's measuring the channel might destroy the encoded message, but this is like classically jamming the channel and cannot be avoided. The point is that *if* Eve measures, then she receives no information about $|\phi\rangle$. It is not hard to see that this is the most general quantum mechanical scenario which allows Bob to recover the message perfectly and at the same time gives Eve zero information.

We formalize this scenario as follows.

Definition 3.1 *Let $\mathcal{S} \subseteq \mathcal{H}_{2^n}$ be a set of pure n-qubit states, $\mathcal{E} = \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}$ be a superoperator where each U_i is a unitary mapping on $\mathcal{H}_{2^m}$, $\sum_{i=1}^{N} p_i = 1$, ρ_a be an $(m-n)$-qubit density matrix, and ρ_0 be an m-qubit density matrix. Then $[\mathcal{S}, \mathcal{E}, \rho_a, \rho_0]$ is called a* Private Quantum Channel (**PQC**) *if and only if for all $|\phi\rangle \in \mathcal{S}$ we have*

$$\mathcal{E}(|\phi\rangle\langle\phi| \otimes \rho_a) = \sum_{i=1}^{N} p_i U_i \left(|\phi\rangle\langle\phi| \otimes \rho_a\right) U_i^{\dagger} = \rho_0.$$

If $n = m$ (i.e. no ancilla), then we omit ρ_a.

Note that by linearity, if the **PQC** works for all pure states in $\mathcal{S}$, then it also works for density matrices over $\mathcal{S}$: applying the **PQC** to a mixture of states from $\mathcal{S}$ gives the same ρ_0 as when we apply it to a pure state. Accordingly, if $[\mathcal{S}, \{\sqrt{p_i}\,U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$ is a **PQC**, then $H(p_1, \ldots, p_N)$ bits of shared randomness are sufficient for Alice to send any mixture ρ of $\mathcal{S}$-states to Bob in a secure way. Alice encodes ρ in a reversible way depending on her key i and Bob can decode because he knows the same i and hence can reverse Alice's operation U_i. On the other hand, Eve has no information about the key i apart from the distribution p_i, so from her point of view the channel is in state $\rho_{Eve} = \rho_0$. This is independent of the ρ that Alice wants to send, and hence gives Eve no information about ρ.

4 Examples of Private Quantum Channels

In this section we exhibit some private quantum channels. The first uses $2n$ bits key to send privately any n-qubit state. The idea is simply to apply a random Pauli matrix to each bit individually. This takes 2 random bits per qubit and it is well known that the resulting qubit is in the completely mixed state. For notational convenience we identity the numbers $\{0, \ldots, 2^{2n}-1\}$ with the set $\{0,1,2,3\}^n$. For $x \in \{0,1,2,3\}^n$ we use $x_i \in \{0,1,2,3\}$ for its ith entry, and we use $\overline{\sigma_x}$ to denote the n-qubit unitary transformation $\sigma_{x_1} \otimes \cdots \otimes \sigma_{x_n}$.

Theorem 4.1 *If* $\mathcal{E} = \{\frac{1}{\sqrt{2^{2n}}}\overline{\sigma_x} \mid x \in \{0,1,2,3\}^n\}$, *then* $[\mathcal{H}_{2^n}, \mathcal{E}, \tilde{I}_{2^n}]$ *is a* **PQC**.

Proof It is easily verified that applying each σ_i with probability $1/4$ to a qubit puts that qubit in the totally mixed state $\tilde{I}_2$ (no matter if it is entangled with other qubits). Operator $\mathcal{E}$ just applies this treatment to each of the n qubits, hence $\mathcal{E}(|\phi\rangle\langle\phi|) = \tilde{I}_{2^n}$ for every $|\phi\rangle \in \mathcal{H}_{2^n}$.

Since the above $\mathcal{E}$ contains 2^{2n} operations and they have uniform probability, it follows that $2n$ bits of private key suffice to privately send any state from $\mathcal{H}_{2^n}$.

The next theorem shows that there is some nontrivial subspace of $\mathcal{H}_{2^n}$ where n bits of private key suffice, namely the set of all tensor products of real-amplitude qubits.

Theorem 4.2 *If* $B = \{\cos(\theta)|0\rangle + \sin(\theta)|1\rangle \mid 0 \leq \theta < 2\pi\}$, $\mathcal{S} = B^{\otimes n}$, *and* $\mathcal{E} = \{\frac{1}{\sqrt{2^n}}\overline{\sigma_x} \mid x \in \{0,2\}^n\}$, *then* $[\mathcal{S}, \mathcal{E}, \tilde{I}_{2^n}]$ *is a* **PQC**.

Proof This is easily verified: applying σ_0 and σ_2, each with probability $1/2$, puts any qubit from B in the totally mixed state. Operator $\mathcal{E}$ does this to each of the n qubits individually.

Note that if we restrict B to classical bits (i.e. $\theta \in \{0, \pi/2\}$) then the above **PQC** reduces to the classical one-time pad: flipping each bit with probability $1/2$ gives information-theoretical security.

In the previous **PQC**s, ρ_0 was the completely mixed state $\tilde{I}_{2^n}$. This is no accident, and holds whenever $n = m$ and $\tilde{I}_{2^n}$ is one of the states that the **PQC** can send:

Theorem 4.3 *If $[\mathcal{S}, \mathcal{E}, \rho_0]$ is a* **PQC** *without ancilla and $\tilde{I}_{2^n}$ can be written as a mixture of S-states, then $\rho_0 = \tilde{I}_{2^n}$.*

Proof If $\tilde{I}_{2^n}$ can be written as a mixture of S-states, then

$$\rho_0 = \mathcal{E}(\tilde{I}_{2^n}) = \sum_{i=1}^{N} p_i U_i \tilde{I}_{2^n} U_i^\dagger = \sum_{i=1}^{N} \frac{p_i}{2^n} U_i U_i^\dagger = \sum_{i=1}^{N} \frac{p_i}{2^n} I_{2^n} = \tilde{I}_{2^n}.$$

In general ρ_0 need not be $\tilde{I}_{2^n}$. For instance, let $\mathcal{S} = \{|0\rangle, \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\}$, $\mathcal{E} = \{\sqrt{p_1} I_2, \frac{\sqrt{p_2}}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}\}$ with $p_1 = p_2 = 1/2$, and $\rho_0 = \begin{pmatrix} \frac{3}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{pmatrix}$. Then it is easily verified that $[\mathcal{S}, \mathcal{E}, \rho_0]$ is a **PQC**.

5 Lower Bound on the Entropy of PQCs

In the previous section we showed that $2n$ bits of entropy suffice for a **PQC** that can send arbitrary n-qubit states. In this section we will show that $2n$ bits are also *necessary* for this. Very recently and independently of our work, this $2n$-bit lower bound was also proven by Boykin and Roychowdhury [BR00] for the special case where the **PQC** is not allowed to use any ancilla qubits. We will first give a shorter version of their proof, basically by observing that a large part of it can be replaced by a reference to the unitary equivalence of identical superoperators stated at the end of Section 2.1.

Theorem 5.1 *If* $[\mathcal{H}_{2^n}, \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \tilde{I}_{2^n}]$ *is a* **PQC**, *then* $H(p_1, \ldots, p_N) \geq 2n$.

Proof Let $\mathcal{E} = \{\sqrt{p_i}U_i\}$, $\mathcal{E}' = \{\frac{1}{\sqrt{2^{2n}}}\overline{\sigma_x} \mid x \in \{0,1,2,3\}^n\}$ be the superoperator of Theorem 4.1, and let $K = \max(2^{2n}, N)$. Since $\mathcal{E}(\rho) = \mathcal{E}'(\rho) = \tilde{I}_{2^n}$ for all n-qubit states ρ, we have that $\mathcal{E}$ and $\mathcal{E}'$ are unitarily related in the way mentioned in Section 2.1: there exists a unitary $K \times K$ matrix A such that for all $1 \leq i \leq N$ we have

$$\sqrt{p_i}U_i = \sum_{x \in \{0,1,2,3\}^n} A_{ix} \frac{1}{\sqrt{2^{2n}}} \overline{\sigma_x}.$$

We can view the set of all $2^n \times 2^n$ matrices as a 2^{2n}-dimensional vector space, with inner product $\langle M, M' \rangle = \mathrm{Tr}(M^\dagger M')/2^n$ and induced matrix norm $||M|| = \sqrt{\langle M, M \rangle}$ (as done in [BR00]). Note that $||M|| = 1$ if M is unitary. The set of all $\overline{\sigma_x}$ forms an orthonormal basis for this vector space, so we get:

$$p_i = ||\sqrt{p_i}U_i||^2 = ||\sum_x A_{ix} \frac{1}{\sqrt{2^{2n}}} \overline{\sigma_x}||^2 = \frac{1}{2^{2n}} \sum_x |A_{ix}|^2 \leq \frac{1}{2^{2n}}.$$

Hence $N \geq 2^{2n}$ and $H(p_1, \ldots, p_N) \geq 2n$.

However, even granted this result it is still conceivable that a **PQC** might require less randomness if it can "spread out" its encoding over many ancilla qubits — it is even conceivable that those ancilla qubits can be used to *establish* privately shared randomness using some variant of quantum key distribution. The general case with ancilla is not addressed in [BR00], and proving that the $2n$-bit lower bound extends to this case requires more work. The next few theorems will do this. These show that a **PQC** that can transmit any n-qubit state requires $2n$ bits of randomness, no matter how many ancilla qubits it uses. Thus Theorem 4.1 exhibits an optimal quantum one-time pad, analogous to the optimal classical one-time pad mentioned in the introduction.

We will use the notation $\mathcal{C}_k = \{|i\rangle \mid 0 \leq i \leq k-1\}$ for the set of the first k classical states. The next theorem states that a **PQC** that privately conveys n qubits using m bits of key, can be transformed into a **PQC** that privately conveys any state from $\mathcal{C}_{2^{2n}}$, still using only m bits of key.

Theorem 5.2 *If there exists a* **PQC** $[\mathcal{H}_{2^n}, \mathcal{E} = \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$, *then there exists a* **PQC** $[\mathcal{C}_{2^{2n}}, \mathcal{E}' = \{\sqrt{p_i}U_i' \mid 1 \leq i \leq N\}, \rho_a, \tilde{I}_{2^n} \otimes \rho_0]$.

Proof For ease of notation we assume without loss of generality that $\mathcal{E}$ uses no ancilla, so we assume ρ_0 is an n-qubit state and omit ρ_a (this does not affect the proof in any way). We first show that $\mathcal{E}(|x\rangle\langle y|) = 0$ whenever $x, y \in \mathcal{C}_{2^{2n}}$ and $x \neq y$ ($\mathcal{E}(|x\rangle\langle y|)$ is well-defined but somewhat of an abuse of notation, since the matrix $|x\rangle\langle y|$ is not a density matrix). This is implied by the following 3 equalities:

$$\rho_0 = \mathcal{E}\left(\frac{1}{2}(|x\rangle\langle x| + |y\rangle\langle y|)\right) = \frac{1}{2}\left(\mathcal{E}(|x\rangle\langle x|) + \mathcal{E}(|y\rangle\langle y|)\right).$$

$$\rho_0 = \mathcal{E}\left((\frac{1}{\sqrt{2}}(|x\rangle + |y\rangle))(\frac{1}{\sqrt{2}}(\langle x| + \langle y|))\right) = \frac{1}{2}\left(\mathcal{E}(|x\rangle\langle x|) + \mathcal{E}(|y\rangle\langle y|) + \mathcal{E}(|x\rangle\langle y|) + \mathcal{E}(|y\rangle\langle x|)\right).$$

$$\rho_0 = \mathcal{E}\left((\frac{1}{\sqrt{2}}(|x\rangle + i\,|y\rangle))(\frac{1}{\sqrt{2}}(\langle x| - i\,\langle y|))\right) = \frac{1}{2}\left(\mathcal{E}(|x\rangle\langle x|) + \mathcal{E}(|y\rangle\langle y|) - i\mathcal{E}(|x\rangle\langle y|) + i\mathcal{E}(|y\rangle\langle x|)\right).$$

The first and second equality imply $\mathcal{E}(|x\rangle\langle y|) + \mathcal{E}(|y\rangle\langle x|) = 0$, the first and third equality imply $\mathcal{E}(|x\rangle\langle y|) - \mathcal{E}(|y\rangle\langle x|) = 0$. Hence $\mathcal{E}(|x\rangle\langle y|) = \mathcal{E}(|y\rangle\langle x|) = 0$.

We now define $\mathcal{E}'$ and show that it is a **PQC**. Intuitively, $\mathcal{E}'$ will map every state from $\mathcal{C}_{2^{2n}}$ to a tensor product of n Bell states by mapping pairs of bits to one of the four Bell states.[2] The second bits of the pairs are then moved to the second half of the state and randomized by applying $\mathcal{E}$ to them. Because of the entanglement between the two halves of each Bell state, the resulting $2n$-qubit density matrix will be $\tilde{I}_{2^n} \otimes \rho_0$. More specifically, define

$$U|x\rangle = (\overline{\sigma_x} \otimes I_{2^n}) \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle |i\rangle,$$

with $\overline{\sigma_x} = \sigma_{x_1} \otimes \cdots \otimes \sigma_{x_n}$ as in Theorem 4.1. Also define $U_i' = (I_{2^n} \otimes U_i)U$. It remains to show that $\mathcal{E}'(|x\rangle\langle x|) = \tilde{I}_{2^n} \otimes \rho_0$ for all $|x\rangle \in \mathcal{C}_{2^{2n}}$:

$$
\begin{aligned}
& \mathcal{E}'(|x\rangle\langle x|) \\
= & \sum_{i=1}^{N} p_i (I_{2^n}\otimes U_i)\left[(\overline{\sigma_x}\otimes I_{2^n})\left(\frac{1}{\sqrt{2^n}}\sum_{y=0}^{2^n-1}|y\rangle|y\rangle\right)\left(\frac{1}{\sqrt{2^n}}\sum_{z=0}^{2^n-1}\langle z|\langle z|\right)(\overline{\sigma_x}\otimes I_{2^n})^\dagger\right](I_{2^n}\otimes U_i)^\dagger \\
= & (\overline{\sigma_x}\otimes I_{2^n})\left[\frac{1}{2^n}\sum_{i=1}^{N} p_i (I_{2^n}\otimes U_i)\left(\sum_{y,z\in\{0,2^n-1\}}|y\rangle\langle z|\otimes|y\rangle\langle z|\right)(I_{2^n}\otimes U_i)^\dagger\right](\overline{\sigma_x}\otimes I_{2^n})^\dagger \\
= & (\overline{\sigma_x}\otimes I_{2^n})\left[\frac{1}{2^n}\sum_{y,z\in\{0,2^n-1\}}|y\rangle\langle z|\otimes\left(\sum_{i=1}^{N} p_i U_i|y\rangle\langle z|U_i^\dagger\right)\right](\overline{\sigma_x}\otimes I_{2^n})^\dagger \\
= & (\overline{\sigma_x}\otimes I_{2^n})\left[\frac{1}{2^n}\sum_{y,z\in\{0,2^n-1\}}|y\rangle\langle z|\otimes\mathcal{E}(|y\rangle\langle z|)\right](\overline{\sigma_x}\otimes I_{2^n})^\dagger \\
\overset{(*)}{=} & (\overline{\sigma_x}\otimes I_{2^n})\left[\frac{1}{2^n}\sum_{y=0}^{2^n-1}|y\rangle\langle y|\otimes\mathcal{E}(|y\rangle\langle y|)\right](\overline{\sigma_x}\otimes I_{2^n})^\dagger \\
= & (\overline{\sigma_x}\otimes I_{2^n})\left[\tilde{I}_{2^n}\otimes\rho_0\right](\overline{\sigma_x}\otimes I_{2^n})^\dagger \\
= & \tilde{I}_{2^n}\otimes\rho_0.
\end{aligned}
$$

In the step marked by $(*)$ we used that $\mathcal{E}(|y\rangle\langle z|)=0$ if $y\neq z$.

Before proving a lower bound on the entropy required for sending arbitrary n-qubit states, we first prove a lower bound on the entropy required for sending states from $\mathcal{C}_{2^m}$. Privately sending any state from $\mathcal{C}_{2^m}$ corresponds to privately sending any classical m-bit string. If communication takes place through *classical* channels, then Shannon's theorem implies that m bits of shared key are required to achieve such security. Shannon's classical lower bound does not translate automatically to the quantum world (it is in fact violated if a *two*-way quantum channel is available, see Footnote 1). Nevertheless, if Alice and Bob communicate via a one-way quantum channel, then Shannon's theorem does generalize to the quantum world:

Theorem 5.3 *If* $[\mathcal{C}_{2^m}, \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$ *is a* **PQC**, *then* $H(p_1, \ldots, p_N) \geq m$.

Proof Diagonalize the ancilla as $\rho_a = \sum_{j=1}^{r} q_j |\psi_j\rangle \langle\psi_j|$, so $S(\rho_a) = H(q_1, \ldots, q_r)$. First note that the properties of Von Neumann entropy (Section 2) imply:

$$\begin{aligned} S(\rho_0) &= S\left(\sum_{i=1}^{N} p_i U_i(|0\rangle \langle 0| \otimes \rho_a)U_i^\dagger\right) = S\left(\sum_{i=1}^{N}\sum_{j=1}^{r} p_i q_j U_i(|0\rangle \langle 0| \otimes |\psi_j\rangle \langle\psi_j|)U_i^\dagger\right) \\ &\leq H(p_1q_1, p_1q_2, \ldots, p_Nq_{r-1}, p_Nq_r) = H(p_1, \ldots, p_N) + H(q_1, \ldots, q_r). \end{aligned}$$

Secondly, note that

$$S(\rho_0) = S\left(\sum_{i=1}^{N} p_i U_i(\tilde{I}_{2^m} \otimes \rho_a)U_i^\dagger\right) \geq \sum_{i=1}^{N} p_i S\left(\tilde{I}_{2^m} \otimes \rho_a\right) = \sum_{i=1}^{N} p_i(m + S(\rho_a)) = m + S(\rho_a).$$

Combining these two inequalities gives the theorem.

In particular, for sending arbitrary states from $\mathcal{C}_{2^{2n}}$ we need entropy at least $2n$. Combining Theorems 5.2 and 5.3 we thus obtain:

Corollary 5.4 *If* $[\mathcal{H}_{2^n}, \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$ *is a* **PQC**, *then* $H(p_1, \ldots, p_N) \geq 2n$ *(and hence in particular* $N \geq 2^{2n}$*).*

In relation to Theorem 4.2, note that $\mathcal{C}_{2^n} \subseteq B^{\otimes n}$. Hence another corollary of Theorem 5.3 is the optimality of the **PQC** of Theorem 4.2:

Corollary 5.5 *If* $[B^{\otimes n}, \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$ *is a* **PQC**, *then* $H(p_1, \ldots, p_N) \geq n$ *(and hence in particular* $N \geq 2^n$*).*

6 Randomization of Quantum States

The above concepts and results were motivated by cryptographic goals, namely to enable private transmission of quantum information using a shared classical key. However, our results can also be stated in terms of the problem of "forgetting" or "randomizing" quantum information, as discussed recently by Braunstein, Lo, and Spiller [BLS99].

The *randomization* of a quantum source $\mathcal{S}$ is a procedure that maps any state ρ coming from $\mathcal{S}$ to some fixed constant state ρ_0 (for instance the completely mixed state). The process thus "forgets" what was specific to ρ. To help in this process, we allow the randomizing process to make use of a piece of the environment which is in some fixed state ρ_a (ancilla qubits). We also give it access to some source of classical randomness. Because every quantum operation can be viewed as a unitary transformation on a larger space, we can assume without loss of generality that the randomization process has the following form: it uses the source of randomness to pick some i with probability p_i, then it applies some unitary transformation U_i to ρ and the ancillary environment, and then it forgets i. The resulting mixed state should be ρ_0. At this point it should be clear to the reader that if $[\mathcal{S}, \{\sqrt{p_i}U_i \mid 1 \leq i \leq N\}, \rho_a, \rho_0]$ is a **PQC**, then it also constitutes a randomization procedure, and vice versa.

We are interested in the amount of entropy that such a randomization procedure needs to generate. This is the entropy of forgetting the random classical input i. It quantifies the thermodynamic cost of the process. Braunstein, Lo, and Spiller [BLS99] have shown that 2 bits of entropy are necessary and sufficient for the randomization of 1 qubit. By translating our **PQC**-results to the randomization-context, we can generalize their result to:

Corollary 6.1 *The generation of $2n$ bits of entropy is sufficient and necessary in order to randomize arbitrary n-qubit states.*

Proof Sufficiency follows from Theorem 4.1 and necessity from Corollary 5.4.

For the more limited set of states $\mathcal{S} = B^{\otimes n}$ we have:

Corollary 6.2 *The generation of n bits of entropy is sufficient and necessary in order to randomize arbitrary tensor products of n real-amplitude qubits.*

Proof Sufficiency follows from Theorem 4.2 and necessity from Corollary 5.5.