Fault-tolerant quantum data locking

Zixin Huang [1,2] Pieter Kok,[2] and Cosmo Lupo [2]
[1] *Center for Engineered Quantum Systems, Department of Physics and Astronomy, Macquarie University, NSW 2109, Australia*
[2] *Department of Physics and Astronomy, The University of Sheffield, Sheffield, S3 7RH, United Kingdom*

(Received 20 November 2020; accepted 28 April 2021; published 24 May 2021)

Quantum data locking is a quantum communication primitive that allows the use of a short secret key to encrypt a much longer message. It guarantees information-theoretical security against an adversary with limited quantum memory. Here we present a quantum data locking protocol that employs pseudorandom circuits consisting of Clifford gates only, which are much easier to implement fault tolerantly than universal gates. We show that information can be encrypted into n-qubit code words using order $n - H_{\min}(\mathsf{X})$ secret bits, where $H_{\min}(\mathsf{X})$ is the min-entropy of the plain text, and a min-entropy smaller than n accounts for information leakage to the adversary. As an application, we discuss an efficient method for encrypting the output of a quantum computer.

DOI: 10.1103/PhysRevA.103.052611

I. INTRODUCTION

Quantum data locking (QDL) is a quantum phenomenon that allows us to encrypt a long message using a much shorter secret key with information theoretic security. This yields one of the strongest violations of classical information theory in quantum physics. In fact, a classic result by Shannon [1], which is at the root of the one-time pad encryption, establishes that information theoretic encryption of a message of n bits requires a private key of no less than n bits.

The first QDL protocol was introduced by DiVincenzo *et al.* [2] who showed that a single secret bit is sufficient to obfuscate half of the information contained in n bits, for any n. This was obtained by encoding n bits of classical information into n qubits, where the one bit of information determines which of two mutually unbiased bases is used. Any attempt to measure the n-qubit cipher text without knowledge of the basis allows one to obtain at most $n/2$ bits of information. Further works have strengthened this seminal result [3–8]. The strongest QDL protocols can encrypt n bits of information using an exponentially small private key, with the guarantee that no more than ϵn bits will leak to the adversary. QDL was discussed in the context of quantum communications in Refs. [9–11], applications to secret key expansion and direct secret communication were introduced in Refs. [12–14], and proof-of-principle demonstrations were presented in Refs. [15,16].

In a typical QDL protocol, a short private key is used to secretly agree on a code, for example a set of basis vectors, to encode classical information into a quantum system. To encrypt information, the sender (Alice) applies a unitary transformation to map the computational basis into the chosen basis. To decrypt, the legitimate receiver (Bob) applies the inverse transformation, followed by a measurement in the computational basis. This is schematically shown in Fig. 1. To achieve secure encryption, we require that only a negligible amount of information is obtained by a unauthorized user (Eve) who attempts to measure the quantum cipher text without knowing the private key. The security of QDL holds independently of the computational capacity of Eve, who may have unlimited computational power, as long as they have limited quantum memory. For example, Eve may have no quantum memory, or a quantum memory with bounded storage time [10,12,17]. For applications in quantum cryptography, this puts QDL in the framework of bounded quantum storage [18].

We show that pseudorandom circuits can be used to build QDL protocols that are fault tolerant and robust against information leakage. In particular, we show that QDL can be realized efficiently using only Clifford gates, which can be made fault tolerant much more easily than the full universal gate set [19]. We assume that the users have the ability to apply the nonuniversal set of Clifford gates in a fault-tolerant way. We also assume that the sender can prepare states in the computational basis of n qubits, and the receiver can apply projective measurements in the computational basis. As an application, we argue that our QDL scheme can be used to encrypt the output of a quantum computer, in such a way that it is accessible only by authorized users. This encryption is secure in a scenario where quantum computing is a mature technology but quantum memories are not yet perfect.

The structure of the paper follows: In Sec. II we will review the framework of quantum data locking and introduce our protocol. In Sec. III we discuss the properties of random qubit circuits. We present our security analysis in Sec. IV, which is followed by our results in Sec. V. We discuss in detail the application of QDL to securing the output of a quantum computation in Sec. VI.

II. QUANTUM DATA LOCKING

Our scheme develops along the lines of previous QDL protocols. The protocol involves the legitimate sender Alice and the receiver Bob. The adversary is called Eve. In QDL, one may distinguish two security scenarios. In *weak* QDL, one assumes that Alice and Bob communicate through a noisy quantum channel, and Eve measures the environment of the channel. This is formally described by saying that Eve has access to the output of the conjugate channel of the channel from Alice to Bob. In *strong* QDL, one instead assumes that Eve can access the output of both the channel from Alice to Bob and its conjugate. Here we work in the strong QDL scenario. Furthermore, we consider the case when the channel from Alice to Bob is noiseless. The extension to noisy channels is still an open problem in the general case, with the exception of a handful of examples of noisy channels, including the erasure channel and the loss channel [11–14].

The QDL protocol is as follows:

(1) Alice and Bob share a unconditionally secure secret key of $\log K$ bits.

(2) They publicly agree upon a set of K n-qubit circuits, $\{C_k\}_{k=1,\dots,K}$. These circuits are composed of Clifford gates only.

(3) Alice encodes the n-bit message x into the quantum state $|x\rangle$, which belongs to the n-qubit computational basis.

(4) She then encrypts the code word and sends it to Bob. The encryption is realized by applying the Clifford circuit corresponding to the unique unitary C_k associated with the private key. Thus, the encrypted code word is $C_k|x\rangle$.

(5) Bob, who knows the private key, applies C_k^{-1}, decrypts the code word $C_k|x\rangle$, and measures in the computational basis.

Alice can chose one among $M = 2^n$ possible code words. If they have same prior probability, then the code book has maximal entropy of exactly n bits. If the code words do not have equal probabilities, then it is convenient to quantify the entropy of the code book using the min-entropy [20]

$$H_{\min}(\mathsf{X}) = -\log_2 p_{\max}. \tag{1}$$

where $p_{\max} := \max_x p_{\mathsf{X}}(x)$. A min-entropy smaller than n also describes a situation where some information about the plain text has leaked to Eve.

The security of QDL is established in a specific setting where the adversary has limited quantum storage capability. For example, Eve may have no reliable quantum memory and thus she is forced to measure the quantum state as soon as she obtains it [10]. QDL may also be secure if Eve can store quantum information reliably for a limited time, and Alice and Bob have an upper bound on her memory time [12,17].[1]

A number of QDL protocols and security proofs have been discussed in the literature. Some of them, however, would be limited to the case where $H_{\min}(\mathsf{X}) = n$ [3,5,6,12,13]. For example, Fawzi *et al.* [5,6] showed an explicit and efficient construction that can encrypt n bits of information using a key of $O(\log_2(n)\log_2(n/\epsilon))$ bits, with a leakage of no more than ϵn bits. However, this construction cannot be made fault tolerant [5]. The approach of Dupuis *et al.* [7] can instead account for nonmaximal min-entropy and would yield results similar to this work but it relies on sampling unitaries from the Haar distribution, which requires an exponential number of gates [21]. In contrast, here we are using an approximate 2-design, which can be sampled using Clifford gates only. Finally, the analysis of partial information leakage was also considered in Refs. [22] as well as in Ref. [7], however the scheme of Ref. [22] may be hard to realize in a fault-tolerant way.

[1]If Eve has a memory with a finite time, this weakens the security of the protocol, as she may obtain side information during the storage time, and then leverage it to gain more knowledge about the encrypted computation. Past works have addressed this issue in a quantitative way, assuming a model of quantum memory as a noisy channel that decoheres in time [22]. This approach may be used to quantify the security as a function of the time elapsed between when Eve receives the quantum state and when she measures it.

TABLE I. Summary of key size and circuit requirement for different schemes for encrypting the information encoded in n qubits.

	I_{acc}	Key size	Circuit class
Quantum one-time pad [23]	0	$2n$	Pauli
Approx. quantum one-time pad [3]	ϵn	$n + \log_2 n + \log_2(1/\epsilon^2)$	Pauli
Ref. [3]	$\epsilon n + 3$	$3\log_2 n$	Haar
Ref. [6]	ϵn	$O(\log_2(n/\epsilon)\log_2 n)$	universal
This paper	ϵn	$n - H_{\min}(\mathsf{X}) + O(\log_2 n) + O(\log_2(1/\epsilon))$	Clifford

Table I shows a summary of some previously known QDL protocol, compared with the contribution of this paper, the quantum one-time pad [23], and the approximate quantum one-time pad [3].

III. PSEUDORANDOM QUANTUM CIRCUITS

Unlike other works, which have considered the uniform ensemble of random unitaries induced by the Haar measure (see, e.g., Refs. [3,5,6]), here we apply pseudorandom unitaries from an approximate 2-design. This ensemble of unitaries has also been used in other applications related to information obfuscation, most notably system decoupling [24]. Using genuine Haar-random unitaries provides slightly more efficient security. However, as pointed out in Ref. [25], using unitaries from the Haar measure is prohibitively inefficient for large systems due to the exponential number of two-qubit gates and random bits required.

Recall that, given a Hilbert space of dimensions d and $\delta > 0$, a δ-approximate t-design is an ensemble of unitary operators C such that [25–27]

$$(1-\delta)M_\ell \leqslant \mathbb{E}[|\langle\alpha|C|\beta\rangle|^{2\ell}] \leqslant (1+\delta)M_\ell, \tag{2}$$

for all unit vectors $|\alpha\rangle$ and $|\beta\rangle$ in d dimensions and $\ell \leqslant t$, where $\mathbb{E}$ denotes the expectation value over the t design, and

$$M_\ell = \frac{\ell!(d-1)!}{(\ell+d-1)!} \tag{3}$$

is the ℓth moment of the uniform distribution induced by the Haar measure, i.e., $M_\ell = \mathbb{E}_{\mathrm{Haar}}[|\langle\alpha|C|\beta\rangle|^{2\ell}]$.

Given an n-qubit circuit, a δ-approximate 2-design can be achieved with $O(n(n+\log_2 1/\delta))$ two-qubit Clifford gates [28], or $O(n\log_2^2 n)$ random $U(4)$ gates [21]. There are known codes that implement the Clifford group in a fault-tolerant manner [19,29], whereas supplementing the Clifford group with fault tolerant gates into a universal set of gates is highly nontrivial [30].

The first two moments of the pseudorandom unitaries play an important role in this work, i.e., the first moment $\mathbb{E}[|\langle\alpha|C|\beta\rangle|^2]$, and the second moment $\mathbb{E}[|\langle\alpha|C|\beta\rangle|^4]$. The ratio

$$\gamma := \frac{\mathbb{E}[|\langle\alpha|C|\beta\rangle|^4]}{\mathbb{E}[|\langle\alpha|C|\beta\rangle|^2]^2} \tag{4}$$

quantifies the spread of the random variable $|\langle\alpha|C|\beta\rangle|^2$ around its average. For δ-approximate 2-designs we can bound γ from above as

$$\gamma \leqslant \frac{2d(1+\delta)}{(d+1)(1-\delta)^2} \leqslant 2\frac{1+\delta}{(1-\delta)^2}. \tag{5}$$

This coefficient will play a fundamental role in our analysis of QDL. We use the above bound on γ to estimate the length of the private key.

IV. SECURITY ANALYSIS

Our security analysis builds on, improves, and generalizes techniques previously applied in Refs. [12,13,31].

Different code words correspond to different quantum states that Alice can prepare, denoted as $|x\rangle$ (with $x = 1,\ldots,M$). These vectors are mutually orthogonal. For example, these states can be the vectors in the n-qubit computational basis. Different code words may have different prior probabilities, denoted $p_{\mathsf{X}}(x)$. Therefore, the prior uncertainty in the code words is well quantified by the min-entropy $H_{\min}(\mathsf{X}) = -\log_2 \max_x p_{\mathsf{X}}(x)$.

From the point of view of the legitimate receiver Bob, who knows the private key, the *a priori* description of the output of the computation is given by the statistical mixture

$$\rho_B = \sum_{x=1}^{M} p_{\mathsf{X}}(x)\,|x\rangle\,\langle x|\,. \tag{6}$$

The description of this state is different for Eve, who does not know the private key,

$$\rho_E = \frac{1}{K}\sum_{k=1}^{K}\sum_{x=1}^{M} p_{\mathsf{X}}(x) C_k\,|x\rangle\,\langle x|\,C_k^\dagger. \tag{7}$$

Below we show that, if K is large enough, then Eve can obtain only a negligible amount of information about the code words by measuring ρ_E.

Like other works on QDL [2,3,10,12,13,22,31], we use the accessible information $I_{\rm acc}(\mathsf{X};E)$ to quantify the potential information leakage to Eve. This quantity represents the maximum number of bits of information about the input variable X that can be obtained from a measurement of the state ρ_E. We anticipate that similar results could be obtained using other metrics, see, e.g., Refs. [5,6,8].

A measurement is a map $\mathcal{M}_{E\to\mathsf{Y}}$ that takes the quantum system E as input and has the classical variable Y as output. For any given measurement, one can consider the mutual information $I(\mathsf{X};\mathsf{Y})$ between the input message and the measurement output. Recall that the mutual information between two random variables X and Y is $I(\mathsf{X};\mathsf{Y}) = H(\mathsf{Y}) - H(\mathsf{Y}|\mathsf{X})$, where $H(\mathsf{Y}|\mathsf{X})$ is the conditional Shannon entropy. The mutual information vanishes when X and Y are statistically independent and reaches its maximum when they are perfectly correlated. The accessible information is defined as the maximum mutual information,

$$I_{\rm acc}(\mathsf{X};E) = \max_{\mathcal{M}_{E\to\mathsf{Y}}} I(\mathsf{X};\mathsf{Y}), \tag{8}$$

where the maximization is over all possible measurements $\mathcal{M}_{E\to\mathsf{Y}}$. We require that the accessible information is sufficiently small, i.e., that the information leaking to Eve is negligible not just for one particular measurement, but for all possible measurements she can perform.

The security analysis of the protocol relies on showing that $I_{\rm acc}(\mathsf{X};E)$ can be made arbitrarily small if K is large enough. This also allows us to quantify the minimal length of the private key to ensure secure encryption. To show this, we first write the accessible information as the difference of two entropies,

$$I_{\rm acc}(\mathsf{X};E) = \max_{\mathcal{M}_{E\to\mathsf{Y}}} H(\mathsf{Y}) - H(\mathsf{Y}|\mathsf{X}), \tag{9}$$

and then show that $H(\mathsf{Y}) \simeq H(\mathsf{Y}|\mathsf{X})$ for all measurements $\mathcal{M}_{E\to\mathsf{Y}}$. The proof shows that, for a random choice of K unitaries and for K large enough, one obtains $I_{\rm acc}(\mathsf{X};E) \leqslant 2n\epsilon$ with probability arbitrarily close to 1.

In general, the measurement map $\mathcal{M}_{E\to\mathsf{Y}}$ is characterized by POVM elements Λ_y, such that $\Lambda_y \geqslant 0$, $\sum_y \Lambda_y = \mathbb{I}$. It is known that the optimal measurement has unit rank [2], i.e., the POVM elements take the form $\Lambda_y = \alpha_y|\phi_y\rangle\langle\phi_y|$, where ϕ_y are unit vectors, and α_y are positive numbers such that $\sum_y \alpha_y = 2^n$.

The outcomes of the measurement are distributed according to the probability distribution

$$p_{\mathsf{Y}}(y) = \alpha_y\,\langle\phi_y|\rho_E|\phi_y\rangle\,, \tag{10}$$

with ρ_E as given in Eq. (7). For given x, the conditional probability of a measurement outcome is

$$p_{\mathsf{Y}|\mathsf{X}=x}(y) = \alpha_y\,\langle\phi_y|\rho_E^x|\phi_y\rangle\,, \tag{11}$$

with

$$\rho_E^x = \frac{1}{K}\sum_{k=1}^{K} C_k|x\rangle\langle x|C_k^\dagger. \tag{12}$$

The accessible information in Eq. (9) is then given by

$$\begin{aligned} I_{\rm acc}(\mathsf{X};E) &= \max_{\mathcal{M}_{E\to\mathsf{Y}}}\left\{-\sum_y p_{\mathsf{Y}}(y)\log_2 p_{\mathsf{Y}}(y) + \sum_{xy} p_{\mathsf{X}}(x)p_{\mathsf{Y}|\mathsf{X}=x}(y)\log_2 p_{\mathsf{Y}|\mathsf{X}=x}(y)\right\} \\ &= \max_{\mathcal{M}_{E\to\mathsf{Y}}}\sum_y \alpha_y\left\{-\langle\phi_y|\rho_E|\phi_y\rangle\log_2\langle\phi_y|\rho_E|\phi_y\rangle + \sum_x p_{\mathsf{X}}(x)\langle\phi_y|\rho_E^x|\phi_y\rangle\log_2\langle\phi_y|\rho_E^x|\phi_y\rangle\right\}. \end{aligned} \tag{13}$$

The security proof proceeds by showing that, by increasing K, both p_{Y} and $p_{\mathsf{Y}|\mathsf{X}=x}$ concentrate towards their common expectation value, and that the probability of a deviation larger than ϵ is exponentially suppressed. Therefore both the entropy $H(\mathsf{Y})$ and the conditional entropy $H(\mathsf{Y}|\mathsf{X})$ will approach the same value. We show that both terms in the curly brackets in Eq. (13) are arbitrarily close to

$$\langle\phi_y|\bar{\rho}_E|\phi_y\rangle \log_2 \langle\phi_y|\bar{\rho}_E|\phi_y\rangle \tag{14}$$

for all vectors ϕ_y, where

$$\bar{\rho}_E := 2^{-n}\mathbb{I} \tag{15}$$

is the n-qubit maximally mixed state. The relative minus sign between the terms then implies that $I(\mathsf{X};\mathsf{Y})$ can be made arbitrarily small.

First, we show using the matrix Chernoff bound [32] that ρ_E is close to the n-qubit maximally mixed state $\bar{\rho}_E := 2^{-n}\mathbb{I}$. Assuming K is large enough, with near unit probability we have

$$\rho_E \leqslant (1+\epsilon)\bar{\rho}_E = (1+\epsilon)2^{-n}\mathbb{I}. \tag{16}$$

From this inequality we find that $\langle\phi|\rho_E|\phi\rangle \leqslant (1+\epsilon)2^{-n}$ uniformly in ϕ. For a random choice of the unitaries, the probability that this inequality is violated is smaller than (see Appendix A for details)

$$P_1 := \exp\left\{n\ln 2 - K\frac{\epsilon^2}{4}\frac{2^{-n}}{p_{\max}}\right\}. \tag{17}$$

Next, we apply a tail bound due to Maurer [33]. We show that, for given ϕ and x,

$$\langle\phi|\rho_E^x|\phi\rangle \geqslant (1-\epsilon)\langle\phi|\bar{\rho}_E|\phi\rangle. \tag{18}$$

This inequality needs to be extended to all code words and to (almost) all values of x. In this way we obtain that, for a random choice of the unitaries, the inequality is verified up to a probability smaller than

$$P_2 := \exp\left[2d\ln\left(\frac{20\times 2^n}{\epsilon}\right) + \frac{\epsilon\ln M}{4p_{\max}} - \frac{K\epsilon^3}{128\gamma\, p_{\max}}\right], \tag{19}$$

where γ has been defined in Eq. (4) (see Appendix B for details).

Putting these two results together, we obtain

$$I(\mathsf{X};\mathsf{Y}) \leqslant 2\epsilon\sum_y \alpha_y 2^{-n} n. \tag{20}$$

Since $\sum_y \alpha_y 2^{-n} = 1$, we finally find

$$I(\mathsf{X};\mathsf{Y}) \leqslant 2\epsilon n. \tag{21}$$

This bound on the accessible information holds probabilistically, but the likelihood of failure can be made arbitrary small for large enough K. Specifically, the probability of failure is no larger than $P_1 + P_2$. Therefore, it can be bounded away from 1 by choosing K such that

$$K > \max\left\{\begin{array}{l}\frac{4n\times 2^n p_{\max}\ln 2}{\epsilon^2},\\ \frac{128\gamma}{\epsilon^3}\left[2^{n+1}p_{\max}\ln\left(\frac{20\times 2^n}{\epsilon}\right) + \frac{\epsilon\ln M}{4}\right].\end{array}\right. \tag{22}$$

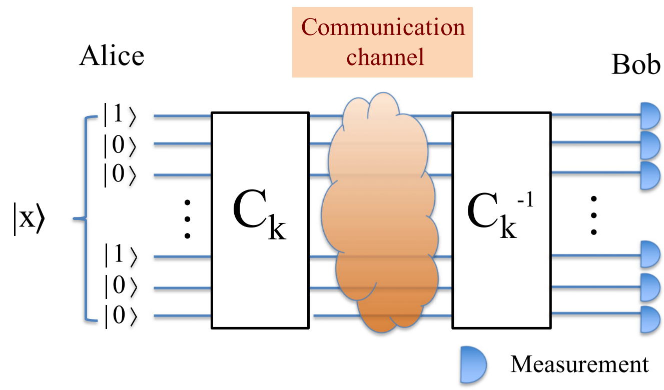

FIG. 1. Circuit layout for the encryption protocol. A useful state $|x\rangle$ is concatenated with the encryption, a pseudorandom quantum circuit C_k. The authorized user applies the unitary c_k^{-1} and correctly decrypts the encryption. An unauthorized user or adversary can attempt to extract information by performing an arbitrary measurement.

V. RESULTS

We have shown that for a random choice of K unitary transformations, the accessible information is upper bounded by a negligible number of bits $2n\epsilon$,

$$I_{\mathrm{acc}}(\mathsf{X};E) \leqslant 2n\epsilon. \tag{23}$$

From Eq. (22), this holds for a private key of length

$$\log_2 K = \log_2\gamma + n - H_{\min}(\mathsf{X}) + O(\log_2 n) + O(\log_2 1/\epsilon). \tag{24}$$

Note that the secret key length depends on the coefficient γ introduced in Eq. (4). For an approximate 2-design using the bound in Eq. (5), we obtain

$$\log_2 K \leqslant n - H_{\min}(\mathsf{X}) + \log_2\frac{1+\delta}{(1-\delta)^2} + O(\log_2 n) + O(\log_2 1/\epsilon). \tag{25}$$

In conclusions, we have shown that QDL achieves secure encryption using order $n - H_{\min}(\mathsf{X})$ secret bits, where $H_{\min}(\mathsf{X})$ is the min-entropy of the code words sent by Alice. We plot Eq. (25) in Fig. 2, where the exact value of K is given by Eq. (22), for $\epsilon = 10^{-8}$ and different values of $H_{\min}$. In the figure, we also compare our protocol with other private-key cryptography methods based on the quantum one-time pad as well as its approximate version [3]. Our protocol is more efficient, in terms of the length of the private key, when $n \gtrsim 50$, and the advantage increases with increasing n.

VI. APPLICATION: SECURING THE OUTPUT OF A QUANTUM COMPUTER

Applications of QDL have been mostly focused on quantum communication. Previous works have applied QDL, for example, to communication through a wiretapped channel. Here, we propose the use of pseudorandom quantum circuits as efficient encryption devices for protecting the output of a quantum computer. This application assumes a scenario where quantum computing is a mature technology but quantum memories are not yet perfect.

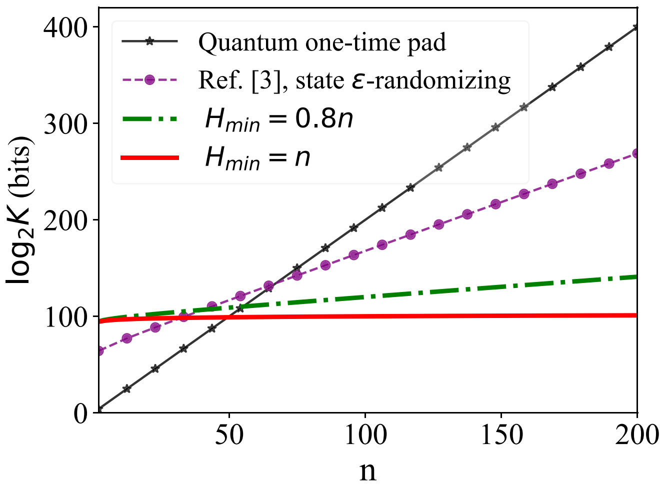

FIG. 2. Number of secret bits [$\log_2 K$ in Eq. (22)] required to lock an n-qubit output of a quantum computer, for $\epsilon = 10^{-8}$ and different values of $H_{\min}$: $H_{\min} = n$ (red solid line), $0.8n$ (green dotted-dashed line), and $0.6n$ (blue dotted line). For comparison, we plot the approximate state-randomization in Ref. [3] (purple dashed line with circles), and the quantum one-time pad [3] (black line with stars).

We imagine quantum computers as devices servicing many distributed users, where the latter may have limited computing capability, or may not know the algorithm that is realized by the server. In this scenario, we anticipate the need to encrypt the output of a quantum computer. To realize this task, we consider a protocol for private-key encryption between a quantum computer and its authorized user. This is schematically shown in Fig. 3. Unlike blind quantum computation [34–37], which is concerned with untrusted hardware and verification, our goal is to prevent unauthorized users from gaining access to the quantum computer's output.

Otherwise, one could encrypt the quantum state $|\psi\rangle$ before the measurement. Perfect encryption obtained with the quantum one-time pad would require a secret key of $2n$ bits [3]. Approximate encryption, one that encrypts the quantum state up to ϵ probability of failure, would instead require a secret key of $O(n) + O(\log_2 1/\epsilon)$ bits [3,4]. These protocols require that the encrypted state be virtually indistinguishable

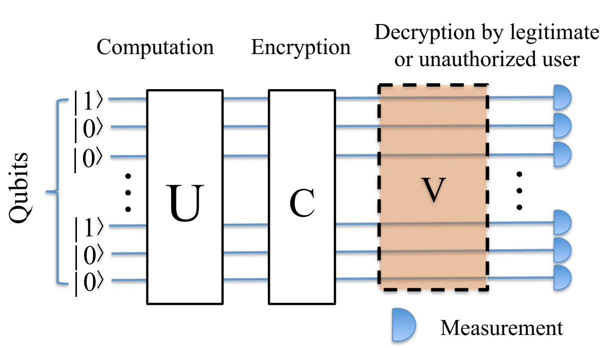

FIG. 3. Circuit layout for the encryption protocol. A useful computation U is concatenated with the encryption, a pseudorandom quantum circuit C. The authorized user applies the unitary $V = C^\dagger$ and correctly decrypts the encryption. An unauthorized user or adversary can attempt to extract information by performing an arbitrary measurement.

from the maximally mixed state. Expressed in terms of the trace norm, $\|\rho - 2^{-n}\mathbb{1}\| \leqslant \epsilon$, for some small ϵ. However, the output of a quantum computation typically contains the answer to a meaningful question. For our purposes, we may simply require that an unauthorized user does not obtain the correct answer. This opens the possibility of performing the encryption in a much more efficient way.

Suppose the quantum computer is used to solve a particular problem whose solution space has cardinality M. Different outputs of the quantum computation correspond to different quantum states, denoted $|x\rangle$ (with $x = 1, \ldots, M$). We develop our analysis within the subspace of fault-tolerant computation that incorporates quantum error correction [38,39]. Therefore, the states $|x\rangle$ are assumed to be quantum error corrected. For example, during transmission of the quantum state, channel loss will erase a subset of the transmitted qubits. Our protocol allows us to include a redundant encoding to mitigate these losses. As long as the error correction is successful, we know that there is no quantum information leakage, and our protocol remains secure.

We further assume that different outputs are associated with a prior probability $p_{\mathsf{X}}(x)$, and that the output states $|x\rangle$ are mutually orthogonal. Therefore, the uncertainty in the measurement outcome is well quantified by the min-entropy $H_{\min}(\mathsf{X}) = -\log_2 \max_x p_{\mathsf{X}}(x)$. Note that this is the prior distribution of the expected outcome of the computation. The value of the min-entropy depends on the particular computation performed by the quantum computer, and it is easy to find examples where $H_{\min}$ is low and where it is high. For example, a parity calculation may have a min-entropy as low as 1, whereas a Grover search may have a min-entropy that is close to maximal.

QDL is particularly efficient when $H_{\min}(\mathsf{X}) \approx n$, this corresponds to the setting when one has little information about the outcome of the computation. In this case $n - H_{\min}(\mathsf{X})$ can be substantially smaller than n, suggesting that the encryption can be implemented much more efficiently than previously thought.

VII. CONCLUSIONS

QDL is a communication primitive that allows us to encrypt, with information theoretic security, a long message with a much shorter private key. This is impossible in classical information theory, where the key needs to be at least as long as the message. When classical information is encoded in a quantum system, the phenomenon of QDL allows for secure encryption against an adversary with limited quantum memory, but unbounded computational power.

In this paper, we have presented a scheme for QDL that employs pseudorandom unitaries for information scrambling. These unitary transformations belong to an approximate unitary 2-design. In particular, the unitaries can be obtained by combining Clifford gates. This is an improvement with respect to previous QDL schemes because fault-tolerant Clifford gates require orders of magnitude fewer physical qubits than universal fault-tolerant quantum computing [30,40]. Furthermore, our QDL protocol allows for partial information leakage to the eavesdropper. This is modeled by the code words having a nonmaximal min-entropy.

We discuss an application of our QDL protocol as a way to encrypt the output of a quantum computer. Unlike blind quantum computation, which is concerned with untrusted hardware and verification, we focus on preventing unauthorized users from gaining access to the output of a quantum algorithm. We have considered a scenario where a server can realize fault-tolerant universal quantum computing, the user is capable only of implementing fault-tolerant Clifford gates and measurements in the computational basis, and the eavesdropper has limited quantum memory.

ACKNOWLEDGMENTS

Z.H. acknowledges Ryan Mann for fruitful discussions. We also thank Earl Campell, Armanda Ottaviano-Quintavalle, Joschka Roffe, Omar Fawzi, Dominik Hangleiter, and the anonymous referees for insightful comments on the paper. This work was supported by the EPSRC Quantum Communications Hub, Grant No. EP/M013472/1.

APPENDIX A: APPLICATION OF THE MATRIX CHERNOFF BOUND

The matrix Chernoff bound states the following (which can be obtained directly from Theorem 19 of Ref. [32]):

Theorem 1. Let $\{X_t\}_{t=1,\dots,K}$ be K independently and identically distributed d-dimensional Hermitian-matrix-valued random variables, with $X_t \approx X$, $0 \leqslant X \leqslant R$, and $\mathbb{E}[X] = 2^{-n}\mathbb{I}$. Then, for any $\epsilon \geqslant 0$,

$$\Pr\left\{\frac{1}{K}\sum_{t=1}^{K} X_t \nleqslant (1+\epsilon)\mathbb{E}[X]\right\} \leqslant d \exp\left\{-KD\left[(1+\epsilon)\frac{2^{-n}}{R}\middle\|\frac{2^{-n}}{R}\right]\right\}, \tag{A1}$$

where $\Pr\{x\}$ denotes the probability that the proposition x is true, and $D[u\|v] = u\ln(u/v) - (1-u)\ln[(1-u)/(1-v)]$. Note that, for $\epsilon < 1$,

$$D\left[(1+\epsilon)\frac{2^{-n}}{R}\middle\|\frac{2^{-n}}{R}\right] \geqslant \frac{\epsilon^2}{4}\frac{2^{-n}}{R}. \tag{A2}$$

We apply the Chernoff bound to the K independent random variables

$$X_k \equiv C_k \sum_{x=1}^{M} p_{\mathsf{X}}(x)|\psi_x\rangle\langle\psi_x|C_k^\dagger. \tag{A3}$$

Note that these operators satisfy $0 \leqslant X_k \leqslant p_{\max} := \max_x p_{\mathsf{X}}(x)$. Therefore, $R \equiv p_{\max}$. Also note that

$$\frac{1}{K}\sum_{k=1}^{K} X_k = \frac{1}{K}\sum_{k=1}^{K} C_k \sum_{x=1}^{M} p_{\mathsf{X}}(x)|\psi_x\rangle\langle\psi_x|C_k^\dagger = \rho_{\mathsf{U}'}, \tag{A4}$$

and $\mathbb{E}[X] = \bar{\rho}_{\mathsf{U}'} = 2^{-n}\mathbb{I}$. By applying the Chernoff bound we then obtain

$$\Pr\{\rho_{\mathsf{U}'} \nleqslant (1+\epsilon)2^{-n}\} \leqslant 2^n \exp\left\{-K\frac{\epsilon^2}{4}\frac{2^{-n}}{p_{\max}}\right\} \tag{A5}$$

$$= \exp\left\{n\ln 2 - K\frac{\epsilon^2}{4}\frac{2^{-n}}{p_{\max}}\right\}. \tag{A6}$$

In conclusion, we have obtained that, up to a probability smaller than

$$P_1 := \exp\left\{n\ln 2 - K\frac{\epsilon^2}{4}\frac{2^{-n}}{p_{\max}}\right\}, \tag{A7}$$

the following matrix inequality holds:

$$\rho_{\mathsf{U}'} \leqslant (1+\epsilon)2^{-n}. \tag{A8}$$

APPENDIX B: APPLICATION OF THE MAURER BOUND

We apply a concentration inequality obtained by Maurer in Ref. [33]:

Theorem 2. Let $\{X_k\}_{k=1,\dots,K}$ be K independent and identically distributed non-negative real-valued random variables, with $X_k \approx X$ and finite first and second moments, $\mathbb{E}[X], \mathbb{E}[X^2] < \infty$. Then, for any $\tau > 0$ we have that

$$\Pr\left\{\frac{1}{K}\sum_{k=1}^{K} X_k < (1-\tau)\mathbb{E}[X]\right\} \leqslant \exp\left(-\frac{K\tau^2\mathbb{E}[X]^2}{2\mathbb{E}[X^2]}\right). \tag{B1}$$

For any given x and ϕ, we apply this bound to the random variables

$$X_k \equiv |\langle\phi|C_k|\psi_x\rangle|^2. \tag{B2}$$

Note that

$$\frac{1}{K}\sum_{k=1}^{K} X_k = \langle\phi|\rho_{\mathsf{U}'}^x|\phi\rangle, \tag{B3}$$

and

$$\mathbb{E}[X] = \bar{\rho}_{\mathsf{U}'} = 2^{-n}\mathbb{I}. \tag{B4}$$

The application of the Maurer tail bound then yields

$$\Pr\{\langle\phi|\rho_{\mathsf{U}'}^x|\phi\rangle < (1-\tau)2^{-n}\} \leqslant \exp\left(-\frac{K\tau^2}{2\gamma}\right), \tag{B5}$$

with γ as defined in Eq. (4).

The probability bound in Eq. (B5) refers to one given value of x. Here we extend it to $\ell < M$ distinct values $x_1, x_2, \dots, x_\ell$. We have

$$\Pr\left\{\forall\, x = x_1, x_2, \dots x_\ell,\ \langle\phi|\rho_{\mathsf{U}'}^x|\phi\rangle < (1-\tau)2^{-n}\right\} \leqslant \exp\left(-\frac{\ell K\tau^2}{2\gamma}\right). \tag{B6}$$

This follows from two observations. First, for different values of x, the random variables $\langle\phi|\rho_{\mathsf{U}'}^x|\phi\rangle$ are identically distributed. Second, these variables are not statistically independent because they obey the subnormalization constraint $\sum_x \langle\phi|\rho_{\mathsf{U}'}^x|\phi\rangle = c \leqslant 1$. If the ℓ random variables $x_1, x_2, \dots, x_\ell$ were statistically independent, then Eq. (B6) would hold. However, Eq. (B6) still holds because the normalization constraint implies that the variables are anticorrelated. Therefore, the probability that they are all small is smaller than if they were statistically independent.

We now extend the concentration inequality to all possible choices of ℓ values of x. This amount to a total of $\binom{M}{\ell}$ events. Applying the union bound we obtain

$$\Pr\{\exists\, x_1, x_2, \dots x_\ell, |\forall\, x = x_1, x_2, \dots x_\ell,\ \langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle < (1-\tau)2^{-n}\} \leqslant \binom{M}{\ell} \exp\left(-\frac{\ell K \tau^2}{2\gamma}\right). \tag{B7}$$

This implies that up to a probability smaller than $\binom{M}{\ell} \exp\left(-\frac{\ell K\tau^2}{2\gamma}\right)$, $\langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \geqslant (1-\tau)2^{-n}$ for at least $M-\ell$ values of x, which yields

$$\sum_{x=1}^{M} p_{\mathsf{X}}(x)\langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \log_2 \langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \leqslant \left(\sum_{x\in S_{M-\ell}} p_{\mathsf{X}}(x)\right)(1-\tau)2^{-n} \log_2 (1-\tau)2^{-n}, \tag{B8}$$

where $S_{M-\ell}$ denotes the set of $M-\ell$ least likely values of x. Note that

$$\sum_{x\in S_{M-\ell}} p_{\mathsf{X}}(x) = 1 - \sum_{x\in L_\ell} p_{\mathsf{X}}(x) \geqslant 1 - \ell p_{\max}, \tag{B9}$$

where L_ℓ is the subset of the ℓ most likely values of x, and $p_{\max} = \max_x p_{\mathsf{X}}(x)$. Putting this into Eq. (B8) yields

$$\begin{aligned}\sum_{x=1}^{M} p_{\mathsf{X}}(x)\langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \log_2 \langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle &\leqslant (1-\ell\, p_{\max})(1-\tau)2^{-n} \log_2 (1-\tau)2^{-n} \\ &\leqslant -(1-\ell\, p_{\max})(1-\tau)2^{-n} n.\end{aligned} \tag{B10}$$

Finally, putting $\ell = \tau/p_{\max}$ we obtain

$$\sum_{x=1}^{M} p_{\mathsf{X}}(x)\langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \log_2 \langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle \leqslant (1-\tau)^2 2^{-n} n = (1-2\tau)2^{-n} n + O(\tau^2). \tag{B11}$$

To extend to all vectors ϕ, we exploit the notion of δ net and closely follows Ref. [3]. In this way we obtain

$$\begin{aligned}&\Pr\{\forall\, \phi\ \exists\, x_1, x_2, \dots x_\ell, |\forall\, x = x_1, x_2, \dots x_\ell,\ \langle\phi|\rho^x_{\mathsf{U}'}|\phi\rangle < (1-2\tau)2^{-n}\} \\ &\leqslant \left(\frac{5\times 2^n}{\tau}\right)^{2d} \binom{M}{\ell} \exp\left(-\frac{\ell K\tau^2}{2\gamma}\right).\end{aligned} \tag{B12}$$

[1] C. E. Shannon, Bell Syst. Tech. J. **28**, 656 (1949).
[2] D. P. DiVincenzo, M. Horodecki, D. W. Leung, J. A. Smolin, and B. M. Terhal, Phys. Rev. Lett. **92**, 067902 (2004).
[3] P. Hayden, D. Leung, P. W. Shor, and A. Winter, Commun. Math. Phys. **250**, 371 (2004).
[4] G. Aubrun, arXiv:0802.4193.
[5] O. Fawzi, Ph.D. thesis, McGill University, School of Computer Science, 2011 (unpublished); M. Berta, O. Fawzi, and S. Wehner, IEEE Trans. Info. Theo. **60**, 1168 (2014).
[6] O. Fawzi, P. Hayden, and P. Sen, J. Assoc. Comput. Mach. **60**, 44 (2013).
[7] F. Dupuis, M. Berta, J. Wullschleger, and R. Renner, Proc. R. Soc. London, Ser. A **469**, 2159 (2013).
[8] R. Adamczak, J. Phys. A: Math. Theor. **50**, 105302 (2017).
[9] S. Lloyd, arXiv:1307.0380.
[10] S. Guha, P. Hayden, H. Krovi, S. Lloyd, C. Lupo, J. H. Shapiro, M. Takeoka, and M. M. Wilde, Phys. Rev. X **4**, 011016 (2014).
[11] A. Winter, J. Cryptol. **30**, 1 (2017).
[12] C. Lupo and S. Lloyd, Phys. Rev. Lett. **113**, 160502 (2014).
[13] C. Lupo and S. Lloyd, New J. Phys. **17**, 033022 (2015).
[14] C. Lupo and S. Lloyd, Phys. Rev. A **92**, 062312 (2015).
[15] Y. Liu, Z. Cao, C. Wu, D. Fukuda, L. You, J. Zhong, T. Numata, S. Chen, W. Zhang, S.-C. Shi, C.-Y. Lu, Z. Wang, X. Ma, J. Fan, Q. Zhang, and J.-W. Pan, Phys. Rev. A **94**, 020301(R) (2016).
[16] D. J. Lum, J. C. Howell, M. S. Allman, T. Gerrits, V. B. Verma, S. W. Nam, C. Lupo, and S. Lloyd, Phys. Rev. A **94**, 022315 (2016).
[17] C. Lupo, Entropy **17**, 3194 (2015).
[18] I. B. Damgaard, S. Fehr, R. Renner, L. Salvail, and C. Schaffner, in *Advances in Cryptology - CRYPTO 2007*, Lecture Notes in Computer Science, edited by A. Menezes (Springer, New York, 2007), Vol. 4622.
[19] B. Eastin and E. Knill, Phys. Rev. Lett. **102**, 110502 (2009).
[20] M. Tomamichel, Ph.D. thesis, ETH Zurich, 2012, arXiv:1203.2142.
[21] A. Harrow and S. Mehraban, arXiv:1809.06957.
[22] C. Lupo, M. M. Wilde, and S. Lloyd, Phys. Rev. A **90**, 022326 (2014).
[23] A. Ambainis, M. Mosca, A. Tapp, and R. De Wolf, in *Proceedings 41st Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, 2000), pp. 547–553.
[24] O. Szehr, F. Dupuis, M. Tomamichel, and R. Renner, New J. Phys. **15**, 053022 (2013).
[25] F. G. Brandao, A. W. Harrow, and M. Horodecki, Commun. Math. Phys. **346**, 397 (2016).

[26] A. Roy and A. J. Scott, Des. Codes Cryptogr. **53**, 13 (2009).
[27] R. L. Mann and M. J. Bremner, arXiv:1711.00686.
[28] A. W. Harrow and R. A. Low, Commun. Math. Phys. **291**, 257 (2009).
[29] B. Zeng, A. Cross, and I. L. Chuang, IEEE Trans. Inf. Theory **57**, 6272 (2011).
[30] E. T. Campbell, B. M. Terhal, and C. Vuillot, Nature (London) **549**, 172 (2017).
[31] Z. Huang, P. P. Rohde, D. W. Berry, P. Kok, J. P. Dowling, and C. Lupo, Quantum **5**, 447 (2021).
[32] R. Ahlswede and A. Winter, IEEE Trans. Inf. Theory **48**, 569 (2002).
[33] A. Maurer *et al.*, J. Inequal. Pure Appl. Math. **4**, 15 (2003).
[34] J. F. Fitzsimons, npj Quantum Inf. **3**, 23 (2017).
[35] T. Morimae and K. Fujii, Nat. Commun. **3**, 1036 (2012).
[36] A. Broadbent, J. Fitzsimons, and E. Kashefi, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (IEEE, 2009), pp. 517–526.
[37] V. Dunjko, J. F. Fitzsimons, C. Portmann, and R. Renner, in *Advances in Cryptology – ASIACRYPT 2014*, edited by P. Sarkar and T. Iwata (Springer, Berlin, Heidelberg, 2014), pp. 406–425.
[38] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge University Press, 2013).
[39] J. Roffe, Contemp. Phys. **60**, 226 (2019).
[40] J. O'Gorman and E. T. Campbell, Phys. Rev. A **95**, 032338 (2017).