Validating and certifying stabilizer states

Amir Kalev,[1,*] Anastasios Kyrillidis,[2,†] and Norbert M. Linke[3,‡]

[1]*Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, Maryland 20742-2420, USA*
[2]*Department of Computer Science, Rice University, Houston, Texas 77005, USA*
[3]*Joint Quantum Institute, University of Maryland, College Park, Maryland 20742-2420, USA*

We propose a measurement scheme that validates the preparation of an n-qubit stabilizer state. The scheme involves a measurement of n Pauli observables, *a priori* determined from the stabilizer state and which can be realized using single-qubit gates. Based on the proposed validation scheme, we derive an explicit expression for the worst-case fidelity, i.e., the minimum fidelity between the stabilizer state and any other state consistent with the measured data. We also show that the worst-case fidelity can be certified, with high probability, using $O(n^2)$ copies of the state.

I. INTRODUCTION

Quantum tomography is the canonical procedure for diagnosing and characterizing quantum states and operations. The main goal of quantum tomography is to obtain detailed information that would allow us to improve the performance of a quantum device. However, extracting this information from experimental data is generally a hard task, even for small-sized systems. In addition, as the resources for quantum tomography scale exponentially with the number of subsystems, fully characterizing the system is experimentally impractical, even for systems with a moderate number of qubits, and under strong assumptions such as purity and unitary dynamics [1–4].

Therefore, rather than focusing on error diagnosis of quantum processors, we are often concerned with the simpler question of validation: e.g., checking how close is the state of the system to a *target* state, where usually closeness is quantified by the fidelity figure of merit. Along this thread, Flammia and Liu [5], and da Silva *et al.* [6], proposed a measurement scheme, known as *direct fidelity estimation*, tailored to estimate the fidelity between the (unknown) state of the system and a target state, without the need to perform full quantum tomography.

In this work, we are interested in cases where the target state is an n-qubit stabilizer state. Stabilizer states constitute an important class of states, used for teleportation-based [7] and measurement-based [8] quantum computation, quantum error-correction codes [9], and quantum self-testing [10]. For such states, the result in [5] and [6] translates into a measurement of $O(\epsilon^{-2}\ln(1/\delta))$ Pauli observables (picked at random from a distribution which depends on the target state), where ϵ and δ are small user-defined quantities related to the estimation error and the failure probability of the

direct fidelity estimation scheme, respectively. Importantly, the number of observables is independent of n. In practice though, to obtain a modest, say 1%, estimation error with 2/3 success probability, we would require a measurement of roughly 10^4 Pauli observables.

Here, we propose a practical scheme for validating an n-qubit stabilizer state by measuring exactly n Pauli observables, which makes it relevant for small- to moderate-sized systems. The measurements can be realized by single-qubit gates, and the observables can be *a priori* chosen based on the target state. We also give an explicit, straightforward formula for the worst-case fidelity, i.e., the minimum fidelity between the target stabilizer state and any state consistent with the measured data. Moreover, we show that the worst-case fidelity can be certified, with high probability, using $O(n\epsilon^{-2}\ln(1/\delta))$ copies of the state of the system for each measured observable. When the worst-case fidelity is close to one, we prove that one can use our scheme to obtain a high-fidelity estimate of the state of the system.

The paper is organized as follows. In the next section, we set up the notation and briefly review the notions of quantum state validation and stabilizer states. In Sec. III, we describe a validation scheme for stabilizer states and prove its certification guarantees. Then in Sec. IV, we apply our scheme to validate the preparation of a three-qubit Greenberger-Horne-Zeilinger (GHZ) state on a trapped-ion quantum computing platform [11] and superconducting quantum computing platform [12]. Finally, we close with a discussion in Sec. V.

II. BACKGROUND AND SETUP

To set up the notions and notation used in this work, we first briefly review the basic idea of state validation and the theory of stabilizer states. In what follows, we refer to a quantum state *validation protocol* as a measurement scheme that, in the noiseless case, certifies with probability one that the state of the system is the target state, *if and only if* this is the case. In the presence of experimental noise, a validation scheme should provide a certification for the worst-case fidelity.

*amirk@umd.edu
†anastasios@rice.edu
‡linke@umd.edu

Unlike quantum tomography, in a validation protocol we may choose the measurement scheme to depend on the target state. For example, given a target state $|\Psi_0\rangle$, we may consider the two-outcome positive-operator valued measure (POVM) $\{E_0 = |\Psi_0\rangle\langle\Psi_0|, E_1 = \mathbb{1} - |\Psi_0\rangle\langle\Psi_0|\}$ as a validation scheme. In this case, the worst-case fidelity is the frequency of occurrence of the outcome E_0. If $|\Psi_0\rangle$ is an entangled state, implementing such a POVM would involve applying a nonlocal unitary to the state of the system. In contrast, in this work we propose a validation scheme for a target n-qubit stabilizer state that consists of experimentally accessible measurements of Pauli observables that can be realized using single-qubit quantum gates.

In a nutshell, an n-qubit stabilizer state $|\Psi_0\rangle$ is the unique eigenstate, with eigenvalue 1, of 2^n commuting n-qubit Pauli operators $\{P_l\}_{l=1}^{2^n}$, where $P_l \in \{\pm 1, \pm i\} \cdot \{I, X, Y, Z\}^{\otimes n}$, and I, X, Y, Z are the identity and the Pauli matrices on one qubit. Hereafter, we set $P_{2^n} = \mathbb{1} \equiv I^{\otimes n}$. The set $\{P_l\}_{l=1}^{2^n}$ forms the *stabilizer group* $\mathcal{S}$. This group is generated by a set of n Pauli operators (their choice is not unique) denoted here by $\mathcal{G} \subset \mathcal{S}$. The stabilizer state can then be written as

$$\rho_0 \equiv |\Psi_0\rangle\langle\Psi_0| = \frac{1}{2^n}\prod_{P_l \in \mathcal{G}}(\mathbb{1} + P_l) = \frac{1}{2^n}\sum_{P_l \in \mathcal{S}} P_l. \quad (1)$$

III. VALIDATION OF STABILIZER STATES

As reflected by Eq. (1), $|\Psi_0\rangle$ is the unique eigenstate, with eigenvalue $+1$ of the n generators $P_l \in \mathcal{G}$. This, in turn, implies that the solution to the following feasibility problem:

$$\begin{aligned} \underset{\rho \in \mathbb{C}^{2^n \times 2^n}}{\text{find}} \quad & \rho \\ \text{subject to} \quad & \text{Tr}(\rho P_l) = 1, \ \forall P_l \in \mathcal{G}, \\ & \rho \succeq 0, \ \text{Tr}(\rho) = 1, \end{aligned} \quad (2)$$

is a singleton ρ_0. To see that, assume that there exists a density matrix $\rho_1 \neq \rho_0$ which is a solution to the feasibility problem above. When writing it in its eigenbasis, $\rho_1 = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$, the feasibility conditions $\text{Tr}(\rho_1 P_l) = 1$ imply that $\sum_j \lambda_j \langle\psi_j|P_l|\psi_j\rangle = 1$, $\forall P_l \in \mathcal{G}$. Since the P_l's have eigenvalues ± 1, we have $-1 \leqslant \langle\psi_j|P_l|\psi_j\rangle \leqslant 1$, and due to the positivity of the λ_j's, we obtain the inequality $-1 = -\sum_j \lambda_j \leqslant \sum_j \lambda_j \langle\psi_j|P_l|\psi_j\rangle \leqslant \sum_j \lambda_j = 1$. The upper bound is obtained when $\langle\psi_j|P_l|\psi_j\rangle = 1$, $\forall P_l \in \mathcal{G}$ and $\forall |\psi_j\rangle$ in the eigenbasis of ρ_1. But since, by definition, $|\Psi_0\rangle$ is the unique pure state for which $\langle\Psi_0|P_l|\Psi_0\rangle = 1$, $\forall P_l \in \mathcal{G}$, we obtain $\rho_1 = \rho_0$, in contradiction to our initial assumption.

Note, however, that while ρ_0 is the unique solution to program (2), there are infinitely many Hermitian matrices with negative eigenvalues (i.e., that do not satisfy the constraint $\rho \succeq 0$ above) for which $\text{Tr}(\rho P_l) = 1$, $\forall P_l \in \mathcal{G}$. All of which have the structure $\frac{1}{2^n}\sum_{P_l \in \mathcal{S}} P_l + \sum_{P_l \notin \mathcal{S}} c_l P_l$ (considering those with trace 1) for some real numbers c_l. Therefore, constraining on density matrices in (2) is crucial to obtain a singleton solution.

Hence, in the absence of noise, the only quantum state that is consistent with the noiseless "data" $\text{Tr}(\rho P_l) = 1$, $\forall P_l \in \mathcal{G}$, is the stabilizer state $|\Psi_0\rangle$. Therefore, given a target n-qubit stabilizer state, we can consider the measurement of the expectation values of a set of n generators as a validation

scheme. Since the generators are mutually commuting, there is, in principle, a measurement scheme to measure them simultaneously.

Consider, for example, the case where the target state is the n-qubit GHZ stabilizer state $|\text{GHZ}_n\rangle = \frac{1}{\sqrt{2}}(|0\rangle^{\otimes n} + |1\rangle^{\otimes n})$, where $\{|0\rangle, |1\rangle\}$ are the eigenbasis of the Z Pauli matrix. To validate that this is indeed the state of the system, we have the freedom to choose a specific set of generators that can be measured. A convenient choice of stabilizer generators is the n Pauli observables, $X^{\otimes n}$ and $Z_k \otimes Z_{k+1}$ for $k = 1, \ldots, n-1$. We can measure the expectation values of these observables with two simple experimental setups, for any n. In the first setup, we measure all the qubits in the X basis, while in the second setup, we measure all the qubits in the computational basis, Z. The expectation values of the generators above can be calculated from the experimental results. Note that these measurements only involve single-qubit gates.

We note that since in the noiseless case the feasibility program (2) has a unique solution, due to convexity, in the presence of small experimental noise, the argument solution to

$$\begin{aligned} \underset{\rho \in \mathbb{C}^{2^n \times 2^n}}{\text{minimize}} \quad & \mathcal{C}(\rho) \\ \text{subject to} \quad & \text{Tr}(\rho P_l) = 1 - \varepsilon_l, \ \forall P_l \in \mathcal{G}, \\ & \rho \succeq 0, \ \text{Tr}(\rho) = 1, \end{aligned} \quad (3)$$

where $\mathcal{C}(\rho)$ is a convex function of ρ, and $\varepsilon_l \geqslant 0$ captures experimental errors, is guaranteed to be close, in fidelity, to ρ_0 (which is assumed to be close to the state of the system). This follows directly from our result, i.e., Proposition 1 below, which implies that in the case of small ε_l's, the argument solution to program (3) has fidelity at least $1 - \frac{1}{2}\sum_{l:P_l \in \mathcal{G}} \varepsilon_l$ with the target state of the system.

Next, we show that the proposed validation scheme leads to an experimentally useful lower bound on the worst-case fidelity to the target stabilizer state. Since ρ_0 is a pure state, the fidelity between ρ_0 and any other state ρ is $F = \text{Tr}(\rho_0 \rho)$. Thus, given the experimental data for the expectation values of the generators, $\tilde{\mu}_1, \ldots, \tilde{\mu}_n$, to find a lower bound on the worst-case fidelity, we can solve the following convex program:

$$\begin{aligned} \underset{\rho \in \mathbb{C}^{2^n \times 2^n}}{\text{minimize}} \quad & \text{Tr}(\rho_0 \rho) \\ \text{subject to} \quad & \text{Tr}(\rho P_l) = \tilde{\mu}_l, \ \forall P_l \in \mathcal{G}, \\ & \rho \succeq 0, \ \text{Tr}(\rho) = 1. \end{aligned} \quad (4)$$

Proposition 1. If $\sum_{l:P_l \in \mathcal{G}} \frac{1 - \tilde{\mu}_l}{2} \leqslant 1$, the solution to program (4) is

$$F_{\min} = 1 - \sum_{l:P_l \in \mathcal{G}} \frac{1 - \tilde{\mu}_l}{2}. \quad (5)$$

Otherwise, if $\sum_{l:P_l \in \mathcal{G}} \frac{1 - \tilde{\mu}_l}{2} > 1$, the solution to program (4) is $F_{\min} = 0$.

Proof. Let us order the stabilizer operators such that $P_0, \ldots, P_{n-1}$ are the measured generators. Then, an n-qubit density matrix that is consistent with the data is given by

$$\rho = \frac{1}{2^n}\left(\mathbb{1} + \sum_{l=0}^{n-1} \tilde{\mu}_l P_l + \sum_{l=n}^{2^n-1} x_l P_l\right). \quad (6)$$

We can add more terms to ρ, in the subspace that lies outside the stabilizer group, and due to the orthogonality property of Pauli observables [$\mathrm{Tr}(P_i P_j) \propto \delta_{i,j}$], the resulting state would still be consistent with the measured data. However, for the same reason, adding such terms will not change the fidelity with $|\Psi_0\rangle$. Therefore, for the purpose of the proof, without loss of generality, we can consider the density matrix of (6) as the most general state consistent with the data.

Since the Pauli observables in the stabilizer group are mutually commuting, $\rho_0 = |\Psi_0\rangle\langle\Psi_0|$ and ρ of Eq. (6) are commuting, and thus can be diagonalized simultaneously. Therefore, it is convenient to rewrite ρ of Eq. (6) as

$$\rho = \lambda_0 |\Psi_0\rangle\langle\Psi_0| + \sum_{k=1}^{2^n-1} \lambda_k |\Psi_k\rangle\langle\Psi_k|, \tag{7}$$

where $\{|\Psi_k\rangle\}_{k=0}^{2^n-1}$ forms an orthonormal basis for the Hilbert space of n qubits, such that for all k and $P_l \notin \mathcal{S}$, $\langle\Psi_k|P_l|\Psi_k\rangle = 0$. In the form of Eq. (7), it is clear that program (4) minimizes the eigenvalue λ_0 or, equivalently, maximizes $\sum_{k=1}^{2^n-1} \lambda_k$, while keeping all the eigenvalues non-negative.

At this point, we have the freedom to choose the basis vectors $\{|\Psi_k\rangle\}_{k=1}^{2^n-1}$. A suitable choice is to define these vectors through the projection operators $\frac{1 \pm P_l}{2}$ associated with the stabilizer generators P_l, $l = 0, \ldots, n-1$, that is,

$$|\Psi_k\rangle\langle\Psi_k| = \prod_{l=0}^{n-1} \frac{1}{2}\left[\mathbb{1} + (-1)^{b_l^{(k)}} P_l\right], \tag{8}$$

for $k = 0, 1, \ldots, 2^n - 1$, and $b_l^{(k)} = \{0, 1\}$ is the lth bit in the binary representation of k, $k = \sum_{l=0}^{n-1} b_l^{(k)} 2^l$. Writing ρ in this basis, we obtain

$$\rho = \sum_{k=0}^{2^n-1} \lambda_k \prod_{l=0}^{n-1} \frac{1}{2}\left[\mathbb{1} + (-1)^{b_l^{(k)}} P_l\right]. \tag{9}$$

The condition that ρ should be consistent with the data implies that $\mathrm{Tr}(\frac{1-P_l}{2}\rho) = \frac{1-\tilde{\mu}_l}{2}$ for all generators $P_l \in \mathcal{G}$. Using the expression for ρ of Eq. (9) together with the relations $\frac{1-P_l}{2}\frac{1+P_l}{2} = 0$ and $(\frac{1-P_l}{2})^2 = \frac{1-P_l}{2}$ yields the set of n constraints,

$$\mathrm{Tr}\left(\frac{\mathbb{1}-P_l}{2}\rho\right) = \frac{1-\tilde{\mu}_l}{2} = \sum_{k:b_l^{(k)}=1} \lambda_k, \tag{10}$$

for $l = 0, 1, \ldots, n-1$. Since the eigenvalues of the Pauli observables are ± 1, the experimental values $\tilde{\mu}_l \in [-1, 1]$ and $\frac{1-\tilde{\mu}_l}{2} \geqslant 0$. Importantly, since for $k = 0$ $b_l^{(0)} = 0$ for all l, the set of equations (10) does not contain λ_0. Moreover, the right-hand side of the set of equations (10), for $l = 0, 1, \ldots, n-1$, contains all of the eigenvalues λ_k, $k = 1 \ldots, 2^n - 1$, with various multiplicities. Therefore, summing Eq. (10) over l, we can write

$$0 \leqslant \sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} = \sum_{k=1}^{2^n-1} \lambda_k + \Lambda(\{\lambda_k\}), \tag{11}$$

where $\Lambda(\{\lambda_k\})$ denotes the sum of all the terms not in $\sum_{k=1}^{2^n-1} \lambda_k$ (its structure is not important for the proof). Since $\forall k \ \lambda_k \geqslant 0$, $\Lambda(\{\lambda_k\})$ must be non-negative as well. Let us

assume that $\sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} \leqslant 1$. Therefore, according to Eq. (11), the maximal value of the sum $\sum_{k=1}^{2^n-1} \lambda_k$ is obtained for $\Lambda(\{\lambda_k\}) = 0$, that is, $\sum_{k=1}^{2^n-1} \lambda_k = \sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2}$. In this case, the minimal value of λ_0 is

$$\lambda_0 = 1 - \sum_{k=1}^{2^n-1} \lambda_k = 1 - \sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} \geqslant 0. \tag{12}$$

If, on the other hand, $\sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} = 1 + \Delta$, for some $\Delta > 0$, then from Eq. (11), it is clear that the maximal value of the sum $\sum_{k=1}^{2^n-1} \lambda_k$ is obtained when $\Lambda(\{\lambda_k\}) = \Delta$, i.e., when $\sum_{k=1}^{2^n-1} \lambda_k = 1$. In this case, $\lambda_0 = 0$ and, thus, the worst-case fidelity with the target state is zero. ∎

We note that the eigenvalue λ_{2^l} appears only in the lth equation of (10). Therefore, for the case where $\sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} \leqslant 1$, a valid solution of Eq. (11) is given by $\lambda_{2^l} = \frac{1-\tilde{\mu}_l}{2} \geqslant 0$, and $\lambda_k = 0$ for all other values $k \neq 0$. This implies that, in this case, one density matrix that minimizes program (4) is given by

$$\hat{\rho} = \left(1 - \sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2}\right) |\Psi_0\rangle\langle\Psi_0| + \sum_{l=0}^{n-1} \left(\frac{1-\tilde{\mu}_l}{2}\right)\left(\frac{\mathbb{1}-P_l}{2}\right)$$
$$\times \prod_{\substack{P_j \in \mathcal{G} \\ j \neq l}} \left(\frac{\mathbb{1}+P_j}{2}\right). \tag{13}$$

The state $\hat{\rho}$ is a good estimate for the state of the system when the fidelity is very close to 1, i.e., when $\tilde{\mu}_l = 1 - \varepsilon_l$ for small ε_l.

Next, we provide a certification for the worst-case fidelity. Since this certification is of interest for large values of $F_{\min}$, we will implicitly assume that $\sum_{l=0}^{n-1} \frac{1-\tilde{\mu}_l}{2} \leqslant 1$.

Proposition 2. Fix the parameters $\epsilon > 0$ and $\delta > 0$, and use $m_l = \lceil \frac{n^2 \ln(2/\delta)}{2\epsilon^2} \rceil$ copies of the state of the system, ϱ, to measure the Pauli generator P_l, for $l = 0, 1, \ldots, n-1$. Then, with probability at least $1 - \delta$, the fidelity $F_{\min}$ lies in the range $[\bar{F}_{\min} - \frac{\epsilon}{2}, \bar{F}_{\min} + \frac{\epsilon}{2}]$, where $\bar{F}_{\min}$ is obtained by replacing $\tilde{\mu}_l$ with $\mu_l = \mathrm{Tr}(\varrho P_l)$ in Eq. (5).

Proof. Let $y_j^{(l)} \in \{-1, +1\}$ be the outcome of measuring the generator P_l on the jth copy of ϱ. The empirical expectation value of P_l is given by $\tilde{\mu}_l = \frac{1}{m_l} \sum_{j=1}^{m_l} y_j^{(l)}$. By Hoeffding's inequality applied to the Bernoulli experiment, the probability that $\tilde{\mu}_l$ is ϵ-close to its mean $\mu_l = \mathrm{Tr}(\varrho P_l)$ is [13]

$$P(|\tilde{\mu}_l - \mu_l| \leqslant \epsilon) \geqslant 1 - 2e^{-2m_l \epsilon^2}. \tag{14}$$

Therefore, taking $m_l = \lceil \frac{\ln(2/\delta)}{2\epsilon^2} \rceil \ \forall l$, with probability at least $1 - \delta$, the data $\tilde{\mu}_l$ lie in the range $[\mu_l - \epsilon, \mu_l + \epsilon]$. Thus, with probability at least $1 - \delta$, we have

$$\bar{F}_{\min} - \frac{n\epsilon}{2} = 1 - \sum_{l=0}^{n-1} \frac{1-\mu_l+\epsilon}{2} \leqslant F_{\min}$$

$$\leqslant 1 - \sum_{l=0}^{n-1} \frac{1-\mu_l-\epsilon}{2} = \bar{F}_{\min} + \frac{n\epsilon}{2}. \tag{15}$$

Taking $\epsilon \to \epsilon/n$ completes the proof. ∎

Following Propositions 1 and 2, we obtain the following corollary:

Corollary 1. Let ϱ be the state of the system and let $m_l = \lceil \frac{n^2 \ln(2/\delta)}{2\epsilon^2} \rceil \; \forall l$. Then, with probability at least $1 - \delta$, the fidelity between ϱ and the target stabilizer state $|\Psi_0\rangle$ is larger than $F_{\min} - \frac{\epsilon}{2}$.

Proof. $\bar{F}_{\min}$ is the solution for program (4), i.e., the minimum fidelity with $|\Psi_0\rangle$, when we have access to the noiseless data $\mu_l = \text{Tr}(\varrho P_l)$. Therefore, the fidelity of the state of the system ϱ and $|\Psi_0\rangle$ is necessarily larger than (or equals to) $\bar{F}_{\min}$. On the other hand, from the right-hand side of inequality (16), with probability $1 - \delta$, we can bound $\bar{F}_{\min}$ from below by $F_{\min} - n\frac{\epsilon}{2}$. Taking $\epsilon \to \epsilon/n$ completes the proof. ∎

We note that since in the proposed procedure we should measure n stabilizer generators, the total sample complexity of our certification protocol is $\lceil \frac{n^3 \ln(2/\delta)}{2\epsilon^2} \rceil$. For comparison, Gottesman [14] and Montanaro [15] have showed that stabilizer states can be identified using only $O(n)$ copies of the state. However, these methods require entangled measurements, while the method proposed in this work uses single-qubit gates. Moreover, the direct fidelity estimation scheme of [5] and [6] requires $O\{\text{poly}[\epsilon^{-1}, \ln(1/\delta)]\}$ copies of the states in total for certification, independent of n. But, as was discussed above, their scheme is favorable for very large systems, while our scheme is applicable for moderate-size systems.

Hoeffding's inequality, used above, does not take into account the information about the variance of the outcome's distribution. Including this information, e.g., by using Bernstein's inequality, can further improve the sample complexity of our scheme. Bernstein's inequality states that given independent random variable $a_1, a_2, \dots, a_m$ with mean μ and variance σ^2, the probability that $\tilde{\mu} = \frac{1}{m} \sum_j a_j$ is ϵ-close to μ is

$$P(|\tilde{\mu} - \mu| \leqslant \epsilon) \geqslant 1 - 2e^{-\frac{m\epsilon^2}{2(\sigma^2 + \epsilon/3)}}. \quad (16)$$

When $\sigma^2 \lesssim \epsilon$, we obtain a tighter lower bound than of Eq. (14) which behaves like $1 - 2e^{-m\epsilon}$ instead of the $1 - 2e^{-m\epsilon^2}$, i.e., in this case, the number of samples (per observable) m scales as $\frac{n \ln(2/\delta)}{\epsilon}$ rather than $\frac{n^2 \ln(2/\delta)}{\epsilon^2}$. This can be useful for our purpose when the state of the system is close to the target stabilizer state since then we expect the measurement outcomes of the n Pauli observables to be narrowly distributed. (The measurement outcomes for the target stabilizer state have zero variance.) However, in our case, since the actual state of the system ϱ is unknown, the variance of the distribution of each measured observable $[1 - \text{Tr}(\varrho P_l)^2]^2$ is unknown and, in practice, we cannot use the Bernstein's inequality as stated above. Therefore, instead, we will use an empirical Bernstein's bound developed by [16,17], which uses a variance calculated from the data, $\tilde{\sigma}^2$, instead of σ^2. The main idea of the empirical Bernstein's bound of [16,17] is to use an online algorithm, called EBStop, which decides when to stop taking data so that the $\tilde{\sigma}^2$ is, with high probability, an upper bound for σ^2. The algorithm should be executed, in our case, for each one of the n measured generators. We refer the reader to [17] for details about the algorithm. It was proven [17,18] that it takes $m = O(\epsilon^{-1} \ln \frac{1}{\delta})$ samples for the above algorithm to stop and to assure the

desire convergence of $\tilde{\mu}$, with probability at least $1 - \delta$. Therefore, by taking $\epsilon \to \epsilon/n$, we find the following result:

Corollary 2. Let ϱ be the state of the system and let $m_l = O(\frac{n \ln(1/\delta)}{\epsilon}) \; \forall l = 0, 1, \dots, n - 1$. Then, with probability at least $1 - \delta$, the fidelity between ϱ and the target stabilizer state $|\Psi_0\rangle$ is larger than $F_{\min} - \frac{\epsilon}{2}$.

IV. EXPERIMENTS

We apply the scheme described above to certify the preparation of a three-qubit GHZ state, $|\text{GHZ}_3\rangle = \frac{1}{\sqrt{2}}(|0\rangle^{\otimes 3} + |1\rangle^{\otimes 3})$, on two different experimental platforms. The first platform is a five-qubit quantum computing system based on trapped ytterbium ions with individual laser beam addressing [11]. There, the state preparation consists of two native entangling XX gates on qubits 1–2 and 2–3, followed by Hadamard gates on all three, as shown in Fig. 1. The XX gate is experimentally realized using the Mølmer-Sørensen scheme [19–21] combined with laser-pulse shaping to control the motional modes in a multi-ion chain [22]. The gate operation is given in Eq. (17). The parameter χ can be varied continuously, $0 < \chi \leqslant \frac{\pi}{4}$, and is set to $\frac{\pi}{4}$ in this case for maximally entangling gates. Single-qubit gates, which are needed to create the Hadamard operations, are generated by driving resonant Rabi rotations of defined phase, amplitude, and duration. Z rotations form an exception as they are applied as phase advances on the classical controllers,

$$XX(\chi) = \begin{pmatrix} \cos(\chi) & 0 & 0 & -i\sin(\chi) \\ 0 & \cos(\chi) & -i\sin(\chi) & 0 \\ 0 & -i\sin(\chi) & \cos(\chi) & 0 \\ -i\sin(\chi) & 0 & 0 & \cos(\chi) \end{pmatrix}. \quad (17)$$

The experimental data consist of 11 000 samples from the distribution associated with measuring each qubit in the Z basis (the computational basis), and 11 000 samples from the distribution associated with measuring each qubit in the X basis (the Hadamard basis). The latter is achieved by applying another round of Hadamard gates before a measurement in the computational basis. The data are corrected for readout errors, which are characterized independently by measuring the state-to-state transfer matrix in a series of reference experiments [11].

For comparison, we apply the validation scheme described above to validate a three-qubit GHZ state prepared on an IBM superconducting quantum computation platform. We used the five-qubit Tenerife processor [12]. The three qubits, labeled 0, 1, and 2, were prepared in a GHZ state by applying
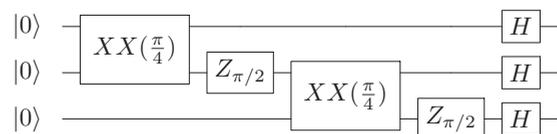

FIG. 1. Circuit used to generate the three-qubit GHZ state in the trapped ion system. The XX-gate action is given in Eq. (17).
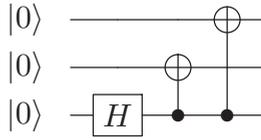
FIG. 2. Circuit used to generate the three-qubit GHZ state on IBM's superconducting quantum computation platform.

the Hadamard transformation to qubit 2, followed by CNOT operations between qubits 2 and 1 and between qubits 2 and 0; see Fig. 2. Here, as well, the experimental data consist of 8192 samples from the distribution associated with measuring each qubit in the Z basis, and 8192 samples from the distribution associated with measuring each qubit in the X basis. The latter is obtained by applying the Hadamard transformation to each qubit followed by a measurement in the computational basis.

In Figs. 3 and 4, we show the results of the two experiments in terms of the occurrence frequencies of the computational basis states. Based on these results, using Corollary 2, we calculate for the two experiments the lower bound on fidelity between the prepared state and the three-qubit GHZ state, $|\text{GHZ}_3\rangle$, for various certificate values. These are given in Table I. For the ion-trap (IBM) experiment, we obtain these quantities by calculating the expected values of the three generators $Z_1 \otimes Z_2$, $Z_2 \otimes Z_3$ ($Z_1 \otimes Z_3$), and $X^{\otimes 3}$. For the observed data, this choice of generators maximizes the lower bound on the fidelity.

V. RELATED WORK AND DISCUSSION

The results in Table I yield a clear difference between the lower fidelity bounds of the two quantum devices investigated. Other comparisons between these platforms have shown only small differences in performance for applications where the limited connectivity of the superconducting system did not have an impact [23]. Consistent with these previous findings, the measurement outcomes on both systems represent the target distributions reasonably well, as the average probability for measuring one of the desired states is only about 15%
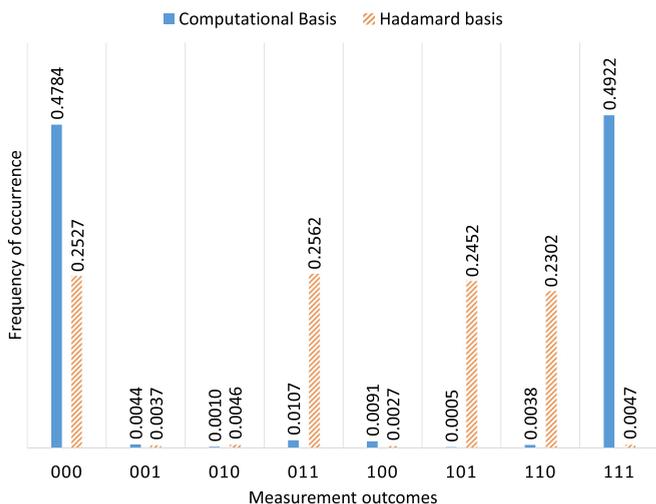

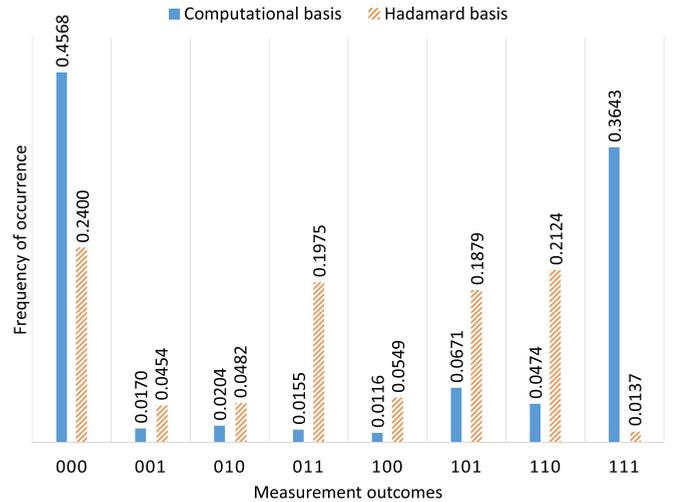FIG. 3. Experimental results for the trapped-ion platform.

FIG. 4. Experimental results for the IBM Q Tenerife processor.

lower on the IBM machine; see Figs. 3 and 4. What our analysis illustrates is that the certified lower-bound fidelity metric punishes these small differences as they disproportionately increase the overlap of the measurement outcomes with other potential distributions consistent with non-GHZ states, and hence resulting in much lower fidelity bounds on the IBM results. Measuring Pauli observables outside a specific set of generators may help to narrow down the set of non-GHZ states that are consistent with the measurement data, and can increase the lower bound on the fidelity.

The protocol we presented does not deal with measurement errors and assumes that we are able to exactly measure the desired set of generators. In a real experimental setup, there could be a bias (either random or systematic) in the measured observables, which can decrease the lower bound on the fidelity. For example, consider a systematic bias in the measurement of a single qubit, where instead of measuring the projection onto $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$, we are, in practice, measuring the POVM $\{(1 - \eta)|0\rangle\langle 0| + \eta|1\rangle\langle 1|, (1 - \eta)|1\rangle\langle 1| + \eta|0\rangle\langle 0|\}$ (in the x, y, and z directions), for small η. Then the lower bound on the fidelity decreases by $O(n^2\eta)$ when the state of the system is close to the stabilizer state. It is, therefore, important to devise state-preparation validation schemes that are robust to measurement errors. We leave that for future work.

Finally, we wish to mention the recent work by Rocchetto [24], who showed that stabilizer states are efficiently probably approximately correct (PAC)-learnable. Based on the

TABLE I. Based on the result shown in Figs. 3 and 4, we conclude that in these two experiments, with probability larger than p_{conf}, the state of the system has fidelity higher than the tabulated value with the three-qubit GHZ state.

	Fidelity lower bound	
p_{conf}	Ion-trap	IBM
0.999	0.905	0.575
0.99	0.913	0.585
0.9	0.923	0.596

results of Aaronson [25], Rocchetto showed that stabilizer states can be learned, in the computational learning theory sense, with only $O(n)$ copies of the state, and proposed a learning procedure which involves an optimization problem that can be solved on a classical computer in polynomial time. The PAC-learnability of stabilizer (GHZ) states was also demonstrated experimentally in [26]. The goal in the PAC-learning methodology is to provide a model for the state that produces (with high probability) good predictions of future experimental results. Nevertheless, even in the case of small experimental noise, in general, the learned state may have poor fidelity with the actual state of the system (which is assumed to be closed to the target state). This is in contrast to the validation method presented here, which guarantees that the estimated state has high fidelity with the state of the system in the case of small experimental noise. In fact, in the aforementioned experiment of [26], the learned state of the system was shown to have a good fidelity with the target GHZ state. The high fidelity with the target state

in [26] can be understood from the point of view of the validation scheme presented here. Specifically, the measurement distribution in [26] was chosen to be uniform over the stabilizer group of the GHZ state and, therefore, dependent on the target state. The authors reported that measuring roughly $1.2n$ stabilizer Pauli observables allowed them to obtain a high-fidelity estimation of the target state by solving a convex program. Indeed, measuring slightly more than n stabilizers assures that, with high probability, n of them are independent and, hence, form a generator group. This, as was discussed above, allows a good reconstruction of the stabilizer state.

[1] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Phys. Rev. Lett. **105**, 150401 (2010).

[2] A. Kalev, R. Kosut, and I. Deutsch, NPJ Quantum Inf. **1**, 15018 (2015).

[3] R. O'Donnell and J. Wright, in *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing, STOC '16* (ACM, New York, 2016), pp. 899–912.

[4] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, IEEE Trans. Inf. Theory **63**, 5628 (2017).

[5] S. T. Flammia and Y.-K. Liu, Phys. Rev. Lett. **106**, 230501 (2011).

[6] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Phys. Rev. Lett. **107**, 210404 (2011).

[7] D. Gottesman and I. L. Chuang, Nature (London) **402**, 390 (1999).

[8] R. Raussendorf and H. J. Briegel, Phys. Rev. Lett. **86**, 5188 (2001).

[9] D. Gottesman, Ph.D. thesis, California Institute of Technology, 1997.

[10] S. Breiner, A. Kalev, and C. A. Miller, in *Proceedings of the 15th International Conference on Quantum Physics and Logic*, edited by P. Selinger and G. Chiribella, Vol. 287 of Electronic Proceedings in Theoretical Computer Science (Open Publishing Association, 2019), pp. 43–66.

[11] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Nature (London) **536**, 63 (2016).

[12] IBM Q 5 (ibmqx2), https://github.com/QISKit/qiskit-backend-information/blob/master/backends/yorktown/V1/README.md (unpublished).

[13] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, 2014).

[14] D. Gottesman, http://pirsa.org/08080052 (unpublished).

[15] A. Montanaro, arXiv:1707.04012.

[16] P. Dagum, R. Karp, M. Luby, and S. Ross, SIAM J. Comput. **29**, 1484 (2000).

[17] V. Mnih, C. Szepesvári, and J.-Y. Audibert, in *Proceedings of the 25th International Conference on Machine Learning, ICML '08* (ACM, New York, 2008), pp. 672–679.

[18] V. Mnih, Ph.D. thesis, University of Alberta, 2008.

[19] K. Mølmer and A. Sørensen, Phys. Rev. Lett. **82**, 1835 (1999).

[20] E. Solano, R. L. de Matos Filho, and N. Zagury, Phys. Rev. A **59**, R2539(R) (1999).

[21] G. Milburn, S. Schneider, and D. James, Fortschr. Phys. **48**, 801 (2000).

[22] T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, Phys. Rev. Lett. **112**, 190502 (2014).

[23] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Proc. Natl. Acad. Sci. USA **114**, 3305 (2017).

[24] A. Rocchetto, Quantum Inf. Comput. **18**, 541 (2018).

[25] S. Aaronson, Proc. R. Soc. A **463**, 3089 (2007).

[26] A. Rocchetto, S. Aaronson, S. Severini, G. Carvacho, D. Poderini, I. Agresti, M. Bentivegna, and F. Sciarrino, Sci. Adv. **5**, eaau1946 (2019).