Practical Characterization of Quantum Devices without Tomography

Marcus P. da Silva,[1,2] Olivier Landon-Cardinal,[2] and David Poulin[2]

[1]*Disruptive Information Processing Technologies Group, Raytheon BBN Technologies, Cambridge, Massachusetts 02138, USA*
[2]*Département de Physique, Université de Sherbrooke, Sherbrooke, Québec, J1K 2R1, Canada*

(Received 11 July 2011; published 16 November 2011)

Quantum tomography is the main method used to assess the quality of quantum information processing devices. However, the amount of resources needed for quantum tomography is exponential in the device size. Part of the problem is that tomography generates much more information than is usually sought. Taking a more targeted approach, we develop schemes that enable (i) estimating the fidelity of an experiment to a theoretical ideal description, (ii) learning which description within a reduced subset best matches the experimental data. Both these approaches yield a significant reduction in resources compared to tomography. In particular, we demonstrate that fidelity can be estimated from a number of simple experiments that is independent of the system size, removing an important roadblock for the experimental study of larger quantum information processing units.

DOI: 10.1103/PhysRevLett.107.210404 PACS numbers: 03.65.Wj, 03.67.Ac

The building blocks for quantum computers have been demonstrated in a number of different physical systems [1–6]. In order to quantify how closely these demonstrations come to the ideal operations, the experiments are fully characterized via either quantum state tomography [7] or quantum process tomography [8]. An important advantage of these methods is that they require only simple local measurements. The main drawbacks however are that tomography fundamentally requires both experimental and data postprocessing resources that increase exponentially with the number of particles n [9].

It is important to realize that the exponential cost of tomography is not a problem restricted to a large number of qubits. For example, recent ion trap experiments characterizing an 8 qubit state required 10 h of measurements, despite collecting only 100 samples per observable [3]. Surprisingly, the postprocessing of the data obtained from these experiments took approximately a week [10]. Under similar time scales, the characterization of a 16 qubit state would take years of measurements, and over a century of data postprocessing. This is clearly a major obstacle in the demonstration of working quantum computers, even at sizes moderately larger than what has been demonstrated to date.

Moreover, one of the key assumptions for the fault-tolerance theorems of quantum computation is that the noise on elementary components does not scale badly with the system size [11]. Therefore, despite the fact that universal quantum computation can be realized with one- and two-qubit elementary operations, it is not sufficient to characterize small gates—larger systems may have significant noise contributions from correlated sources as seen in recent experiments [6]. The characterization of multiqubit states and operations provides crucial information for the verification of these assumptions, and therefore the development of large quantum information processors.

Part of the problem with the usual approach is that tomography often provides more information than what is truly sought. Given an experiment that prepares a quantum state represented by a density operator $\hat{\sigma}$, one usually extracts a complete description for $\hat{\sigma}$ via quantum tomography, and then compares this description to a theoretical state $\hat{\rho}$ by computing the fidelity $F(\hat{\rho}, \hat{\sigma})$—a single number, commonly used as similarity measure. As this example illustrates, we often have an idea of what has been realized in the laboratory, so we are interested in asking for much less information—e.g., we only want to know the distance to some particular theoretical target or to learn the identity of the state or operation within a restricted set of possibilities.

In this Letter, we develop targeted approaches to directly extract the information of interest. Our main results, summarized at Table I, show that it is possible to efficiently characterize a large class of states and operations—including some that are universal resources for quantum computation—without resorting to tomography and using only local measurements and the preparation of product states. Our methods apply to discrete variable systems such as qubits, as well as continuous variable systems such as oscillators. We consider two types of characterization: *certification* and *learning*.

Learning consists of identifying the theoretical description from a restricted set of possibilities that best matches the experimental data. There exists many classes of ''variational'' states in physics that can be specified with a small number of parameters. We provide examples where these parameters can be extracted directly from experiments, circumventing tomography and hence drastically reducing the complexity.

Certification consists of estimating the fidelity between an experimental device and some theoretical target. We demonstrate that certification always requires

TABLE I. Complexity of the characterization of various states and processes. Entries with asterisks are efficient, i.e., require resources that grow at most polynomially with the number of qubits n. The sampling column gives the complexity of the classical processing required to sample from the relevance distribution (C1). The fluctuations column gives the number of measurements required to suppress statistical fluctuations when evaluating the fidelity (C2). The learning column gives the total number of measurements (including repetitions of the same measurement setting) required to learn the state within a restricted set; the classical processing is always a polynomial of that number. When both fidelity estimate and learning are efficient, it is not necessary to assume that the state belongs to a restricted set as fidelity testifies of that assumption. Stabilizer states, Clifford gates, local Hamiltonians and Lindbladians are discussed in the main text. The W state has often been used as an experimental benchmark, e.g., [3]. The $|t_n\rangle$ state plays a key role in linear optics quantum computation [12]. Matrix product states (MPS) accurately describe ground states of 1D quantum systems [13]. An important example of a process with MPS Choi matrix is the approximate quantum Fourier transform [14], key component of Shor's factoring algorithm. Question marks indicate open problems, but they can be no worst than the general states and operations. All schemes require only single qubit measurements, except for stabilizer state and Clifford process learning [28,29].

		Certification		
		Sampling (C1)	Fluctuations (C2)	Learning
States	Stabilizer	$O(n)^*$	$O(1)^*$	poly(n)
	W	$O(n)^*$	$O(n)^*$	$O(n)^*$
	$\|t_n\rangle$	$O(n)^*$	$O(n)^*$	$O(n)^*$
	General MPS	$O(n)^*$	?	$O(n)^*$ [15]
	General pure state	$O(n^2 2^{2n})$	$O(2^n)$	$O(2^{6n})$
Processes	Clifford	$O(1)^*$	$O(1)^*$	poly(n)
	MPS Choi matrix	$O(n)^*$	?	$O(n)^*$
	General unitary	$O(n^2 2^{4n})$	$O(2^{2n})$	$O(2^{12n})$
Evolution	Local Hamiltonian	NA	NA	$O(n)^*$
	Local Lindbladian	NA	NA	$O(n)^*$

drastically less resources than full tomography—in some important cases, it is an exponential reduction in resources. Even in the worst case, our scheme offers four significant advantages for the characterization of quantum states (equivalent statements hold for quantum operations): (i) Its computational cost is bounded by $n^2 4^n$, compared to 4^{3n} required for the simplest tomography procedure based on pseudoinverses. (ii) The number of distinct experimental settings it requires is constant—independent of the system size and depending only on the desired accuracy of the estimate—compared to the 4^n distinct experiments needed by tomography, or the $O(n2^n)$ settings required by compressed sensing techniques [16]. (iii) The total number of measurements (counting repeated measurements used to statistically estimate expectation values) of our scheme is bounded by $O(2^n)$, which is at least a quadratic improvement over what is required by full tomography. (iv) The data postprocessing of our scheme is trivial, while the correct method of processing tomography data is a matter of current debates and different methods produce significantly different results [10].

The rest of this Letter is structured as follows. In the next three sections, we describe the state certification scheme for qubits, show how it extends to continuous variable systems, and the certification of quantum processes. Then, we present concrete examples drawn from Table I.

Monte Carlo state certification.—To estimate the fidelity to some theoretical pure state $\hat{\rho}$, we use the fidelity

$$F(\hat{\rho}, \hat{\sigma}) = \mathrm{tr}\hat{\rho}\,\hat{\sigma} = \sum_i \frac{\rho_i \sigma_i}{d} = \sum_i \frac{\rho_i^2}{d}\frac{\sigma_i}{\rho_i}, \qquad (1)$$

where $\rho_i = \mathrm{tr}\hat{\rho}\hat{P}_i$, $\sigma_i = \mathrm{tr}\hat{\sigma}\hat{P}_i$, d is the dimension of the Hilbert space, and $\hat{P}_i$ is some orthonormal Hermitian operator basis satisfying $\mathrm{tr}\hat{P}_i\hat{P}_j = d\delta_{ij}$. For a system composed of n qubits, the $\hat{P}_i$ could be the 4^n Pauli operators obtained by taking tensor products of the Pauli matrices and the identity. Defining the *relevance distribution* $\Pr(i) = \frac{\rho_i^2}{d}$, we can rewrite the fidelity as $F(\hat{\sigma}, \hat{\rho}) = \sum_i \Pr(i)\frac{\sigma_i}{\rho_i}$, where the sum is taken over only the i with $\rho_i \neq 0$. This expression leads to an experimental procedure to estimate the fidelity based on Monte Carlo methods as follows: one generates N random indices $i_1, i_2, \ldots, i_N$ following the relevance distribution $\Pr(i)$ and estimates $\sigma_{i_k} = \langle \hat{P}_{i_k} \rangle_{\hat{\sigma}}$, the experimental expectation value of the observable $\hat{P}_{i_k}$. With high probability, the fidelity is close to $\frac{1}{N}\sum_{k=1}^{N}\frac{\sigma_{i_k}}{\rho_{i_k}}$ with an uncertainty that decreases as $\frac{1}{\sqrt{N}}$. The total number of distinct experimental settings is at most N, independent of the system size.

There are two important caveats to this technique: (C1) Generating an index i according to the relevance distribution $\Pr(i)$ can in general require an exponential amount of computational resources. (C2) Each σ_{i_k} is estimated within some finite accuracy. To estimate the fidelity with accuracy ϵ therefore requires repeating the measurement of P_{i_k} roughly $(\epsilon\rho_{i_k})^{-2}$ times, which in the worst case grows exponentially with the number of qubits.

These are important limitations, and as a consequence our method will not scale polynomially for all quantum states and operations, but nevertheless always does significantly better than tomography. In addition, there are important classes of states and operations which avoid these two problems (see Table I and the Supplemental Material [17] for complete details).

Continuous variables systems—For infinite dimensional systems, such as a harmonic oscillator or a single optical mode in a cavity, it is more convenient to describe a state $\hat{\rho}$ by its Wigner functions $W_{\hat{\rho}}(\alpha)$ [18] (other indicator functions could also be used). Equation (1) becomes

$$F(\hat{\rho}, \hat{\sigma}) = \frac{1}{\pi} \int_{\mathbb{C}} d^2\alpha p(\alpha) \frac{W_{\hat{\sigma}}(\alpha)}{W_{\hat{\rho}}(\alpha)}, \tag{2}$$

where the relevance density $p(\alpha) = W^2_{\hat{\rho}}(\alpha)$ is defined as the square of the Wigner function of the theoretical state, whose purity guarantees once again that $p(\alpha)$ is well defined as a probability density. The Wigner function of the experimental state $\hat{\sigma}$ can be measured by interactions with an atom and measurements of the atom's state [19]. Points in the complex plane can be selected according to $p(\alpha)$ using simple methods such as rejection sampling. As an example, we simulated this proposed method to estimate the fidelity between a quantum superposition of two harmonic oscillator states—a "cat" state $\frac{1}{\sqrt{2}}(|\alpha\rangle + |-\alpha\rangle)$—and the probabilistic mixture of those two classical states. For the given choice of parameters, this fidelity is $\frac{1}{2}(1 + e^{-2\alpha^2}) \approx 0.5$, and Fig. 1 clearly demonstrates a close agreement between the Monte Carlo estimate and the exact theoretical value, as the absolute error decreases like the square-root of the number of samples of the Wigner function. As expected, the error in the fidelity estimate does not depend on the state itself (e.g. average number of photons, amplitude, etc.) but only on the number of samples. We emphasize once again that no estimate of the Wigner function of the experimental state is ever made, so there is no need for maximum-likelihood fits to the data, or Radon transforms.

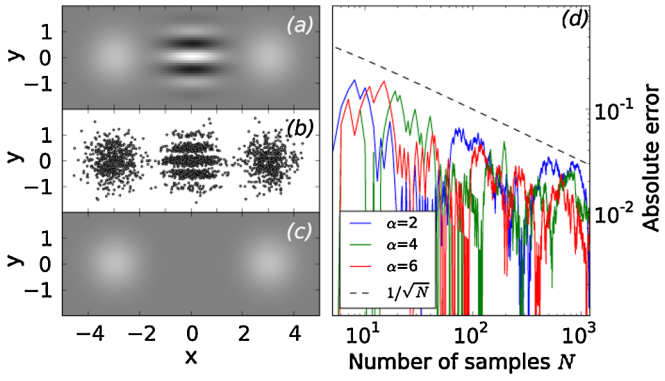

FIG. 1 (color online). (a) Wigner function representation of a harmonic oscillator in the superposition $|\psi\rangle = |\alpha\rangle + |-\alpha\rangle$ for $\alpha = 3$, (b) 10^3 samples of points in the complex plane drawn according to the relevance density of $|\psi\rangle$, (c) Wigner function representation of a harmonic oscillator in the incoherent mixture of $|\alpha\rangle$ and $|-\alpha\rangle$, corresponding to the preparation of $\hat{\sigma}$, (d) absolute error in successive estimates of the fidelity $F(\hat{\rho}, \hat{\sigma})$ for 5 different runs with 10^3 samples each.

Monte Carlo process certification.—The Choi-Jamiołkowski isomorphism [20] associates to every quantum operation $\mathcal{E}$ on a d-dimensional space a density operator $\hat{\rho}_{\mathcal{E}}$ on a d^2-dimensional space via $\hat{\rho}_{\mathcal{E}} = (\mathrm{id} \otimes \mathcal{E})(|\phi\rangle\langle\phi|)$ where $|\phi\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^{d} |i\rangle \otimes |i\rangle$ and id is the identity operation. As with state certification, our goal is to compare a target unitary $\mathcal{U}$ to its experimental realization $\widetilde{\mathcal{U}}$. A good figure of merit in that case is the average output fidelity $\bar{F}(\mathcal{U}, \widetilde{\mathcal{U}})$, defined as the fidelity between the output states produced by $\mathcal{U}$ and $\widetilde{\mathcal{U}}$, averaged uniformly over all pure input states. It can be shown that $\bar{F}(\mathcal{U}, \widetilde{\mathcal{U}}) = \frac{dF(\hat{\rho}_{\mathcal{U}}, \hat{\rho}_{\widetilde{\mathcal{U}}}) + 1}{d+1}$ [21], reducing the problem of comparing two processes $\mathcal{U}$ and $\widetilde{\mathcal{U}}$ to the problem of comparing two states $\hat{\rho}_{\mathcal{U}}$ and $\hat{\rho}_{\widetilde{\mathcal{U}}}$. This problem is solved by the Monte Carlo state certification presented above.

While this derivation makes use of the maximally entangled state $|\phi\rangle$, the experimental realization of the protocol requires only the preparation of product states. A direct implementation of the quantum Monte Carlo state certification would prepare a maximally entangled state $|\phi\rangle$, apply $\widetilde{\mathcal{U}}$ to half of the system, and then measure random Pauli operators on all qubits. A more practical approach consists of preparing the complex conjugate of random product of eigenstates of local Pauli operators (corresponding to the resulting state after half of the entangled state is measured destructively), applying the transformation $\widetilde{\mathcal{U}}$ to the system, and finally measuring a random Pauli operator on each qubit. This simplification, based on the identity $(|\mu\rangle\langle\mu| \otimes \mathrm{id})|\phi\rangle = |\mu\rangle \otimes |\mu\rangle^*$, generates the same statistics as the direct scheme [22].

Computation via teleportation.—Some of the most promising approaches to universal and scalable quantum computation are teleportation-based quantum computation [23] and measurement-based quantum computation [24]. Both these approaches rely heavily on the preparation of stabilizer states [25] and the application of quantum operations known as the Clifford group [23], which map stabilizer states to stabilizer states. Stabilizer states are also important for quantum computation in general because of their close relationship to a large class of quantum error correction codes known as *stabilizer codes*. Many of the experimental demonstrations of state preparation to date have been of stabilizer states, such as states encoded into stabilizer codes [2], cluster states [4], and the GHZ state $|00\cdots0\rangle + |11\cdots1\rangle$ [5,6].

We first describe how to *certify* these states and operations. Stabilizer states are defined to be $+1$ eigenstates of

some set of commuting Pauli operators $\hat{S}_j$ that generate the stabilizer group, i.e., $\hat{S}_j|\psi\rangle = |\psi\rangle$ for all $j = 1, \ldots, n$. It follows that $\Pr(i) = 1/d$ if either of $\pm\hat{P}_i$ is in the stabilizer group and 0 otherwise. Sampling from $\Pr(i)$ thus amounts to generating an index i uniformly between 1 and d, avoiding the problem associated with caveat (C1). For the same reasons, $\rho_i^2 = 1$ for all i with $\Pr(i) \neq 0$, so that the uncertainty in the estimation of σ_i is not amplified, avoiding the problem associated with caveat (C2). It also follows that the fidelity $F(\hat{\sigma}, \hat{\rho})$ to a stabilizer state $\hat{\rho}$ can be estimated with error ϵ using $N = O(\frac{1}{\epsilon^4})$ experiments involving only local projective measurements, independently of the system size and without any prior knowledge of the experimental state $\hat{\sigma}$. Since this result relies only on local measurements, it can immediately be generalized to states which are locally equivalent to stabilizer states.

This result carries over directly to the certification of Clifford operation because their Choi-Jamiołkowski density operators are stabilizer states. In the case of Clifford transformations similar results can be obtained using ''twirling'' experiments [26] or by the selective measurement of matrix elements of the Choi matrix [22], although the Monte Carlo approach described here generalizes to other cases.

While operations in the Clifford group are not sufficient to perform universal computation [23], single-qubit rotations can be used to reach universality, and these can be certified efficiently thanks to local equivalence of either operations (if the rotation is applied directly) or state preparation (if the rotation is applied via ''magic state'' teleportation [23,27]).

Stabilizer states can also be *learned* efficiently, as pointed out by Aaronson and Gottesman [28], although the known method for efficient stabilizer learning requires collective measurements. Aside from the direct generalization of the stabilizer approach, Clifford group operations can be learned efficiently [29] if one has access to Bell measurements and the inverse of the operation being learned—or, equivalently, applying the algorithm by Aaronson and Gottesman to appropriately chosen output states from the unitary in question. The problem of performing these tasks efficiently with strictly local measurements and without the need for the inverse remains open.

Local Hamiltonians and Lindbladians.—Models of universal quantum computation exist where the idea of discrete gates is not a natural fit. Instead, the system evolves in a continuous way, governed by some dynamical equation $\frac{\partial}{\partial t}\hat{\rho} = \mathcal{G}\hat{\rho}$. The most direct way to determine how accurately these dynamics can be realize is to estimate the time evolution generator $\mathcal{G}$ of the system, and explicitly check how it compares against the ideal target generator. Important examples include local Hamiltonians and Lindbladians that are universal for adiabatic quantum computation [30] and dissipation-driven quantum computation [31], respectively.

In what follows we demonstrate how to learn such local $\mathcal{G}$ using only (i) the preparation of initial product states, (ii) the simultaneous measurement of a constant number of single-qubit operator, (iii) a number of experimental settings that grows linearly with the system size, (iv) and classical postprocessing of complexity n^3 (inverting an $cn \times cn$ matrix for some constant c); improving on [32].

Consider the case of coherent evolution generated by some Hamiltonian H. For a short time t, the expectation value of any observable $\hat{A}$ evolves as

$$\langle\hat{A}(t)\rangle_{\hat{\rho}} - \mathrm{tr}\hat{A}\,\hat{\rho} = it\langle[\hat{H}, \hat{A}]\rangle_{\hat{\rho}} + O(\|\hat{H}\|^2 t^2). \quad (3)$$

By experimentally measuring this expectation value, we obtain one linear constraint on the Hamiltonian. Varying over different observables $\hat{A}_i$ and initial states $\hat{\rho}_j$, we obtain more linear constraints that we can write as $W_{ij} = \langle\hat{A}_i(t)\rangle_{\hat{\rho}_j} - \mathrm{tr}\hat{A}_i\hat{\rho}_j = it\langle[\hat{H}, \hat{A}_i]\rangle_{\hat{\rho}_j}$ where we have dropped the higher order terms $O(\|\hat{H}\|^2 t^2)$. Writing $\hat{H}$ in an operator basis $\hat{H} = \sum_l h_l \hat{P}_l$, we obtain the linear equation $W_{ij} = \sum_l T_{ij,l} h_l$ where $T_{i,j,l} = it\mathrm{tr}\hat{\rho}_j[\hat{P}_l, \hat{A}_i]$. The Hamiltonian can be learned by inverting this linear equation [32].

There are in general a number important caveats to this approach, although all of these disappear when the Hamiltonian is *local*, which is nonetheless sufficient to achieve universal quantum computation [30,31]. The Lieb-Robinson bound [33] shows that only the Hamiltonian $\hat{H}_R$ in a region R a distance $d \approx vt$ of the local observable $\hat{A}$ contributes to its evolution, i.e., $e^{i\hat{H}t}\hat{A}e^{-i\hat{H}t} \approx e^{i\hat{H}_R t}\hat{A}e^{-i\hat{H}_R t}$ (for details of the proof see the Supplemental Material [17]). This fact solves all the problems associated to the proposal of [32]: (i) The error $O(\|H\|^2 t^2)$ appearing in Eq. (3) becomes $O(\|H_R\|^2 t^2) = O(\|\hat{A}\|^2 t^4)$, independent of the system size. Thus, it is not necessary to decrease the evolution time t as the system size increases to achieve a given accuracy. (ii) Because the Hamiltonian is local, the number of nonzero terms h_l is proportional to the number of particles in any finite dimension. Thus, in the linear equation for W_{ij}, the range of the index l increases only linearly with the number of particles, as opposed to the exponential growth for generic Hamiltonians. (iii) Because the dynamics is local, $T_{ij,l} = T_{ij',l}$ when $\hat{\rho}_j$ and $\hat{\rho}_{j'}$ differ only outside a region of radius k away from the local observable $\hat{A}_i$. In addition, the T become linearly dependent—and thus redundant—when the input states are linearly dependent. For each observable $\hat{A}_i$, we only need to vary the initial state locally, so the total number of observable-state pairs (ij) grows linearly with the number of particles. Thus, learning the Hamiltonian—or equivalently the h_l—amounts to inverting the linear-size linear equation $W_{ij} = \sum_l T_{ij,l} h_l$. (iv) Product input states form a complete operator basis, so they are sufficient to gain all information about the Hamiltonian. Thus $\mathrm{tr}\hat{A}_i\hat{\rho}_j$ can be easily computed since

$\hat{A}_i$ is local and $\hat{\rho}_j$ is a product state. The quantity $\mathrm{tr}\hat{\rho}_j[\hat{P}_l, \hat{A}_i]$ can also be evaluated efficiently because the commutator of two k-local operators is at most $2k$ local, and $\hat{\rho}_j$ is a product state.

We thank P. Sémon for many instructive discussions about Monte Carlo methods. This work is partially funded by FQRNT, NSERC, and numerical calculations were performed using resources from RQCHP.

After this work was made public, some of our findings were independently derived by Flammia and Liu [34]

[1] T. Yamamoto, Y. A. Pashkin, O. Astafiev, Y. Nakamura, and J. S. Tsai, Nature (London) **425**, 941 (2003).
[2] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, Nature (London) **432**, 602 (2004).
[3] H. Haffner, W. Hansel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Korber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Guhne, W. Dur, and R. Blatt, Nature (London), **438** 643 (2005).
[4] P. Walther, K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger, Nature (London) **434**, 169 (2005).
[5] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, R. Reichle, and D. J. Wineland, Nature (London) **438**, 639 (2005).
[6] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, Phys. Rev. Lett. **106**, 130506 (2011).
[7] K. Vogel and H. Risken, Phys. Rev. A **40**, 2847 (1989).
[8] J. F. Poyatos, J. I. Cirac, and P. Zoller, Phys. Rev. Lett. **78**, 390 (1997).
[9] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, Phys. Rev. A **77**, 032322 (2008).
[10] R. Blume-Kohout, New J. Phys. **12**, 043034 (2010).
[11] P. Aliferis, D. Gottesman, and J. Preskill, Quantum Inf. Comput. **6**, 97 (2006).
[12] E. Knill, R. Laflamme, and G. J. Milburn, Nature (London) **409**, 46 (2001).
[13] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).
[14] A. Barenco, A. Ekert, K.-A. Suominen, and P. Törmä, Phys. Rev. A **54**, 139 (1996).
[15] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Nature Commun. **1**, 149 (2010).
[16] D. Gross , Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Phys. Rev. Lett. **105**, 150401 (2010).
[17] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.107.210404 for details.
[18] E. P. Wigner, Phys. Rev. **40**, 749 (1932); K. E. Cahill and R. J. Glauber, Phys. Rev. **177**, 1882 (1969).
[19] D. Leibfried, D. M. Meekhof, B. E. King, C. Monroe, W. M. Itano, and D. J. Wineland, Phys. Rev. Lett. **77**, 4281 (1996); L. G. Lutterbach and L. Davidovich, Phys. Rev. Lett. **78**, 2547 (1997); P. Bertet, A. Auffeves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. M. Raimond, and S. Haroche, Phys. Rev. Lett. **89**, 200402 (2002); M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, J. Wenner, J. M. Martinis, and A. N. Cleland, Nature (London) **459**, 546 (2009).
[20] A. Jamiołkowski, Rep. Math. Phys. **3**, 275 (1972).M.-D. Choi, Linear Algebra Appl. **10**, 285 (1975).
[21] M. Horodecki, P. Horodecki, and R. Horodecki, Phys. Rev. A **60**, 1888 (1999).
[22] A. Bendersky, F. Pastawski, and J. P. Paz, Phys. Rev. Lett. **100**, 190403 (2008).
[23] D. Gottesman and I. L. Chuang, Nature (London) **402**, 390 (1999).
[24] R. Raussendorf and H. J. Briegel, Phys. Rev. Lett. **86**, 5188 (2001).
[25] D. Gottesman, Stabilizer Codes and Quantum Error Correction, Ph.D. thesis, California Institute of Technology, 1997.
[26] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Science **317**, 1893 (2007); O. Moussa, M. P. da Silva, and R. Laflamme (2011) (unpublished).
[27] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).
[28] D. Gottesman "Identifying Stabilizer States", Perimeter Institute Seminar, http://pirsa.org/08080052 (2008).
[29] R. A. Low, Phys. Rev. A **80**, 052314 (2009).
[30] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, and S. Lloyd, SIAM J. Comput. **37**, 166 (2007).
[31] F. Verstrate, M. M. Wolf, and J. I. Cirac, Nature Phys. **5**, 633 (2009).
[32] A. Shabani, M. Mohseni, S. Lloyd, R. L. Kosut, and H. Rabitz, Phys. Rev. A **84**, 012107 (2011).
[33] E. H. Lieb and D. W. Robinson, Commun. Math. Phys. **28**, 251 (1972); M. B. Hastings, Phys. Rev. Lett. **93**, 140402 (2004); D. Poulin, Phys. Rev. Lett. **104**, 190401 (2010).
[34] S. T. Flammia and Y.-K. Liu, Phys. Rev. Lett. **106**, 230501 (2011).