Quantum state restoration and single-copy tomography

Edward Farhi,* David Gosset,† Avinatan Hassidim,‡ and Andrew Lutomirski§
Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139

Daniel Nagaj¶
Research Center for Quantum Information, Institute of Physics, Slovak Academy of Sciences, Dúbravská cesta 9, 845 11 Bratislava, Slovakia

Peter Shor**
Department of Mathematics and Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139
(Dated: December 18, 2009)

Given a single copy of a quantum state $|\psi\rangle$, the no cloning theorem greatly limits the amount of information which can be extracted from it. On the other hand, given only a procedure which verifies the state, for example access to a measurement $M = \{\mathbb{I} - |\psi\rangle\langle\psi|, |\psi\rangle\langle\psi|\}$, even distinguishing $\mathbb{I} - |\psi\rangle\langle\psi|$ from the identity takes exponential time. In this paper, we consider the scenario in which we are given both a single copy of $|\psi\rangle$and the ability to verify it. We show that taken together, these primitives enable us to efficiently learn about $|\psi\rangle$. In particular, for any POVM (even with non commuting operators) we give an algorithm which estimates its statistics on $|\psi\rangle$, in time polynomial in the number of operators. We show how this algorithm puts severe limitations on possible quantum money shcemes.

I. INTRODUCTION

Quantum mechanics places constraints on what can be done with a single copy of an unknown state. For example, the no-cloning theorem says that it is impossible to copy such a state. Another implication of the theorem is that measuring an observable on an unknown state generally damages it. As an extreme example, learning the full description of the state or even the description of a small piece of the state cannot be done with a single copy of it.

In this paper, we are interested in the additional power given by the ability to verify a state. Given a single copy of an unknown quantum state $|\psi\rangle$ and a verifier—that is, a black-box which measures the operator $P = |\psi\rangle\langle\psi|$—the no-cloning theorem no longer applies, and new possibilities arise. In this setting, we present novel algorithms that can copy small parts of the state and make measurements on $|\psi\rangle$ without damaging the state. One situation where such a verifier exists is when $|\psi\rangle$ is known to be the ground state of a particular gapped local Hamiltonian. In this case, we can use phase estimation to measure the ground state energy E_0, and we can subsequently verify a state by measuring the energy again and confirming that it has the same value.[12]

We call our first algorithm quantum state restoration. This algorithm allows one to set aside any small subsystem of an unknown state $|\psi\rangle$, and efficiently restore it using only the verifier. That is, if $|\psi\rangle$ lives in the Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, our algorithm takes as input the part of $|\psi\rangle$ that lives in subsystem A (described by the reduced density matrix ρ_A) and produces as output the state $|\psi\rangle$. This can be used to copy small subsystems of $|\psi\rangle$: if $|\psi\rangle$ has the reduced density matrix ρ_B on a small subsystem B, we can run our algorithm on subsystem A, producing as output both $|\psi\rangle$ and the mixed state ρ_B. If we use this to obtain multiple copies of ρ_B, we can perform tomography on subsystem B. We call this application single-copy tomography, and we give two more specialized algorithms to do the same thing. All these algorithms have running time polynomial in the dimension of subsystem B. We also give a reduction from estimating the statistics of a general POVM measurement (even if it includes non-commuting operators) to single-copy tomography, with running time polynomial in the number of POVM operators.

Our original motivation for developing these algorithms was to understand the security of a class of public-key quantum money schemes. Public-key quantum money is a quantum state that can be produced by a bank and

*Electronic address: farhi@mit.edu
†Electronic address: dgosset@mit.edu
‡Electronic address: avinatan@mit.edu
§Electronic address: luto@mit.edu
¶Electronic address: daniel.nagaj@savba.sk
**Electronic address: shor@math.mit.edu

verified by anyone—ideally, the verification algorithm is a projector onto the state in question [1, 2, 7]. The definition of quantum money requires that no one other than the bank can efficiently produce states that pass verification, and when a state passes verification it is returned undamaged by the procedure. Whether or not secure quantum money protocols exist is an open question. However, algorithms such as quantum state restoration and single-copy tomography rule out a large class of possible quantum money schemes.

The simplest example of a quantum money scheme that is broken by our algorithm is based on product states. The bank chooses a string of n uniformly random angles θ_i between 0 and 2π. This string is a classical secret known only to the bank. Using these angles, the bank generates the state $|\psi\rangle = \otimes_i |\theta_i\rangle$ where $|\theta_i\rangle = \cos\theta_i|0\rangle + \sin\theta_i|1\rangle$ and chooses a set of (say) 4-local projectors $\{P_i\}$ which are all orthogonal to $|\psi\rangle$. This set is chosen to be large enough so that $|\psi\rangle$ is the only state in the intersection of the zero eigenspaces of all of the projectors. The quantum money consists of the state $|\psi\rangle$ and a classical description of the projectors[13]. The bank must choose a new set of angles $\{\theta_i\}$ for each quantum money state it produces; otherwise standard tomography can break this protocol. Anyone can verify the money by measuring the projectors. Since a good money state is an eigenstate of the projectors, the measurement passes along good money undamaged.

At first glance, product state quantum money seems promising. First, given only the state $|\psi\rangle$, the no-cloning theorem prevents anyone from making a second copy. In general, given only a set of 4-local projectors, the problem of finding the corresponding angles (if they exist) is NP-complete (although in our case the projectors are chosen from a specific distribution and there is a planted solution, so the problem may be easier). However, given both the state $|\psi\rangle$ and the projectors, $|\psi\rangle$ can be efficiently copied using quantum state restoration. We use the quantum money's verifier as our projector $P = |\psi\rangle\langle\psi|$. We can then copy the qubits one at a time. To copy the first qubit, a simplified version of quantum state restoration proceeds as follows:

1. Set aside the first qubit. We are left with the state $|\theta_2\rangle \cdots |\theta_n\rangle$.

2. Add a new register at the beginning containing a random one-qubit state. We now have a state which can be written as

$$\left(\alpha\,|\theta_1\rangle + \beta\,\left|\theta_1^{\perp}\right\rangle\right)|\theta_2\rangle \cdots |\theta_n\rangle$$

   where $\langle\theta_i|\theta_i^{\perp}\rangle = 0$ and α and β are unknown random variables.

3. Verify the quantum money. This produces either the desired state $|\psi\rangle$ or an invalid quantum money state

$$\left|\theta_1^{\perp}\right\rangle|\theta_2\rangle \cdots |\theta_n\rangle$$

   with equal probability (averaged over the choice of the random state in step 2). If we have produced the desired state, then we have cloned the first qubit: we have both the copy in the $|\psi\rangle$ and the qubit that we set aside in step 1. If not, then we discard the qubit $|\theta_1^{\perp}\rangle$ and go back to step two and repeat until we get $|\theta_1\rangle$.

Repeating this procedure for each qubit allows us to clone the state $|\psi\rangle$ in linear time.

The algorithm we just described can copy the full n-qubit state $|\psi\rangle$ because $|\psi\rangle$ is a product state. We can think of this algorithm as first removing a subsystem of $|\psi\rangle$ (step 1) and then recovering the state $|\psi\rangle$ from the part that remains (steps 2 and 3). Surprisingly, a small modification of steps 2 and 3 leads to an algorithm that efficiently restores small missing subsystems, even on entangled states: this is quantum state restoration.

Our paper is structured as follows. In section II, we present the quantum state restoration algorithm and analyze its running time. In section III, we present two alternative algorithms for single-copy tomography, one of which is asymptotically faster than quantum state restoration. Finally, we give several scenarios in which new algorithms could be developed using the techniques in this paper.

II. QUANTUM STATE RESTORATION

Quantum state restoration takes as input a large subsystem of a state $|\psi\rangle$ (this subsystem could be, for example, the first $n-k$ qubits of the n qubit state $|\psi\rangle$) and, using the ability to measure the projector $P = |\psi\rangle\langle\psi|$, reconstructs the full state $|\psi\rangle$.

Theorem 1. *Suppose that $|\psi\rangle$ is an unknown quantum state in a Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ and we are given oracle access to a coherent measurement of $P = |\psi\rangle\langle\psi|$ (that is, the oracle performs the operation $P \otimes \mathbb{I} + (1-P) \otimes \sigma_x$ on the original Hilbert space plus a single-qubit ancilla). Then there exists an efficient quantum algorithm that takes as input a mixed state in $\mathcal{H}_A$ with density matrix $\mathrm{Tr}_B\,|\psi\rangle\langle\psi|$ and outputs $|\psi\rangle$. This algorithm makes an expected number $O\left((\dim\mathcal{H}_B)^2\right)$ of calls to the measurement oracle.*

The idea is that any state $|\psi\rangle$ on a Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ (where d is the dimension of $\mathcal{H}_B$) can be Schmidt decomposed as

$$|\psi\rangle = \sum_{i=1}^{\chi} \sqrt{p_i}|u_i\rangle|v_i\rangle$$

where χ is the Schmidt rank of $|\psi\rangle$ (note that $\chi \leq d$). If we start with the state $|\psi\rangle$ and set aside the part that lives on $\mathcal{H}_B$, then we are left with the mixed state $\rho_A = \mathrm{Tr}_B |\psi\rangle\langle\psi|$, which has all of its support on the Schmidt basis span $\{|u_i\rangle\}$. From ρ_A, we can construct the state $\rho_A \otimes \frac{I}{d}$ on $\mathcal{H}_A \otimes \mathcal{H}_B$. We now measure the projector P. If we obtain the outcome 1, then we are left with the state $|\psi\rangle$, If not, we discard (i.e. trace out) $\mathcal{H}_B$, leaving a state on $\mathcal{H}_A$ that *still* has all of its support on the Schmidt basis. We then try again until we obtain the outcome 1. If all the p_i are equal, then each attempt succeeds with probability $\frac{1}{\chi d}$, and the entire algorithm finishes in an expected number of iterations χd. For general values $\{p_i\}$, the expected running time is still exactly χd, although the distribution of the running time becomes more complicated.

We now summarize the quantum state restoration algorithm.

1. Start with the state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ and set aside the part of $|\psi\rangle$ that lives in subsystem B. We are left with the mixed state

$$\rho_A = \operatorname*{Tr}_B |\psi\rangle\langle\psi|.$$

2. Add a random state on subsystem B. The state is now

$$\rho_A \otimes \frac{\mathbb{I}}{d}.$$

3. Measure the projector $P = |\psi\rangle\langle\psi|$. If the outcome is $+1$ then you are done: you still have the original copy of subsystem B that you set aside and you have recovered the state $|\psi\rangle$. If not, discard subsystem B and repeat from step 2.

We now show that the expected running time of this algorithm is $\chi \cdot d \leq d^2$ (measured in number of uses of P).

A. Running Time of Quantum State Restoration

In the simple case where all of the p_i are equal, then the initial state ρ_A is the fully mixed state over the span $\{|u_i\rangle\}$. In this case, if you measure 0 in step 3, the density matrix left in register A after discarding register B is unchanged. The algorithm terminates with probability $\frac{1}{\chi \cdot d}$ on each iteration, finishing in an expected number of iterations $\chi \cdot d$. If the p_i are not all equal, then the algorithm can reach bad states where most of the weight is on low-weight elements of the Schmidt basis. When this happens, the chance of success on any given iteration drops (see figure 1 for an extreme example), but the probability of reaching these bad states decreases with the corresponding p_i. Surprisingly, these effects exactly cancel, and the expected number of iterations required to restore the state is $\chi \cdot d$ regardless of the values of the p_i.

To prove this, we define two maps

$$F_0(\sigma) = \operatorname*{Tr}_B \left[(1 - |\psi\rangle\langle\psi|) \left(\sigma \otimes \frac{\mathbb{I}}{d} \right) (1 - |\psi\rangle\langle\psi|) \right]$$
$$F_1(\sigma) = \operatorname*{Tr}_B \left[|\psi\rangle\langle\psi| \left(\sigma \otimes \frac{\mathbb{I}}{d} \right) |\psi\rangle\langle\psi| \right].$$

Here $F_b(\sigma)$ is the unnormalized density matrix obtained by measuring P on the state given by the density matrix σ, conditioned on the measurement outcome $b \in \{0, 1\}$. The probability of obtaining a sequence of measurement outcomes $b_1, b_2, \ldots, b_m$, starting with the state σ is then given by

$$\Pr\left[\{b_1, b_2, b_3, ..., b_m\} \,\middle|\, \sigma\right] = \mathrm{Tr}[F_{b_m} \circ \cdots \circ F_{b_1}(\sigma)], \tag{1}$$

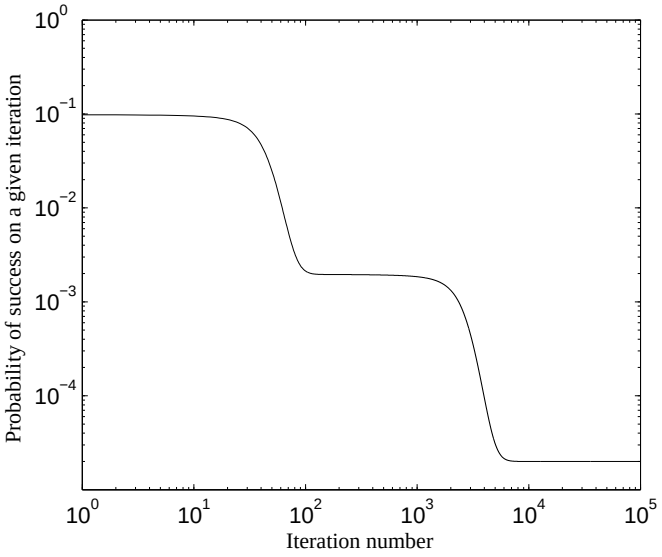

Figure 1: Probability of restoring the state on a given iteration conditioned on all previous iterations failing. Conditioned on failing every time, the first two flat regions are metastable states and the third is stable. In this graph, $|\psi\rangle = \sqrt{1-10^{-2}-10^{-4}}|0\rangle_A|0\rangle_B + \sqrt{10^{-2}}|1\rangle_A|1\rangle_B + \sqrt{10^{-4}}|2\rangle_A|2\rangle_B$, $\dim \mathcal{H}_B = 10$, and the expected number of iterations required is 30.

which can be seen by induction:

$$\begin{aligned}
\Pr\left[\{b_1, b_2, b_3, ..., b_m\} \,\middle|\, \sigma\right] &= \Pr\left[b_m \,\middle|\, \sigma, \{b_1, b_2, b_3, ..., b_{m-1}\}\right] \\
&\quad \times \Pr\left[\{b_1, b_2, b_3, ..., b_{m-1}\} \,\middle|\, \sigma\right] \\
&= \operatorname{Tr} F_{b_m}\left(\frac{F_{b_{m-1}} \circ \cdots \circ F_{b_1}(\sigma)}{\operatorname{Tr}\left(F_{b_{m-1}} \circ \cdots \circ F_{b_1}(\sigma)\right)}\right) \\
&\quad \times \operatorname{Tr}\left(F_{b_{m-1}} \circ \cdots \circ F_{b_1}(\sigma)\right) \\
&= \operatorname{Tr} F_{b_m}\left(F_{b_{m-1}} \circ \cdots \circ F_{b_1}(\sigma)\right).
\end{aligned}$$

We can use this equation to write an explicit formula for the expected number of measurements $T(\sigma)$, starting with the state σ:

$$\begin{aligned}
T(\sigma) &= \sum_{k=1}^{\infty} k \cdot \Pr\left[\{\underbrace{0, 0, \ldots, 0}_{k-1}, 1\} \,\middle|\, \sigma\right] \\
&= \sum_{k=1}^{\infty} k \cdot \operatorname{Tr}\left[F_1 \circ F_0 \circ \cdots \circ F_0(\sigma)\right]. \qquad (2)
\end{aligned}$$

As written, this formula is difficult to evaluate, but we can see that it is linear in σ. We are interested in the quantity $T(\rho_A)$, which we expand as

$$T(\rho_A) = \sum_{i=1}^{\chi} p_i T(|u_i\rangle\langle u_i|). \qquad (3)$$

We expand $T(|u_i\rangle\langle u_i|)$ by conditioning on the outcome of the first measurement:

$$\begin{aligned} T\left(|u_i\rangle\langle u_i|\right) &= \Pr\left[1 \,\middle|\, |u_i\rangle\langle u_i|\right] + \Pr\left[0 \,\middle|\, |u_i\rangle\langle u_i|\right] \left(1 + T\left(\frac{F_0(|u_i\rangle\langle u_i|)}{\Pr\left[0 \,\middle|\, |u_i\rangle\langle u_i|\right]}\right)\right) \\ &= 1 + T\left(F_0(|u_i\rangle\langle u_i|)\right) \\ &= 1 + T\left(|u_i\rangle\langle u_i| - 2\frac{p_i}{d}|u_i\rangle\langle u_i| + \frac{p_i}{d}\sum_{j=1}^{\chi} p_j |u_j\rangle\langle u_j|\right) \\ &= 1 + \left(1 - 2\frac{p_i}{d}\right) T\left(|u_i\rangle\langle u_i|\right) + \frac{p_i}{d}\sum_{j=1}^{\chi} p_j T\left(|u_j\rangle\langle u_j|\right). \end{aligned}$$

Using (3), this can be transformed into

$$2p_i T\left(|u_i\rangle\langle u_i|\right) - p_i T(\rho_A) = d.$$

Summing both sides over $i = 1, \ldots, \chi$ using $\sum p_i = 1$ and (3) again, we obtain

$$T(\rho_A) = \chi \cdot d,$$

which is the desired result. This proves Theorem 1.

III. SINGLE-COPY TOMOGRAPHY AND ESTIMATION OF MEASUREMENT STATISTICS

We expect that quantum state restoration will most commonly be used to perform tomography on a single copy of a verifiable quantum state. We can perform several different types of tomography, and we give algorithms for some types that are faster than quantum state restoration.

a. General tomography on a subsystem In the simplest case, we have a single copy of an unknown state $|\psi\rangle$ and access to the measurement $P = |\psi\rangle\langle\psi|$ and we would like to estimate properties of the density matrix $\rho_B = \mathrm{Tr}_A\,|\psi\rangle\langle\psi|$ for a subsystem B. We can do this by using quantum state restoration to prepare many unentangled states, each with (independent) density matrixes ρ_B. We can then use any standard state tomography algorithm on these states.

b. Measurement of a subsystem in an orthogonal basis For many applications, it is sufficient to estimating the probabilities $q_i = \mathrm{Tr}\left[|i\rangle_B\langle i|_B|\psi\rangle\langle\psi|\right]$ of obtaining the outcome i if one were to measure subsystem B of $|\psi\rangle$ in the orthonormal basis $\{|i\rangle_B\}$. Quantum state restoration can sample these probabilities directly. We discuss this application in section III A.

In sections III A and III B, we present two other specialized algorithms to compute these probabilities. Both algorithms measure the q_i one at a time by considering the statistics of the two-outcome measurements $\{|i\rangle_B\langle i|_B, \mathbb{I} - |i\rangle_B\langle i|_B\}$, and both are based on previously presented schemes for amplifying QMA verifiers [6, 8].

In each case, we fix a precision $\delta > 0$ and an error probability $\epsilon > 0$ and compute the running time to produce estimates q_i^{est} such that

$$|q_i^{\mathrm{est}} - q_i| < \delta$$

for all i with probability at least $1 - \epsilon$.

c. Estimation of the statistics of any POVM We can use any of our algorithms to estimate the statistics of a general measurement (on the complete state, not just a subsystem). This is because a general POVM measurement can be reduced to a measurement of a subsystem in an orthogonal basis, as we now review. Given an efficiently implementable POVM $\{E_i\}$ where $i \in \{1, \ldots, d\}$, we can implement a unitary operator U such that

$$U\left(|\phi\rangle_A|1\rangle_B\right) = \sum_{i=1}^{d}\left(\sqrt{E_i}|\phi\rangle_A\right)|i\rangle_B$$

for any state $|\phi\rangle$. If we work in a two-register Hilbert space, where register A can hold $|\phi\rangle$ and register B has dimension d, then the probability of measurement outcome i when the POVM is measured on $|\phi\rangle$ is equal to

$$\langle\phi|E_i|\phi\rangle = \mathrm{Tr}\left[\rho_B|i\rangle_B\langle i|_B\right]$$

where $\rho_B = \mathrm{Tr}_A \left[U|\phi\rangle_A|1\rangle_B\langle 1|_B\langle\phi|_A U^\dagger\right]$. If we define

$$|\psi\rangle = U|\phi\rangle_A|1\rangle_B$$
$$P' = |\psi\rangle\langle\psi| = UPU^\dagger$$

then $|\psi\rangle$ can be efficiently prepared (given $|\phi\rangle$) and P' can be efficiently measured. Now we can use any of the algorithms to estimate the measurement statistics of subsystem B of $|\psi\rangle$ using the projector P' in the computational basis (that is, any of the algorithms below) to estimate the probabilities $\langle\phi|E_i|\phi\rangle = \mathrm{Tr}\left[|i\rangle\langle i|\rho'_B\right]$. After estimating the probabilities, we uncompute U to recover the initial state $|\phi\rangle$. We summarize this ability with the following theorem.

Theorem 2. *Suppose that $|\phi\rangle$ is an unknown quantum state and we are given oracle access to a coherent measurement of $P = |\phi\rangle\langle\phi|$ (that is, the oracle performs the operation $P\otimes\mathbb{I}+(1-P)\otimes\sigma_x$ on the original Hilbert space plus a single-qubit ancilla). Fix $0 < \epsilon < 1$, $\delta > 0$, and an efficiently implementable d-outcome POVM given by operators $\{E_i\}$. Then there exists an efficient quantum algorithm that takes as input a single copy of $|\phi\rangle$ and outputs an undamaged copy of $|\phi\rangle$ along with estimates q_i^{est} such that*

$$|q_i^{est} - \langle\phi|E_i|\phi\rangle| < \delta$$

for all i with probability at least $1-\epsilon$. This algorithm uses an expected number $O\left(\frac{d}{\delta}\log\left(\frac{d}{\epsilon}\right)\right)$ calls to the measurement oracle and the POVM.

The algorithm which achieves this running time is given in section III B 2.

If we want to perform tomography on a subsystem of $|\phi\rangle$, we can use theorem 2 to estimate an informationally complete POVM on that subsystem.

A. Using quantum state restoration to estimate measurement statistics

In this section we consider the running time of estimating the probabilities $q_i = \mathrm{Tr}\left[\rho_B|i\rangle_B\langle i|_B\right]$ on a given state $|\psi\rangle$ using quantum state restoration. We do this by repeatedly measuring register B and then restoring the state. Let m_i be the number of times we observe outcome i in N trials. Our estimate of q_i is

$$q_i^{\mathrm{est}} = \frac{m_i}{N}.$$

For the j^{th} observation, let $x_{i,j} \in \{0,1\}$ indicate whether the outcome of that observation was i. For fixed i, the $x_{i,j}$ are independent. To obtain a bound on the error $|q_i^{\mathrm{est}} - q_i|$, we use Hoeffding's inequality [4], which for a sequence of N independent and identically distributed random bits $x_{i,j}$ with mean value $\mathbb{E}_j[x_{i,j}] = q_i$ implies that

$$\Pr\left[\left|\frac{1}{N}\sum_{j=1}^{N} x_{i,j} - q_i\right| \geq \delta\right] \leq 2e^{-2N\delta^2}, \text{ for any } \delta > 0. \tag{4}$$

So

$$\Pr\left[|q_i^{\mathrm{est}} - q_i| \geq \delta\right] \leq 2e^{-2N\delta^2}$$

for each i individually, and, by a union bound,

$$\Pr\left[|q_i^{\mathrm{est}} - q_i| \geq \delta \text{ for any } i\right] \leq 2de^{-2N\delta^2}.$$

Choosing $N = \left\lceil \frac{1}{2\delta^2}\ln\frac{2d}{\epsilon}\right\rceil$ makes the right hand side $\leq \epsilon$. Each of the N repetitions of quantum state restoration takes an expected time $\chi\cdot d$, so the total expected number $\mathbb{E}[M_{\mathrm{SR}}]$ (where the subscript stands for "state restoration") of uses of P is

$$\mathbb{E}[M_{\mathrm{SR}}] = \chi\cdot d\left\lceil\frac{1}{2\delta^2}\ln\frac{2d}{\epsilon}\right\rceil.$$

B. Improved algorithms to estimate measurement statistics

In this section we describe two other algorithms which can be used for single copy tomography. Both of these approaches are based on Jordan's lemma [5]. The algorithms we discuss in this section are based on the QMA amplification schemes of Marriott and Watrous [6] and Nagaj et al. [8].

To use these algorithms, we fix $i \in \{1, \ldots, d\}$ and we will estimate

$$q_i = \mathrm{Tr}\left[\rho_B |i\rangle_B \langle i|_B\right].$$

We repeat this for each value of i.

We begin by defining the projector

$$Q_i = |i\rangle_B \langle i|_B$$

and the states

$$\begin{aligned} |v_i\rangle &= \frac{1}{\sqrt{q_i}} Q_i |\psi\rangle, \\ |v_i^\perp\rangle &= \frac{1}{\sqrt{1-q_i}} \left(1 - Q_i\right) |\psi\rangle. \end{aligned}$$

Note that we can write

$$|\psi\rangle = \sqrt{q_i}|v_i\rangle + \sqrt{1-q_i}|v_i^\perp\rangle. \tag{5}$$

We also define the state

$$|\psi_i^\perp\rangle = -\sqrt{1-q_i}|v_i\rangle + \sqrt{q_i}|v_i^\perp\rangle. \tag{6}$$

We can then use the above expressions to write $|v_i\rangle$ and $|v_i^\perp\rangle$ in terms of $|\psi\rangle$ and $|\psi_i^\perp\rangle$

$$\begin{aligned} |v_i\rangle &= \sqrt{q_i}|\psi\rangle - \sqrt{1-q_i}|\psi_i^\perp\rangle \\ |v_i^\perp\rangle &= \sqrt{1-q_i}|\psi\rangle + \sqrt{q_i}|\psi_i^\perp\rangle. \end{aligned} \tag{7}$$

The principal angle $\theta_i \in [0, \frac{\pi}{2}]$ between the two bases $\{|\psi\rangle, |\psi_i^\perp\rangle\}$ and $\{|v_i\rangle, |v_i^\perp\rangle\}$ is defined by

$$\cos^2\theta_i = |\langle v_i|\psi\rangle|^2 = \langle\psi|v_i\rangle\langle v_i|\psi\rangle = \langle\psi| Q_i |\psi\rangle = q_i. \tag{8}$$

Having defined the two bases $\{|v_i\rangle, |v_i^\perp\rangle\}$ and $\{|\psi\rangle, |\psi_i^\perp\rangle\}$, we are now ready to describe two algorithms for computing the expectation value q_i more efficiently than by using quantum state restoration. For any chosen ϵ and δ, each of these algorithms will generate an estimate q_i^{est} such that $|q_i^{\mathrm{est}} - q_i| < \delta$ with probability at least $1 - \frac{\epsilon}{d}$. Repeating for each i, we have $|q_i^{\mathrm{est}} - q_i| < \delta$ for all i with probability at least $1 - \epsilon$ by a union bound. The running times of these algorithms as a function of δ and ϵ are summarized in table I.

1. *Alternating Projections*

This algorithm is an application of the scheme of Marriott and Watrous [6] which was originally proposed for witness-reusing amplification of the complexity class QMA. Observe from (5), (6) and (7) that when performing the measurement P on the state $|v_i\rangle$, the probability of obtaining 1 (and the state $|\psi\rangle$) is q_i. Similarly, when measuring P on the state $|v_i^\perp\rangle$, the probability of obtaining 0 (and the state $|\psi_i^\perp\rangle$) is also q_i. We can estimate q_i by performing many alternating measurements of P and Q_i and counting the number of transitions $|v_i\rangle \leftrightarrow |\psi\rangle$ or $|v_i^\perp\rangle \leftrightarrow |\psi_i^\perp\rangle$. Let us now present the algorithm and compute its complexity measured by the expected number of measurements of P, as a function of the desired precision δ and error probability $\frac{\epsilon}{d}$.

1. Start with the state $|\psi\rangle$. Fix $N = \left\lceil \frac{1}{2} + \frac{\ln \frac{2d}{\epsilon}}{4\delta^2} \right\rceil$.
2. Repeat for $t = 1, \ldots, N$
   (a) Measure Q_i and record the measurement outcome as a bit $a_{2t-1} \in \{0, 1\}$. This produces one of the two states $|v_i\rangle$ or $|v_i^\perp\rangle$.

(b) Measure the projector $P = |\psi\rangle\langle\psi|$ and record the result $a_{2t} \in \{0, 1\}$. This produces either the state $|\psi\rangle$ or $|\psi_i^\perp\rangle$.

3. If the state is not currently $|\psi\rangle$ (because the last measurement in step 2b gave a 0), then the state is $|\psi_i^\perp\rangle$. In this case alternate measuring Q_i and P until you recover $|\psi\rangle$.

4. From the list $(a_1, \ldots, a_{2N})$, compute the list of differences $(\Delta_1, \Delta_2, ..., \Delta_{2N-1})$ where $\Delta_j = a_{j+1} \oplus a_j$. Let m denote the number of zeroes in this list of differences. Then the estimate of q is given by

$$q_i^{\text{est}} \equiv \frac{m}{2N-1}. \tag{9}$$

As discussed above, the probability of getting a measurement outcome (1 or 0) which is the same as the previous measurement outcome is q_i. So the number of zeroes which appear in the list $(\Delta_1, \Delta_2, ..., \Delta_{2N-1})$ is a binomial random variable with mean $q_i(2N-1)$. This is why (9) gives an estimator for the value of q_i.

We now show that the estimate q_i^{est} from (9) has the required precision δ, with probability at least $1-\epsilon$. To show this, we again use Hoeffding's inequality (4). Applying this to the case at hand with $q_k = 1 \oplus \Delta_k$ for $k \in \{1..., 2N-1\}$, we obtain

$$\Pr\left[|q_i^{\text{est}} - q_i| \geq \delta\right] \leq 2e^{-2(2N-1)\delta^2}.$$

The choice $N = \left\lceil \frac{1}{2} + \frac{\log \frac{2d}{\epsilon}}{4\delta^2} \right\rceil$ guarantees that the RHS is $\leq \frac{\epsilon}{d}$. Thus we have shown that the desired precision δ is achieved by our scheme with probability at least $1 - \frac{\epsilon}{d}$.

We now derive the expected number $\mathbb{E}[M_{\text{AP}}^{(i)}]$ (AP stands for alternating projections) of uses of P in the above algorithm. The random variable $M_{\text{AP}}^{(i)}$ is N plus the number of additional uses of P in step 3. The operation comprised of measuring Q_i and then measuring P is an update of a symmetric random walk on the two states $\left\{|\psi\rangle, |\psi_i^\perp\rangle\right\}$. Let $w(r)$ be the probability of transitioning from $|\psi\rangle$ to $|\psi_i^\perp\rangle$ in r steps. Then with probability $1 - w(N)$ step 3 does not use P at all and, with probability $w(N)$ it uses an expected number $\frac{1}{w(1)}$ invocations of P. Thus the expected running time of the algorithm is

$$\begin{aligned}\mathbb{E}\left[M_{\text{AP}}^{(i)}\right] &= N + w(N)\frac{1}{w(1)} \\ &\leq 2N.\end{aligned}$$

In the last line, we used the fact that $w(N)$ is less than or equal to the probability of at least one transition occuring in N steps, which is at most $Nw(1)$ by a union bound.

Hence

$$\mathbb{E}[M_{\text{AP}}^{(i)}] \leq 2\left(\left\lceil \frac{1}{2} + \frac{1}{4\delta^2}\ln\frac{2d}{\epsilon}\right\rceil\right).$$

Repeating this procedure to obtain estimates of each q_i (which are all within the desired precision δ with probability at least $1-\epsilon$) takes expected running time

$$\mathbb{E}[M_{\text{AP}}] \leq 2d\left(\left\lceil \frac{1}{2} + \frac{1}{4\delta^2}\ln\frac{2d}{\epsilon}\right\rceil\right).$$

2. An improved algorithm using phase estimation

In this section we will give an improved algorithm for single copy tomography using phase estimation, based on a fast QMA amplification scheme given in [8]. Its advantage over the previous two algorithms is that it requires quadratically fewer measurements of P . The results of this section will prove Theorem 2.

As in the previous section, we estimate the q_i one at a time for $i \in \{1, \ldots, d\}$.

We begin by defining the unitary operator

$$W_i = (2P - \mathbb{I})(2Q_i - \mathbb{I}),$$

which is a product of two reflections. Note that if we can implement P so that it coherently xors its measurement outcome into an ancilla register (as in the assumption of theorem 2), then we can implement the operator $(2P - \mathbb{I})$ by first initializing that ancilla to $|-\rangle$ and applying the measurement.

Within the 2D subspace S_i spanned by the vectors $|\psi\rangle$ and $|\psi_i^\perp\rangle$ (7), the operator W_i is a rotation

$$W_i\big|_{S_i} \;=\; e^{-2i\theta_i\sigma_y}, \tag{10}$$

where θ_i is the principal angle as defined in (8) (and σ_y refers to the Pauli matrix).

We now describe how to obtain $q_i = \langle\psi| Q_i |\psi\rangle = \cos^2\theta_i$ by running phase estimation of the operator W_i on the state $|\psi\rangle$. The eigenvectors of W_i are

$$|\phi_i^\pm\rangle = \frac{1}{\sqrt{2}}\left(|\psi\rangle \pm i|\psi_i^\perp\rangle\right). \tag{11}$$

and correspond to eigenvalues $e^{\mp i2\pi\phi_i}$, where $\phi_i = \frac{\theta_i}{\pi}$ so that $0 < \phi_i < \frac{1}{2}$. After running phase estimation of W_i on the input state $|\psi\rangle$, we will likely measure a good approximation to either ϕ_i or $1 - \phi_i$. Note that either outcome provides a good estimate of

$$q_i = \cos^2(\pi\phi_i) = \cos^2(\pi(1-\phi_i)).$$

This is the idea of the algorithm we present in this section. Our algorithm must have a failure probability lower than that obtained by a single use of phase estimation, and we must recover the state $|\psi\rangle$ at the end of the algorithm.

Our algorithm begins by defining

$$\begin{aligned} t &= \left\lceil \log_2\left(\frac{3\pi}{\delta}\right)\right\rceil + 2. \\ r &= \left\lceil \frac{1}{\log_2\left(\frac{2}{\sqrt{3}}\right)} \log_2\left(\frac{d}{2\epsilon}\right)\right\rceil. \end{aligned} \tag{12}$$

We proceed as follows:

1. Start in the state $|\psi\rangle|0\rangle^{\otimes t}$.
2. Repeat for $j = 1, ..., r$
   (a) Reset the t qubits of the second register to the state $|0\rangle^{\otimes t}$. Perform phase estimation of the operator W_i on the state of the first register, computing the phase using the t ancillae in the second register. Define
   $$q_i^{(j)} = \cos^2\left(\pi\phi_i^{(j)}\right)$$
   where $\phi_i^{(j)}$ is the measured phase.
   (b) Measure the projector $P = |\psi\rangle\langle\psi|$ on the first register.
3. If the state is not currently $|\psi\rangle$ (because the last measurement in step 2b gave a 0), then the state is $|\psi_i^\perp\rangle$. In this case repeat phase estimation followed by measurement of P until you measure a 1 for P, recovering the state $|\psi\rangle$.
4. Let q_i^{est} be the median of the values $\{q_i^{(j)}\}$ for $j \in \{1, ...r\}$.

We now determine the expected runtime of this algorithm, and then we will show that the resulting estimate q_i^{est} achieves the desired precision with high enough probability. Our analysis of the runtime is based on the observation that each iteration of phase estimation followed by measurement of P is an update of a random walk on the two states $\{|\psi\rangle, |\psi_i^\perp\rangle\}$. If we start in state $|\psi\rangle$ of the first register then after applying phase estimation (but before measuring the phase) we obtain a state

$$|\Psi_i\rangle = \frac{1}{\sqrt{2}}\left(|\phi_i^+\rangle|\gamma\rangle + |\phi_i^-\rangle|\mu\rangle\right).$$

where $|\gamma\rangle$ and $|\mu\rangle$ are t-qubit states. So the probability of measuring 1 in step 2b is

$$\Pr\left[|\psi\rangle \to |\psi\rangle\right] = \mathrm{Tr}\left[\left(|\psi\rangle\langle\psi| \otimes \mathbb{I}\right)|\Psi_i\rangle\langle\Psi_i|\right]$$

in which case the resulting state of the first register is $|\psi\rangle$. The probability of measuring a zero in this step is

$$\Pr\left[|\psi\rangle \to \left|\psi^{\perp}\right\rangle\right] = \mathrm{Tr}\left[\left(|\psi^{\perp}\rangle\langle\psi^{\perp}| \otimes \mathbb{I}\right)|\Psi_i\rangle\langle\Psi_i|\right] = 1 - \Pr\left[|\psi\rangle \to |\psi\rangle\right]$$

in which case the resulting state of the first register is $|\psi^{\perp}\rangle$. Similarly, one can compute the transition probabilities starting from the state $|\psi^{\perp}\rangle$ of the first register. These satisfy

$$\begin{aligned}\Pr\left[\left|\psi^{\perp}\right\rangle \to \left|\psi^{\perp}\right\rangle\right] &= \Pr\left[|\psi\rangle \to |\psi\rangle\right] \\ \Pr\left[\left|\psi^{\perp}\right\rangle \to |\psi\rangle\right] &= \Pr\left[|\psi\rangle \to \left|\psi^{\perp}\right\rangle\right]\end{aligned}$$

so the random walk is symmetric. We can then directly apply our analysis of the previous section to show that

$$\mathbb{E}[\#\text{ of uses of phase estimation followed by measurement of } P] \leq 2r.$$

Each time we use phase estimation with t ancillae, we use the gate W_i less than 2^t times [9]. So each time we repeat phase estimation followed by measurement of P we use less than $2^t + 1$ measurements of P so the expected total number of times $\mathbb{E}[M_{\mathrm{PE}}^{(i)}]$ (PE stands for phase estimation) that we use the measurement of P is

$$\begin{aligned}\mathbb{E}[M_{\mathrm{PE}}^{(i)}] &< 2r \cdot \left(2^t + 1\right) \\ &\leq 2r\left(\frac{12\pi}{\delta} + 1\right) \\ &= 2\left\lceil \frac{1}{\log_2\left(\frac{2}{\sqrt{3}}\right)} \log_2\left(\frac{d}{2\epsilon}\right)\right\rceil \left(\frac{12\pi}{\delta} + 1\right).\end{aligned}$$

Repeating this procedure to obtain estimates of each q_i takes expected running time

$$\mathbb{E}[M_{\mathrm{PE}}] < 2d\left\lceil \frac{\log_2\left(\frac{d}{2\epsilon}\right)}{\log_2\left(\frac{2}{\sqrt{3}}\right)}\right\rceil \left(\frac{12\pi}{\delta} + 1\right). \tag{13}$$

We now show that the probability that all the estimates q_i^{est} obtained by using the above algorithm satisfy

$$|q_i^{est} - q_i| < \delta$$

is at least $1 - \epsilon$. Our choice of t was designed so that the output of phase estimation of W_i on the state $|\phi_i^+\rangle$ using t ancillae is a state $|\phi_i^+\rangle|\gamma\rangle$ such that a measurement of the t-qubit state $|\gamma\rangle$ in the computational basis produces a phase $\tilde{\phi}$ that satisfies

$$|\tilde{\phi} - \phi_i| \leq \frac{\delta}{3\pi}$$

with probability at least $\frac{3}{4}$[9]. Similarly the output of phase estimation of W_i on the state $|\phi_i^-\rangle$ using t ancillae is a state $|\phi_i^-\rangle|\mu\rangle$ such that a measurement of the t-qubit state $|\mu\rangle$ in the computational basis produces a phase $\tilde{\phi}$ that satisfies

$$|\tilde{\phi} - (1 - \phi_i)| \leq \frac{\delta}{3\pi}$$

with probability at least $\frac{3}{4}$. In step 2a of our algorithm we perform phase estimation on either the state $|\psi\rangle$ or the state $|\psi_i^{\perp}\rangle$. In either case, the reduced density matrix of the t-qubit ancilla register after applying the phase estimation (but before measuring the phase) is

$$\frac{1}{2}\left(|\gamma\rangle\langle\gamma| + |\mu\rangle\langle\mu|\right)$$

State Restoration	Alternating Projectors	Phase Estimation
$\mathbb{E}[M_{\mathrm{SR}}] = \mathrm{O}\left(\frac{\chi \cdot d}{\delta^2} \log \frac{d}{\epsilon}\right)$	$\mathbb{E}[M_{\mathrm{AP}}] = O\left(\frac{d}{\delta^2} \log \frac{d}{\epsilon}\right)$	$\mathbb{E}[M_{\mathrm{PE}}] = O\left(\frac{d}{\delta} \log\left(\frac{d}{\epsilon}\right)\right)$

Table I: Scaling of the expected number of measurements of $P = |\psi\rangle\langle\psi|$ used by each algorithm as a function of the desired precision δ and error probability ϵ.

which is an equal probabilistic mixture of $|\gamma\rangle$ and $|\mu\rangle$. So, with probability at least $\frac{3}{4}$ (regardless of whether we started in $|\psi\rangle$ or $|\psi^{\perp}\rangle$), the phases $\phi_i^{(j)}$ measured in step 2 of the algorithm satisfy either

$$|\phi_i^{(j)} - \phi_i| \leq \frac{\delta}{3\pi}$$

or

$$|\phi_i^{(j)} - (1 - \phi_i)| \leq \frac{\delta}{3\pi}.$$

Using the inequality

$$|\cos^2(\pi\alpha) - \cos^2(\pi\beta)| \leq 2\pi|\alpha - \beta|$$

and the fact that $\cos^2(\pi x) = \cos^2(\pi(1-x))$ it follows that the estimates $q_i^{(j)}$ each (independently) satisfy

$$|q_i^{(j)} - q_i| < \delta$$

with probability at least $\frac{3}{4}$. The median lemma of [8] says in this case that the probability that the median of the r independent measured values $q_i^{(j)}$ falls outside the interval $(q_i - \delta, q_i + \delta)$ is upper bounded as $p_{\mathrm{fail}} \leq \frac{1}{2}\left(\frac{\sqrt{3}}{2}\right)^r$. Plugging in our choice of r from equation 12 gives

$$|q_i^{\mathrm{est}} - q_i| < \delta$$

for each i with probability at least $1 - \frac{\epsilon}{d}$. So the probability that the above inequality is satisfied for all of the $i \in \{1, ...d\}$ is at least $1 - \epsilon$.

C. Performance comparison for estimating measurement statistics

These three algorithms for estimating the probabilities $q_i = \mathrm{Tr}\left[\rho_B |i\rangle_B\langle i|_B\right]$ give estimates $\{q_i^{\mathrm{est}}\}$ (for i from 1 to d) which are all within δ of the correct values with probability at least $1 - \epsilon$. Their running times are summarized in table I.

State restoration is conceptually the simplest of the three algorithms, and we expect that it will be sufficient for most purposes. It is also the slowest as a function of d and δ (assuming χ is increasing as a function of d). The state restoration algorithm has the advantage that we can drop in different tomography schemes that may improve performance.

In the absence of a better tomography scheme, however, both other algorithms outperform state restoration as a function of d. Phase estimation also performs quadratically better than both other algorithms as $\delta \to 0$.

IV. APPLICATIONS OF QUANTUM STATE RESTORATION AND SINGLE-COPY TOMOGRAPHY

In this section we describe applications of quantum state restoration and single-copy tomography.

A. Breaking quantum money

As we discussed in the introduction, quantum money is the idea of using a state as money—that is, something that can be passed around but not forged. The money consists of a quantum state and a verification procedure which should succeed with high probability on valid money issued by the bank but should fail with high probability for any

efficiently forgeable state. The first quantum money protocols [3, 11] required the verification procedure to be secret, so only the bank (i.e. the issuer of the money) could verify money states. There is recent interest in publicly verifiable quantum money [1, 2, 7], in which everyone, including a would-be forger, has access to the verification procedure. In the introduction, we showed that quantum state restoration breaks quantum money based on product states. More generally, as a corollary of Theorem 2, any quantum money protocol in which the verifier is a projector must be designed to withstand attacks based on single-copy tomography. If the verifier is a projector, then an adversary can use single-copy tomography to learn the measurement statistics of any efficiently implementable measurement with a small number of outcomes on the quantum money state $|\psi\rangle$. The money must be designed so that these statistics are not useful to a counterfeiter.

B. Studying ground states of many-body Hamiltonians

Quantum computers offer potentially exponential speedups in simulating quantum mechanics, but some problems are still hard. For example, preparing ground states of many-body systems generically takes exponential time in the number of particles. Nonetheless, for sufficiently small systems with large enough energy gaps, algorithms such as [10] may run quickly enough to prepare a single copy of the ground state, and phase estimation can be used to verify the ground state. Single-copy tomograpy allows us to make multiple tomographic measurements (even of non-commuting operators) on small numbers of particles without having to prepare multiple copies of the ground state. This gives a large speedup over traditional tomography.

Single-copy tomography could also be useful to characterize the ground state during adiabatic evolution. This information could even be used in real time to guide the choice of path for an adiabatic algorithm.

V. CONCLUSIONS

It is strongly believed that the ability to verify an unknown state $|\psi\rangle$ does not give the ability to produce that state efficiently. Without the ability to verify a state, mere posession of that state confers little power. As we have shown, the combination of a verifier and single copy of $|\psi\rangle$ is more powerful that either one alone. We hope that the methods that we present in this paper can be used as tools to develop new algorithms beyond those that we have suggested.

VI. ACKNOWLEDGEMENTS

This work was supported in part by funds provided by the U.S. Department of Energy under cooperative research agreement DE-FG02-94ER40818, the W. M. Keck Foundation Center for Extreme Quantum Information Theory, the U.S. Army Research Laboratory's Army Research Office through grant number W911NF-09-1-0438, the National Science Foundation through grant number CCF-0829421, the NDSEG fellowship, the Natural Sciences and Engineering Research Council of Canada, and Microsoft Research.

[1] S. Aaronson. Quantum copy-protection and quantum money. In *Computational Complexity, Annual IEEE Conference on*, pages 229–242, 2009.

[2] Scott Aaronson, Edward Farhi, David Gosset, Avinatan Hassidim, Jon Kelner, Andrew Lutomirski, and Peter Shor. Breaking and making quantum money: Toward a new quantum cryptographic protocol. *Innovations in Computer Science ICS2010*.

[3] C.H. Bennett, G. Brassard, S. Breidbart, and S. Wiesner. Quantum cryptography, or unforgeable subway tokens. In *Advances in Cryptology–Proceedings of Crypto*, volume 82, pages 267–275, 1983.

[4] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), 1963.

[5] C. Jordan. *Bulletin de la S. M. F.*, 3:103, 1875.

[6] Chris Marriott and John Watrous. Quantum arthur-merlin games. *Computational Complexity*, 14(2):122–152, 2005.

[7] Michele Mosca and Douglas Stebila. Quantum coins, 2009.

[8] Daniel Nagaj, Pawel Wocjan, and Yong Zhang. Fast amplification of qma. *Quantum Information & Computation*, 9(11&12):1053–1068, 2009.

[9] Michael A. Nielsen and Isaac L. Chuang. *Quantum Information and Computation*. Cambridge University Press, Cambridge, UK, 2000.

[10] David Poulin and Pawel Wocjan. Preparing ground states of quantum many-body systems on a quantum computer. *Physical Review Letters*, 102(13):130503, 2009.
[11] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983.
[12] Verification is not the same as measuring the energy. One way to verify the state is to apply phase estimation, compute an indicator of whether the energy has the right value, uncompute the phase estimation step, and measure the indicator.
[13] The bank must also digitally sign the description of the projectors using a classical digital signature protocol which is a secure against quantum adversaries. Such protocols are believed to exist.