Improved Stabilizer Estimation via Bell Difference Sampling

Sabee Grewal
sabee@cs.utexas.edu
University of Texas at Austin
Austin, Texas, USA

Vishnu Iyer
vishnu.iyer@utexas.edu
University of Texas at Austin
Austin, Texas, USA

William Kretschmer
kretsch@berkeley.edu
Simons Institute for the Theory of Computing
Berkeley, California, USA

Daniel Liang
dl88@rice.edu
Rice University
Houston, Texas, USA

ABSTRACT

We study the complexity of learning quantum states in various models with respect to the stabilizer formalism and obtain the following results:

- We prove that a linear number of T-gates are necessary for any Clifford+T circuit to prepare computationally pseudorandom quantum states, an exponential improvement over the previously known bound. This bound is asymptotically tight if linear-time quantum-secure pseudorandom functions exist.
- Given an n-qubit pure quantum state $|\psi\rangle$ that has fidelity τ with some stabilizer state, we give an algorithm that outputs a succinct description of a stabilizer state that witnesses fidelity at least $\tau - \varepsilon$. The algorithm uses $O(n/(\varepsilon^2\tau^4))$ samples and $\exp\left(O(n/\tau^4)\right)/\varepsilon^2$ time. In the regime of τ constant, this algorithm estimates stabilizer fidelity substantially faster than the naïve $\exp(O(n^2))$-time brute-force algorithm over all stabilizer states.
- In the special case of $\tau > \cos^2(\pi/8)$, we show that a modification of the above algorithm runs in polynomial time.
- We exhibit a tolerant property testing algorithm for stabilizer states.

The underlying algorithmic primitive in all of our results is Bell difference sampling. To prove our results, we establish and/or strengthen connections between Bell difference sampling, symplectic Fourier analysis, and graph theory.

CCS CONCEPTS

• **Theory of computation** → *Sample complexity and generalization bounds*; *Pseudorandomness and derandomization*; *Design and analysis of algorithms*; *Quantum complexity theory*; *Cryptographic primitives*; • **Mathematics of computing** → *Probabilistic inference problems.*

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649738

KEYWORDS

Bell difference sampling, stabilizer formalism, Clifford group, pseudorandom quantum states, quantum learning theory, property testing, quantum algorithms

ACM Reference Format:
Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. 2024. Improved Stabilizer Estimation via Bell Difference Sampling. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3618260.3649738

1 INTRODUCTION

A central goal in quantum information is to understand which quantum states are efficiently learnable. While many quantum state learning algorithms are extremely efficient in sample complexity [2, 19, 39], fewer classes of time-efficiently-learnable quantum states are known. One such example is the class of *stabilizer states*, which are n-qubit states that are stabilized by a group of 2^n commuting Pauli matrices.[1] Stabilizer states are well-studied because of their broad importance and widespread applications throughout quantum information, including in quantum error correction [20, 26, 60], efficient classical simulation of quantum circuits [16, 18], randomized benchmarking [44], and measurement-based quantum computation [57], to name a few examples.

The first computationally efficient algorithm for learning a complete description of an unknown stabilizer state was given by Montanaro [53].[2] Given copies of a stabilizer state $|\phi\rangle$, Montanaro's algorithm utilizes the algebraic properties of Pauli matrices and Bell-basis measurements to efficiently learn the generators of the stabilizer group of $|\phi\rangle$, which suffices to determine $|\phi\rangle$. More specifically, Montanaro (implicitly) introduced *Bell difference sampling*, which, at a high level, is an algorithmic primitive that takes copies of some state and induces a measurement distribution on Pauli matrices. Bell difference sampling was studied more thoroughly in [33] and has seen extended success in the development of algorithms for stabilizer states and states that are close to stabilizer states [30, 33, 36, 49, 53].

[1] Some other examples of state classes that admit time-efficient learning algorithms include matrix product states [21], non-interacting fermion states [5], and certain classes of phase states [10].
[2] In 2008, Gottesman gave a short video lecture explaining how to learn stabilizer states, based on joint work with Aaronson [4]. However, the details of this algorithm were never published.

In this work, we extend the use of Bell difference sampling to give faster, more general, and otherwise improved algorithms for learning properties of quantum states related to the stabilizer formalism. By understanding how these properties affect the Bell difference sampling distribution, we are able to find relevant certificates of these properties faster than the previous state-of-the-art.

1.1 Our Results

Tight Pseudorandomness Bounds. Pseudorandom states are a quantum cryptographic primitive that have recently attracted much attention in quantum cryptography and complexity theory. They can be thought of as a quantum analogue of pseudorandom generators, with the main difference being that pseudorandom states mimic the Haar measure over n-qubit states, rather than the uniform distribution over n-bit strings. For a formal definition, see [42].

Pseudorandom states suffice to build a wide range of cryptographic primitives, including quantum commitments, secure multiparty computation, one-time digital signatures, and more [6, 11, 31, 37, 42, 54]. The language of pseudorandom states has also been found to play a key role in resolving some paradoxes at the heart of black hole physics [13, 14]. Finally, and perhaps most surprisingly, there is recent evidence to suggest that pseudorandom states can be constructed without assuming the existence of one-way functions [46, 47].

Collectively, these results have motivated recent works that seek to characterize what computational properties or resources are required of pseudorandom states. For example, [3] investigates the possibility of building pseudorandom quantum states with limited entanglement, and prove the existence of pseudorandom state ensembles with entanglement entropy substantially smaller than n, assuming the existence of quantum-secure one-way functions.

Analogously, Grewal, Iyer, Kretschmer, and Liang [30] study quantum pseudorandomness from the perspective of stabilizer complexity. They treat the number of non-Clifford gates in a circuit as a resource, similar to size or depth. The main result of [30] shows that states having fidelity at least $\frac{1}{\mathrm{poly}(n)}$ with a stabilizer state cannot be computationally pseudorandom. As a consequence, they deduce that $\omega(\log n)$ non-Clifford gates are necessary for a family of circuits to yield an ensemble of pseudorandom quantum states.

We give an exponential improvement on this lower bound:[3]

Theorem 1.1 (Informal version of Corollary 3.9). *Any family of Clifford circuits that produces an ensemble of n-qubit computationally pseudorandom quantum states must use at least $n/2$ auxiliary non-Clifford single-qubit gates.*

In the special case that the non-Clifford gates are all diagonal (e.g. T-gates), our lower bound improves to n.

Under plausible computational assumptions, Theorem 1.1 is tight up to constant factors. In particular, the existence of linear-time quantum-secure pseudorandom functions implies the existence of linear-time constructible pseudorandom states [15, 30], which of course have at most $O(n)$ non-Clifford gates. Note that linear-time *classically*-secure pseudorandom functions are strongly believed to exist [23, 41], and it seems conceivable that these constructions remain secure against quantum adversaries.

We remark that Theorem 1.1 bears analogy to a recent result of Leone, Oliviero, Lloyd, and Hamma [51] that information scrambled by an n-qubit unitary implemented with Clifford gates and $t < n$ T-gates can be efficiently unscrambled. In particular, both Theorem 1.1 and [51] establish different forms of non-pseudorandomness (for states and unitaries, respectively) in the same parameter regime of non-Cliffordness.

Faster Stabilizer State Approximation. As noted earlier, one of the prominent applications of stabilizer states is in classical simulation algorithms of quantum circuits. Such algorithms work by modeling the output state of a quantum circuit as a decomposition of stabilizer states (e.g., as a linear combination) [16]. The runtime of these algorithms then scale with respect to one of several measures of the "amount of non-stabilizerness" in this decomposition. These measures are sometimes called *magic monotones* [64, Definition 3] [34, Definition 3], because they are non-increasing under Clifford operations. Typically, magic monotones increase exponentially as non-Clifford gates are applied.[4] Examples of well-known magic monotones include the stabilizer rank, stabilizer extent, and inverse of stabilizer fidelity [16].

A series of recent and simultaneous works have explored the question of whether magic monotones can be estimated efficiently, or whether states with low magic are efficiently learnable. For example, recall that [30] showed that states with non-negligible stabilizer fidelity are weakly learnable, in the sense that they are efficiently distinguishable from random. [27, 28, 35, 50] proved that states with bounded stabilizer nullity are efficiently learnable, and [27] also gave an efficient property tester for stabilizer nullity. [34] showed that various magic monotones *cannot* be estimated efficiently in certain parameter regimes, by constructing states with low magic that are cryptographically indistinguishable from states with large magic. Finally, [8, 10] raised the question of whether states of bounded stabilizer rank are efficiently learnable.

Our second result is a further contribution towards understanding the learnability of low-magic states: we give an algorithm that finds stabilizer state approximations of states with non-negligible *stabilizer fidelity*. As its name suggests, stabilizer fidelity (denoted $F_{\mathcal{S}}(|\psi\rangle)$) measures how close a state $|\psi\rangle$ is to a stabilizer state: it is simply the maximum of $|\langle\phi|\psi\rangle|^2$ over all stabilizer states $|\phi\rangle$. Hence, it is not hard to see that the inverse of stabilizer fidelity is a magic monotone. Assuming $|\psi\rangle$ has stabilizer fidelity at least τ, our algorithm returns a stabilizer state that witnesses overlap at least $F_{\mathcal{S}}(|\psi\rangle) - \varepsilon$ with $|\psi\rangle$.

Theorem 1.2 (Informal version of Theorem 4.9). *Fix $\tau > \varepsilon > 0$. There is an algorithm that, given copies of an n-qubit pure state $|\psi\rangle$ with $F_{\mathcal{S}}(|\psi\rangle) \geq \tau$, returns a stabilizer state $|\phi\rangle$ that satisfies $|\langle\phi|\psi\rangle|^2 \geq F_{\mathcal{S}}(|\psi\rangle) - \varepsilon$ with high probability. The algorithm uses $O(n/(\varepsilon^2\tau^4))$ copies of $|\psi\rangle$ and $\exp\left(O(n/\tau^4)\right)/\varepsilon^2$ time.*

[3] We remark that while the result of [30] is not tight in terms of the number of non-Clifford gates, recent work [3] shows that [30]'s bound in terms of stabilizer fidelity is optimal up to polynomial factors, because [3] constructs pseudorandom state ensembles with any inverse-superpolynomial stabilizer fidelity (assuming quantum-secure one-way functions exist).

[4] Some authors prefer to work with the logarithm of the monotone, so that they scale linearly as non-Clifford gates are applied.

To our knowledge, this is the first nontrivial algorithm to approximate an arbitrary quantum state with a stabilizer state.[5] Indeed, we are not aware of any prior algorithm better than a brute-force search over all stabilizer states, which takes $2^{O(n^2)}$ time and $O(n^2)$ samples.[6] Thus our algorithm offers a substantial improvement in the regime of $\tau = \omega(n^{-1/4})$. Arguably, the most interesting setting of parameters is constant τ, in which case we have a quadratic improvement in sample complexity and a superpolynomial improvement in time complexity.

Observe that, because we output a witness of stabilizer fidelity at least $\tau - \varepsilon$ with high probability, assuming a state with fidelity τ exists, our algorithm can additionally be used as a subroutine to *estimate* stabilizer fidelity and, moreover, find a stabilizer state that witnesses this. More precisely, if the goal is to estimate stabilizer fidelity to accuracy $\pm\varepsilon$, then one can break $[0, 1]$ into intervals of width ε and perform a binary search procedure using our algorithm. Overall, this takes $O(n/\varepsilon^6)$ samples and $\exp(O(n/\varepsilon^4))$ time.

As an application, our stabilizer state approximation algorithm could be used to search for better stabilizer decompositions of magic states. Recall that magic states are states that, when injected into Clifford circuits, allow for universal quantum computation [17]. The best-known algorithms for simulating quantum circuits dominated by Clifford gates use decompositions of magic states into linear combinations of stabilizer states and have a runtime that scales polynomially in the complexity of the decomposition, either in terms of the stabilizer rank or stabilizer extent [16]. Hence, better stabilizer decompositions of magic states yield faster algorithms. These decompositions are often obtained by writing the tensor product of a small number of magic states (usually on the order of 10 qubits) as linear combination of a slightly larger number of stabilizer states [18, 45]. Therefore, if a classical simulation of our algorithm could be made practical for (say) $n \approx 15$ qubits, there is reason to believe that running this algorithm on magic states, combined with a meta-algorithm such as matching pursuit [52], could find better stabilizer decompositions of magic states and, as a result, improve the runtime of near-Clifford simulation.

Finally, we remark that the problem we solve is similar in spirit to the agnostic probably approximately correct (PAC) learning framework [43, 62]. In the agnostic PAC model, a learner is given labeled training data $\{(x_1, y_1), \ldots, (x_m, y_m)\}$ from some unknown distribution $\mathcal{D}$, as well as some concept class C to choose a hypothesis from. The goal of the learner is to find a hypothesis function $h \in C$ that approximates the best fit for the training data, even though no function in C will necessarily fit the training data perfectly. In an analogous fashion, our algorithm finds a stabilizer state $|\phi\rangle$ that approximates the best fit for $|\psi\rangle$ over the set of stabilizer states, which need not contain $|\psi\rangle$. We note that Aaronson studied PAC learning of quantum states in the so-called realizable setting [1]. However, agnostic PAC learning of quantum states has not yet appeared in the literature.

[5]We thank David Gosset (personal communication) for bringing this barrier to our attention.

[6]The polynomial sample complexity follows from a straightforward application of the classical shadows framework [39]. See [32, Corollary 21] for a proof that there are $2^{\Theta(n^2)}$ many stabilizer states.

Bounded-Distance Stabilizer Approximation. Although our stabilizer state approximation algorithm significantly improves upon brute force, it still requires exponential time in general. One might wonder whether this exponential runtime is necessary. For example, is it possible that finding stabilizer state approximations is computationally hard, even for states whose distance to the nearest stabilizer state is bounded by some small constant? *A priori*, this might even be expected, because in other contexts, learning stabilizer states with a constant rate of noise can be as hard as the Learning Parities with Noise (LPN) problem [25, 38], which is believed to be hard. What if the stabilizer fidelity is large enough to guarantee the existence of a *unique* closest stabilizer state? Our third result shows that in this regime, a modification of the algorithm from Theorem 1.2 is computationally efficient. In particular, this modification works when the stabilizer fidelity is larger than $\cos^2(\pi/8) \approx 0.8536$, which is precisely threshold above which $|\psi\rangle$ is guaranteed to have a unique closest stabilizer state.

Theorem 1.3 (Informal version of Theorem 5.7). *Fix $\gamma > 0$. There is an algorithm that, given copies of an n-qubit pure state $|\psi\rangle$ that has fidelity at least $\cos^2(\pi/8) + \gamma$ with some stabilizer state $|\phi\rangle$, returns $|\phi\rangle$ with high probability. The algorithm uses $O\left(n + \frac{\log n}{\gamma^2}\right)$ copies of $|\psi\rangle$ and $O\left(n^3 + \frac{n^2 \log n}{\gamma^2}\right)$ time.*

Note that, unlike Theorem 1.2, this algorithm finds the stabilizer state $|\phi\rangle$ witnessing the maximum fidelity $F_{\mathcal{S}}(|\psi\rangle)$, rather than a (possibly different) state witnessing fidelity $F_{\mathcal{S}}(|\psi\rangle) - \varepsilon$.

Tolerant Stabilizer Testing. Our final result is a tolerant property testing algorithm for stabilizer states. In the tolerant property testing model [56], which generalizes ordinary property testing [24, 58], a tester must accept objects that are at most ε_1-far from having some property ("completeness") and reject objects that are at least ε_2-far from having that same property ("soundness") for $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$. The standard property testing model is recovered when $\varepsilon_1 = 0$, and the relaxed completeness condition generally makes tolerant testing a much harder problem. Nonetheless, the tolerant testing model is natural to consider in certain error models, such as in the presence of imprecise quantum gates.

Our result extends work by Gross, Nezami, and Walter [33], who gave a property tester (hereafter, the "GNW algorithm") for stabilizer states. When combined with the prior work of [30], we deduce the existence of a *tolerant* property testing algorithm for stabilizer states. Our algorithm takes copies of an n-qubit quantum state $|\psi\rangle$ and decides whether $|\psi\rangle$ has stabilizer fidelity at least α_1 or at most α_2, promised that one of these is the case. Note that we have taken $\alpha_1 := 1 - \varepsilon_1$ and $\alpha_2 := 1 - \varepsilon_2$ for notational simplicity.

Theorem 1.4 (Informal version of Theorem 6.1). *Fix $\alpha_1, \alpha_2 \in [0, 1]$ such that $\alpha_2 < \frac{4\alpha_1^6 - 1}{3}$, and define $\gamma := \alpha_1^6 - \frac{3\alpha_2 + 1}{4}$. There is an algorithm that uses $O(1/\gamma^2)$ copies of a quantum state $|\psi\rangle$, $O(n/\gamma^2)$ time, and decides whether $|\psi\rangle$ has stabilizer fidelity at least α_1 or at most α_2, promised that one of these is the case.*

While our algorithm does not work for all settings of ε_1 and ε_2—giving such an algorithm is an open problem—our algorithm does significantly improve over prior work. In Section 6.2, we compare

the parameter regimes in which our algorithm works to the existing literature and show those regimes visually in Fig. 1.

1.2 Our Techniques

The unifying tool in our work is *Bell difference sampling*, a measurement primitive that has recently found applications in a variety of algorithms related to stabilizer states [30, 33, 53]. We defer a full definition of Bell difference sampling to Section 2.2, but note some of its important properties here. Bell difference sampling involves measuring pairs of qubits of $|\psi\rangle^{\otimes 2}$ in the Bell basis, repeating again with $|\psi\rangle^{\otimes 2}$, and combining the measurements to interpret the result as corresponding to an n-qubit Pauli operator. Overall, this consumes four copies of $|\psi\rangle$, though it only performs measurements across two copies of $|\psi\rangle$ at a time. It will be most convenient to parameterize the sampled Pauli operators by strings in $\mathbb{F}_2^{2n}$, which we do as follows. For $x = (a, b) \in \mathbb{F}_2^{2n}$, where a and b are the first and last n bits of x, respectively, we define the *Weyl operator* W_x as

$$W_x := i^{a\cdot b} X^{a_1} Z^{b_1} \otimes \cdots \otimes X^{a_n} Z^{b_n}.$$

Importantly for us, the Weyl operators form an orthogonal basis for $\mathbb{C}^{2^n \times 2^n}$, and so they give rise to the *Weyl expansion* of a quantum state $|\psi\rangle$ as

$$|\psi\rangle\langle\psi| = \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^{2n}} \langle\psi|W_x|\psi\rangle\, W_x.$$

For pure states, the squared coefficients in this expansion sum to 1, and therefore form a distribution over $\mathbb{F}_2^{2n}$. We denote this distribution by $p_\psi(x) := 2^{-n} \langle\psi|W_x|\psi\rangle^2$.[7]

Gross, Nezami, and Walter [33] give an explicit form for the distribution obtained by performing Bell difference sampling. In particular, they showed that Bell difference sampling a quantum pure state $|\psi\rangle$ is equivalent to sampling from the following distribution:

$$q_\psi(x) := \sum_{a \in \mathbb{F}_2^{2n}} p_\psi(a) p_\psi(a + x),$$

i.e., the convolution of p_ψ with itself. At a high level, we establish our results by proving some structure on q_ψ and p_ψ for certain quantum states.

Tight Pseudorandomness Bounds. To prove our lower bound on the number of non-Clifford gates required to prepare pseudorandom states, we give an algorithm that distinguishes Haar-random states from quantum states prepared by circuits with fewer than $n/2$ non-Clifford single-qubit gates. The key insight is that if $|\psi\rangle$ is the output of such a circuit, then q_ψ is concentrated on a proper subspace of $\mathbb{F}_2^{2n}$, whereas for Haar-random states, q_ψ is anticoncentrated on all such subspaces with overwhelming probability over the Haar measure. Proving these properties of q_ψ reveals a simple algorithm: draw a linear number of samples from q_ψ and compute the number of linearly independent vectors in the sample. Haar-random states will have $2n$ such vectors with high probability and, otherwise, there will be strictly less than $2n$ such vectors.

Faster Stabilizer State Approximation. Our algorithms for stabilizer approximation also rely on proving anticoncentration properties of q_ψ. We begin by showing that if $|\psi\rangle$ has large fidelity with some stabilizer state $|\phi\rangle$, then q_ψ is well-supported on the n-dimensional subspace $\mathrm{Weyl}(|\phi\rangle) := \{x \in \mathbb{F}_2^{2n} : W_x |\phi\rangle = \pm|\phi\rangle\}$ of Weyl operators that stabilize $|\phi\rangle$ (up to sign). Next, we establish that if $|\phi\rangle$ is the state that maximizes stabilizer fidelity, then the mass of q_ψ on $\mathrm{Weyl}(|\phi\rangle)$ is not too concentrated on any proper subspace. Hence, by sampling from q_ψ enough times, we can be guaranteed that with high probability, $\mathrm{Weyl}(|\phi\rangle)$ will be generated by some subset of the sampled Weyl operators. By iterating through all mutually commuting subsets of the sampled Weyl operators, we compile a list of candidate stabilizer states $|\phi\rangle$ that must contain the fidelity-maximizing $|\phi\rangle$. Therefore, our algorithm reduces to estimating the fidelity of $|\psi\rangle$ with each candidate $|\phi\rangle$. We further improve the time efficiency via an algorithm for finding maximal cliques, due to [61], by observing that the candidate subsets must correspond to maximal cliques in the graph of commutation relations.[8] We also improve the sample complexity by using the classical shadows protocol [39] to estimate all of the fidelities with candidate states $|\phi\rangle$ efficiently. For more details on these improvements, see Section 4.2.

Bounded-Distance Stabilizer Approximation. In the case where stabilizer fidelity is bounded below by $\cos^2(\pi/8)$, we follow the same approach, but use a different and more efficient subroutine for determining which of the sampled Weyl operators generate $\mathrm{Weyl}(|\phi\rangle)$. In particular, we show that there is a simple statistical test for this purpose: if $|\langle\phi|\psi\rangle|^2 > \cos^2(\pi/8)$, then for any $x \in \mathbb{F}_2^{2n}$, $x \in \mathrm{Weyl}(|\phi\rangle)$ if and only if $\langle\psi|W_x|\psi\rangle^2 > \frac{1}{2}$ (Corollary 5.4). This allows us to eschew the maximal clique algorithm entirely, and we instead directly estimate $\langle\psi|W_x|\psi\rangle^2$ to determine whether W_x belongs to $\mathrm{Weyl}(|\phi\rangle)$. We further improve upon the sample complexity of this subroutine by making use of an algorithm due to Huang, Kueng, and Preskill [40] for estimating the expectation of m different Weyl operators from only $O(\log m)$ samples.

Symplectic Fourier Analysis. An essential tool for proving the above results is symplectic Fourier analysis, wherein the Fourier transform over real-valued functions is defined with respect to the symplectic product on $\mathbb{F}_2^{2n}$. To give a sense of the usefulness of symplectic Fourier analysis in our work, we showcase two powerful theorems whose proofs are symplectic-Fourier-analytic. In what follows, for a subspace $T \subseteq \mathbb{F}_2^{2n}$ identified with a set of Weyl operators $\{W_x : x \in T\}$, the subspace $T^\perp$ denotes the set of Weyl operators that commute with T.

Theorem 1.5. *Let $T \subseteq \mathbb{F}_2^{2n}$ be a subspace, and let $|\psi\rangle$ be an n-qubit quantum pure state. Then*

$$\sum_{a \in T} p_\psi(a) = \frac{|T|}{2^n} \sum_{x \in T^\perp} p_\psi(x),$$

and

$$\sum_{a \in T} q_\psi(a) = |T| \sum_{x \in T^\perp} p_\psi(x)^2.$$

In words, Theorem 1.5 shows that p_ψ and q_ψ exhibit a strong duality property with respect to the commutation relations among

[7] Here is an easy proof that p_ψ is a distribution: $\sum_x p_\psi(x) = \sum_x 2^{-n} \langle\psi|W_x|\psi\rangle^2 = \mathrm{tr}(|\psi\rangle\langle\psi|\psi\rangle\langle\psi|) = 1$, where the second step follows from using the Weyl expansion of $|\psi\rangle$.

[8] I.e., the graph whose edges connect nodes corresponding to commuting Weyl operators.

Weyl operators. In particular, the first part shows that the mass of p_ψ on a subspace T of Weyl operators is directly proportional to the mass on the subspace $T^\perp$ of Weyl operators that commute with T. Theorem 1.5 is especially powerful when the subspace T is very large, because T and $T^\perp$ always have inversely proportional size (see Fact 2.4). Hence, using our duality theorems, we can convert summations over high-dimensional subspaces into summations over just a few terms.

Full Version. Due to space constraints, proof details and some formal definitions have been omitted. These can be found in the full version of this manuscript [29].

2 PRELIMINARIES

We introduce notation and background that is central to our work. We write $[n] := \{1, \ldots, n\}$. For $x = (a, b) \in \mathbb{F}_2^{2n}$, a and b always denote the first and last n bits of x, respectively. For a probability distribution $\mathcal{D}$ on a set S, we denote drawing a sample $s \in S$ according to $\mathcal{D}$ by $s \sim \mathcal{D}$. We denote drawing a sample $s \in S$ uniformly at random by $s \sim S$. In an undirected graph G, a clique is a complete subgraph of G. A maximal clique is a clique that is not a proper subgraph of another clique. For quantum pure states $|\psi\rangle, |\phi\rangle$, let $d_{\mathrm{tr}}(|\psi\rangle, |\phi\rangle) = \sqrt{1 - |\langle\psi|\phi\rangle|^2}$ denote the trace distance and $F(|\psi\rangle, |\phi\rangle) = |\langle\psi|\phi\rangle|^2$ denote the fidelity.

The n-qubit Pauli group $\mathcal{P}_n$ is the set $\{\pm 1, \pm i\} \times \{I, X, Y, Z\}^{\otimes n}$, where I, X, Y, Z are the standard Pauli matrices. We refer to unitary transformations in the Clifford group as Clifford circuits (equivalently, Clifford circuits are quantum circuits comprised only of Clifford *gates*, namely, the Hadamard, Phase, and CNOT gates). Clifford gates with the addition of any single-qubit non-Clifford gate form a universal gate set. The T-gate is often the non-Clifford gate of choice, where the T-gate is defined by $T := |0\rangle\langle 0| + e^{i\pi/4} |1\rangle\langle 1|$. We denote the set of n-qubit stabilizer states by $\mathcal{S}_n$. One way to measure the "stabilizer complexity" of a quantum state is the stabilizer fidelity.

Definition 2.1 (Stabilizer fidelity, [16, Definition 4])**.** Suppose $|\psi\rangle$ is a pure n-qubit state. The *stabilizer fidelity* of $|\psi\rangle$, denoted $F_{\mathcal{S}}$, is

$$F_{\mathcal{S}}(|\psi\rangle) := \max_{|\phi\rangle \in \mathcal{S}_n} |\langle\phi|\psi\rangle|^2.$$

2.1 Symplectic Vector Spaces

We work extensively with $\mathbb{F}_2^{2n}$ as a symplectic vector space by equipping it with the symplectic product.

Definition 2.2 (Symplectic product)**.** For $x, y \in \mathbb{F}_2^{2n}$, we define the *symplectic product* as $[x, y] = x_1 \cdot y_{n+1} + x_2 \cdot y_{n+2} + \ldots + x_n \cdot y_{2n} + x_{n+1} \cdot y_1 + x_{n+2} \cdot y_2 + \ldots + x_{2n} \cdot y_n$, where all operations are performed over $\mathbb{F}_2$.

The symplectic product gives rise to the notion of a *symplectic complement*, much like the orthogonal complement for the standard inner product.

Definition 2.3 (Symplectic complement)**.** Let $T \subseteq \mathbb{F}_2^{2n}$ be a subspace. The *symplectic complement* of T, denoted by $T^\perp$, is defined by

$$T^\perp := \{a \in \mathbb{F}_2^{2n} : \forall x \in T,\ [x, a] = 0\}.$$

We present the following useful facts about the symplectic complement, many of which are similar to that of the more familiar orthogonal complement.

Fact 2.4. *Let S and T be subspaces of $\mathbb{F}_2^{2n}$. Then:*

- *$T^\perp$ is a subspace.*
- *$(T^\perp)^\perp = T$.*
- *$|T| \cdot |T^\perp| = 4^n$, or equivalently $\dim T + \dim T^\perp = 2n$.*
- *$T \subseteq S \iff S^\perp \subseteq T^\perp$.*

A subspace $T \subset \mathbb{F}_2^{2n}$ is *isotropic* when for all $x, y \in T$, $[x, y] = 0$. A subspace $T \subset \mathbb{F}_2^{2n}$ is *Lagrangian* when $T^\perp = T$. Lagrangian subspaces can equivalently be defined as isotropic subspaces with dimension n.

Many of our proofs (which are omitted from this version) use *symplectic* Fourier analysis, which is similar to Boolean Fourier analysis (see e.g., [55]), except the Fourier characters are defined with respect to the symplectic product.

2.2 Weyl Operators and Bell Difference Sampling

For $x = (a, b) \in \mathbb{F}_2^{2n}$, the *Weyl operator* W_x is defined as

$$W_x := i^{a' \cdot b'} (X^{a_1} Z^{b_1}) \otimes \cdots \otimes (X^{a_n} Z^{b_n}),$$

where $a', b' \in \mathbb{Z}^n$ are the embeddings of a, b into $\mathbb{Z}^n$. Each Weyl operator is a Pauli operator, and every Pauli operator is a Weyl operator up to a phase. Because the Clifford group normalizes the Pauli group, Clifford circuits induce an action on $\mathbb{F}_2^{2n}$ by conjugation of the corresponding Weyl operators (up to phase). That is, for every Clifford circuit C and $x \in \mathbb{F}_2^{2n}$, there exists a unique $y \in \mathbb{F}_2^{2n}$ and phase $\alpha \in \{\pm 1\}$ such that $C W_x C^\dagger = \alpha W_y$. In a slight abuse of notation, we denote this action on $\mathbb{F}_2^{2n}$ by $C(x) = y$.

There is clearly a bijection between $\mathbb{F}_2^{2n}$ and the set of Weyl operators, so any subset of $\mathbb{F}_2^{2n}$ corresponds to a subset of Weyl operators. Importantly, commutation relations between Weyl operators are determined by the symplectic product. In particular, for $x, y \in \mathbb{F}_2^{2n}$, the Weyl operators W_x, W_y commute when $[x, y] = 0$ and anticommute when $[x, y] = 1$. So, if $T \subseteq \mathbb{F}_2^{2n}$ is a subspace, then T is isotropic if and only if $\{W_x : x \in T\}$ is a set of mutually commuting Weyl operators. Similarly, T is Lagrangian if and only if $\{W_x : x \in T\}$ is a set of 2^n mutually commuting Weyl operators.

Definition 2.5 (Unsigned stabilizer group)**.** Let $\mathrm{Weyl}(|\psi\rangle) := \{x \in \mathbb{F}_2^{2n} : W_x |\psi\rangle = \pm |\psi\rangle\}$ denote the unsigned stabilizer group of $|\psi\rangle$.

It is not hard to show that, as a consequence of the uncertainty principle, $\mathrm{Weyl}(|\psi\rangle)$ is an isotropic subspace of $\mathbb{F}_2^{2n}$. Additionally, if $T \subset \mathbb{F}_2^{2n}$ is a Lagrangian subspace, then the set of states $\{|\varphi\rangle : \mathrm{Weyl}(|\varphi\rangle) = T\}$ forms an orthonormal basis of the n-qubit Hilbert space. Moreover, since each basis state $|\varphi\rangle$ is stabilized by 2^n Weyl operators (up to phase), every basis state is a stabilizer state. Conversely, observe that for any stabilizer state $|\varphi\rangle$, $\mathrm{Weyl}(|\varphi\rangle)$ is a Lagrangian subspace.

We now define a new stabilizer complexity measure based on the unsigned stabilizer group.

Definition 2.6 (Stabilizer dimension)**.** Let $|\psi\rangle$ be an n-qubit pure state. The *stabilizer dimension of* $|\psi\rangle$ is the dimension of $\mathrm{Weyl}(|\psi\rangle)$ as a subspace of $\mathbb{F}_2^{2n}$.[9]

The stabilizer dimension of a stabilizer state is n, which is maximal, and, for most states, the stabilizer dimension is 0.

The Weyl operators collectively form an orthogonal basis for $2^n \times 2^n$ matrices with respect to the inner product $\langle A, B\rangle = \mathrm{tr}(A^\dagger B)$. This gives rise to the so-called *Weyl expansion* of a quantum state.

Definition 2.7 (Weyl expansion)**.** Let $|\psi\rangle \in \mathbb{C}^{2^n}$ be an n-qubit quantum pure state. The Weyl expansion of $|\psi\rangle$ is

$$|\psi\rangle\langle\psi| = \frac{1}{\sqrt{2^n}} \sum_{x \in F_2^{2n}} c_\psi(x) W_x,$$

where $c_\psi(x) = \frac{1}{\sqrt{2^n}} \langle\psi|W_x|\psi\rangle$.

Squaring the $c_\psi(x)$'s gives rise to a distribution over $\mathbb{F}_2^{2n}$ and therefore over the Weyl operators (see Footnote 7 for a proof). We denote this distribution by $p_\psi(x) \coloneqq c_\psi(x)^2$ and refer to it as the *characteristic distribution*. Note that, for all x, $p_\psi(x) \in [0, 2^{-n}]$.

A significant algorithmic primitive in our work is *Bell difference sampling* [33, 53]. Let $|\Phi^+\rangle \coloneqq \frac{|00\rangle+|11\rangle}{\sqrt{2}}$. Then, the set of quantum states $\{|W_x\rangle \coloneqq (W_x \otimes I)|\Phi^+\rangle : x \in \mathbb{F}_2^2\}$ forms an orthonormal basis of $\mathbb{C}^2 \otimes \mathbb{C}^2$, which we call the *Bell basis*. Bell difference sampling an n-qubit state $|\psi\rangle$ just means the following. First, take two copies of a pure state $|\psi\rangle$. Take the first qubit in each copy and measure them in the Bell basis. Repeat this for each remaining pair of qubits. Let (a_i, b_i) denote the two-bit measurement outcome from measuring the ith pair of qubits. Then, we denote the measurement outcome on the two copies by $x = (a_1, \ldots, a_n, b_1, \ldots, b_n) \in \mathbb{F}_2^{2n}$. Repeat this once more with two fresh copies of $|\psi\rangle$ to obtain a string $y \in \mathbb{F}_2^{2n}$. Finally, output $x + y$.[10] Historically, Bell difference sampling has found use in algorithms for stabilizer states. However, Gross, Nezami, and Walter proved that Bell difference sampling is meaningful for all quantum states.

Lemma 2.8 (Bell difference sampling, [33, Theorem 3.2])**.** *Let $|\psi\rangle$ be an arbitrary n-qubit pure state. Bell difference sampling corresponds to drawing a sample from the following distribution:*

$$q_\psi(x) \coloneqq 4^n (p_\psi * p_\psi)(x) = \sum_{y \in \mathbb{F}_2^{2n}} p_\psi(y) p_\psi(x+y),$$

and uses four copies of $|\psi\rangle$. We refer to $q_\psi(x)$ as the Weyl distribution.

3 PSEUDORANDOMNESS LOWER BOUNDS

We prove that the output state of any Clifford circuit augmented with fewer than $n/2$ non-Clifford single-qubit gates can be efficiently distinguished from Haar random.[11] As a result, any circuit family that prepares an ensemble of n-qubit pseudorandom quantum states must use at least $\Omega(n)$ non-Clifford single-qubit gates.

[9] The stabilizer dimension is closely related to the stabilizer nullity [12] (in fact, for n-qubit states, the stabilizer dimension is simply n minus the stabilizer nullity).

[10] Even when $|\psi\rangle$ is a stabilizer stabilizer state, measuring two copies of $|\psi\rangle$ in the Bell basis returns $x \in \mathbb{F}_2^{2n}$ with probability $p_\psi(x+a)$, where $a \in \mathbb{F}_2^{2n}$ is an unwanted shift. Bell difference sampling essentially cancels out this unwanted shift a. See [33, 53] for more detail.

[11] If we fix the non-Clifford gate to be a T-gate, then $n/2$ can be improved to n.

The key idea is that Haar-random states have minimal stabilizer dimension (Definition 2.6) with overwhelming probability. By contrast, for a quantum circuit that acts on a stabilizer state (which has stabilizer dimension n), each single-qubit non-Clifford gate decreases the stabilizer dimension by at most 2.

We introduce the following definition to simplify the exposition, borrowing terminology from [51].

Definition 3.1 (t-doped Clifford circuits)**.** A t-doped Clifford circuit is a quantum circuit comprised only of Clifford gates (i.e., Hadamard, Phase, and CNOT) and at most t single-qubit non-Clifford gates that starts in the state $|0^n\rangle$.

3.1 Quantum Circuits With Few Non-Clifford Gates

To begin, we show that the output state $|\psi\rangle$ of a t-doped Clifford circuit, where $t < n/2$, induces a distribution q_ψ that is supported over a subspace of dimension at most $2n - 2$.

Lemma 3.2. *Let $|\psi\rangle$ be the output state of a t-doped Clifford circuit. Then the stabilizer dimension of $|\psi\rangle$ is at least $n - 2t$.*

We remark that the stabilizer dimension lower bound in Lemma 3.2 can be improved to $n - t$ in the case that all of the non-Clifford gates are diagonal (for example, if all of the non-Clifford gates are T-gates). This is because diagonal gates commute with both I *and* Z.

Lemma 3.3. *The supports of p_ψ and q_ψ are contained in* $\mathrm{Weyl}(|\psi\rangle)^\perp$.

Corollary 3.4. *Let $|\psi\rangle$ be the output state of a t-doped Clifford circuit. Then the support of q_ψ is a subspace of dimension at most $n + 2t$.*

3.2 Anticoncentration of Haar-Random States

Now we show that if $|\psi\rangle$ is Haar-random, then q_ψ is well-supported over the entirety of $\mathbb{F}_2^{2n}$ in the sense that every proper subspace of $\mathbb{F}_2^{2n}$ contains a bounded fraction of the q_ψ mass. This implies that sampling from q_ψ gives $2n$ linearly independent elements of $\mathbb{F}_2^{2n}$ after a reasonable number of iterations.

We first require the following lemma, which shows that the Weyl measurements are concentrated around 0. Proved in [30], this is a consequence of Lévy's lemma.

Lemma 3.5 ([30, Corollary 22])**.** *Let $|\psi\rangle$ be a Haar-random n-qubit state. Then*

$$\Pr\left[\exists x \neq 0 : |\langle\psi|W_x|\psi\rangle| \geq \varepsilon\right] \leq 2^{2n+1} \exp\left(-\frac{2^n \varepsilon^2}{36\pi^3}\right).$$

Combining with the fact (Theorem 1.5) that the q_ψ mass on a subspace is proportional to its p_ψ^2 mass on the symplectic complement, we obtain the following.

Lemma 3.6. *Let $|\psi\rangle$ be a Haar-random n-qubit state. Then all subspaces $T \subseteq \mathbb{F}_2^{2n}$ of dimension $2n - 1$ simultaneously satisfy*

$$\sum_{x \in T} q_\psi(x) \leq \frac{2}{3},$$

except with probability at most

$$2^{2n+1}\exp\left(-\frac{2^n}{36\sqrt{3}\pi^3}\right).$$

3.3 Distinguishing From Haar-Random

We are now ready to state and analyze our algorithm that, given copies of $|\psi\rangle$, efficiently distinguishes whether $|\psi\rangle$ is (i) Haar-random or (ii) a state prepared by a $(n/2-1)$-doped Clifford circuit, promised that one of these is the case.

While the analysis is not so trivial, the algorithm itself is straightforward: Bell difference sample $O(n)$ times, and, with high probability, we will have a set of Weyl operators that span $\mathbb{F}_2^{2n}$ when $|\psi\rangle$ is Haar-random. On the other hand, if $|\psi\rangle$ is the output of an t-doped Clifford circuit, for $t < n/2$, this can never happen because q_ψ is supported on a subspace of dimension at most $n+2t$ (which we proved in Corollary 3.4).

Algorithm 1: Distinguishing output of an $(n/2-1)$-doped Clifford circuit from Haar-Random

Input: $24n + 18\log(2/\delta)$ copies of $|\psi\rangle$
Promise: $|\psi\rangle$ is Haar-random or the output of an $(n/2-1)$-doped Clifford circuit
Output: 0 if $|\psi\rangle$ is Haar-random and 1 otherwise, with probability at least $1-\delta$

Let $m = 6n + \frac{9}{2}\log(2/\delta)$
Let $T = \{\}$
repeat m **times**
  Perform Bell difference sampling to obtain $x \in \mathbb{F}_2^{2n}$
  Add x to T
Compute the dimension d of the span of T using Gaussian elimination.
return 0 *if* $d = 2n$ *and* 1 *otherwise.*

To prove the correctness of Algorithm 1, we need the following lemma.

Lemma 3.7. *Let $|\psi\rangle$ be an n-qubit Haar-random quantum state and fix $\delta > 0$. Taking $6n + \frac{9}{2}\log(2/\delta)$ samples from q_ψ suffices to sample $2n$ linearly independent elements of $\mathbb{F}_2^{2n}$ with probability at least $1-\delta$ over both the Haar measure and the sampling process.*

Theorem 3.8. *Algorithm 1 succeeds with probability at least $1-\delta$, and it uses $O(n+\log(1/\delta))$ copies of the input state and $O(n^3 + n^2\log(1/\delta))$ time.*

Our distinguishing algorithm immediately implies a lower bound on the number of non-Clifford gates needed to prepare computationally pseudorandom quantum states.

Corollary 3.9. *Any family of t-doped Clifford circuits that produces an ensemble of n-qubit computationally pseudorandom quantum states must satisfy $t \geq n/2$.*

Note that this lower bound can be improved by a factor of 2 in the special case that all of the non-Clifford gates are diagonal (e.g. T-gates), because of the improved lower bound on stabilizer dimension in Lemma 3.2 for this case.

Corollary 3.10. *Any family of Clifford+T circuits that produces an ensemble of n-qubit computationally pseudorandom quantum states must use at least n T-gates.*

4 STABILIZER STATE APPROXIMATIONS

We state and analyze our algorithm that, given copies of an n-qubit quantum pure state $|\psi\rangle$ that has stabilizer fidelity at least τ, outputs a succinct description of a stabilizer state that witnesses fidelity at least $\tau - \varepsilon$.

Our presentation is split into two parts. First, in Section 4.1, we prove a useful lemma regarding q_ψ on $S^* = \mathrm{Weyl}(|\phi\rangle)$, where $|\phi\rangle$ is the stabilizer state that maximizes stabilizer fidelity with $|\psi\rangle$. At a high level, we argue that any sample from q_ψ has a good enough chance of "making progress" towards learning a complete set of generators for S^*. Formally, we prove that the q_ψ-mass on S^* is not heavily concentrated on proper subspaces of S^*, so that when we sample an element of S^*, we obtain an element of S^* that is linearly independent of the previous samples with a reasonable probability. Second, in Section 4.2, we state our algorithm, prove its correctness, and analyze its sample and time complexities.

4.1 Bell Difference Sampling Makes Progress

The next lemma gives a way to argue that in many of our proofs, we can suppose without loss of generality that $|0^n\rangle$ maximizes stabilizer fidelity.

Lemma 4.1. *Given an n-qubit stabilizer state $|\psi\rangle$, let $S = \mathrm{Weyl}(|\psi\rangle)$ be its unsigned stabilizer group, and let $T \subseteq S$ be a subspace of dimension $n-t$. Then there exists a Clifford circuit C such that $C|\psi\rangle = |0^n\rangle$, $C(S) = 0^n \times \mathbb{F}_2^n$, and $C(T) = 0^{n+t} \times \mathbb{F}_2^{n-t}$.*

Now, we show that the p_ψ-mass on S^* is bounded below by the squared stabilizer fidelity of $|\psi\rangle$.

Lemma 4.2. *Given an n-qubit state $|\psi\rangle$, let $|\phi\rangle$ be a stabilizer state that maximizes the stabilizer fidelity, and let $S^* = \mathrm{Weyl}(|\phi\rangle)$. Then*

$$\sum_{x\in S^*} p_\psi(x) \geq F_{\mathcal{S}}(|\psi\rangle)^2.$$

We can generalize this result to arbitrary subspaces of S^*.

Corollary 4.3. *Given an n-qubit state $|\psi\rangle$, let $|\phi\rangle$ be a stabilizer state that maximizes the stabilizer fidelity, and let $S^* = \mathrm{Weyl}(|\phi\rangle)$. Let $T \subseteq S^*$ be a subspace of S^*. Then*

$$\sum_{x\in T} p_\psi(x) \geq \frac{|T|}{2^n} F_{\mathcal{S}}(|\psi\rangle)^2.$$

Now we show a series of anticoncentration lemmas on proper subspaces of S^*. For these next lemmas, we will find it more convenient to assume without loss of generality (because of Lemma 4.1) that the state maximizing fidelity is $|0^n\rangle$, which conceptually simplifies the computations.

Lemma 4.4. *Let $|\psi\rangle$ be an n-qubit state. Suppose the fidelity $|\langle\psi|\phi\rangle|^2$ is maximized by $|\phi\rangle = |0^n\rangle$ over stabilizer states $|\phi\rangle$. Let $S^* = 0^n \times \mathbb{F}_2^n = \mathrm{Weyl}(|0^n\rangle)$, and let $T = 0^{n+1} \times \mathbb{F}_2^{n-1}$ be a maximal subspace of S^*. Then*

$$\sum_{x\in S^*\setminus T} c_\psi(x) \geq 2^{\frac{n}{2}-1}\left(\sqrt{3}-1\right) F_{\mathcal{S}}(|\psi\rangle).$$

Lemma 4.5. *Given an n-qubit state $|\psi\rangle$, let $|\phi\rangle$ be a stabilizer state that maximizes the stabilizer fidelity, and let $S^* = \mathrm{Weyl}(|\phi\rangle)$. Let $T \subset S^*$ be a proper subspace of S^*. Then*

$$\sum_{x \in S^* \setminus T} q_\psi(x) \geq \frac{2-\sqrt{3}}{2} F_{\mathcal{S}}(|\psi\rangle)^4.$$

4.2 The Algorithm

Our algorithm for stabilizer state approximations uses the powerful classical shadows framework [39] to improve its sample complexity.

Theorem 4.6 (Classical shadows algorithm [39]). *Let ρ be an unknown n-qubit mixed state. Then there exists a quantum algorithm that first performs $m_{\mathsf{shadow}} = O(\log(K/\delta)/\varepsilon^2)$ random Clifford measurements on independent copies of ρ. Then, later given K different observables $O_1, O_2, \ldots, O_K$ in an online fashion, where each O_i is a rank-1 projector, the algorithm uses the measurement results to output estimates $\widehat{o}_1, \ldots, \widehat{o}_K$, such that with probability at least $1-\delta$, for every $i \in [K]$, $|\widehat{o}_i - \mathrm{tr}(O_i\rho)| \leq \varepsilon$. Moreover, if O_i is a projector onto a stabilizer state, then each $\widehat{o}_i$ can be computed from the measurement results by a classical algorithm that takes time $O(n^2 m_{\mathsf{shadow}})$.*

For the "moreover" part of Theorem 4.6, see the remarks on Page 1053 of [39].[12]

We also require an algorithm, due to [61], for computing all of the maximal cliques in a graph.

Theorem 4.7 (Computing maximal cliques [61]). *Given an undirected graph G with n vertices, there is a classical algorithm that outputs a list of all of the maximal cliques in G in time $O(3^{n/3})$.*

Note that this implies that the number of maximal cliques is at most $O(3^{n/3})$.

We are now ready to describe the stabilizer state approximation algorithm. At a high level, it uses Bell difference sampling to obtain a list of candidate Lagrangian subspaces generated by the sampled Weyl operators. Then, it iterates through the candidate groups to find the stabilizer state with largest fidelity, using classical shadows to perform the estimation.

We first argue that with high probability, one of the maximal cliques generates the Lagrangian subspace corresponding to a state that maximizes stabilizer fidelity.

Lemma 4.8. *Given an n-qubit state $|\psi\rangle$, let $|\phi\rangle$ be a stabilizer state that maximizes the stabilizer fidelity, and let $S^* = \mathrm{Weyl}(|\phi\rangle)$. Suppose $|\langle\phi|\psi\rangle|^2 \geq \tau$. Then choosing $m_{\mathsf{clique}} \geq \frac{8+4\sqrt{3}}{\tau^4}(n + \log(1/\delta))$ is sufficient to guarantee that with probability at least $1-\delta$, the Bell difference sampling step of Algorithm 2 samples a complete set of generators for S^*.*

Now we have everything needed to prove the correctness of Algorithm 2.

Theorem 4.9. *Let $|\psi\rangle$ be an n-qubit state with $F_{\mathcal{S}}(|\psi\rangle) \geq \tau$. Then choosing*

$$m_{\mathsf{clique}} = \frac{8+4\sqrt{3}}{\tau^4}(n + \log(2/\delta)) \quad m_{\mathsf{shadow}} = O\left(\frac{n+\log(1/\delta)}{\varepsilon^2\tau^4}\right)$$

[12] This is page 1053 of *Nature Physics* Volume 16. Alternatively, see page 5 of the arXiv version.

Algorithm 2: Stabilizer State Approximation

Input: $m_{\mathsf{shadow}} + 4m_{\mathsf{clique}}$ copies of $|\psi\rangle$
Promise: $|\psi\rangle$ has stabilizer fidelity at least τ
Output: A stabilizer state $|\phi\rangle$ such that $|\langle\phi|\psi\rangle|^2 \geq F_{\mathcal{S}}(|\psi\rangle) - \varepsilon$ with probability at least $1-\delta$

```
Initialize an empty graph G
repeat m_clique times
    Using 4 copies of |ψ⟩, perform Bell difference sampling to obtain x ∈ F_2^{2n}
    Add a vertex for x in G and connect it to all vertices y in G such that [x, y] = 0.
repeat m_shadow times
    Choose a random Clifford circuit U
    Measure U|ψ⟩ in the computational basis and store the result
foreach maximal clique (v_1, ..., v_k) ∈ G computed using Theorem 4.7 do
    Compute S := ⟨v_1, ..., v_k⟩ via Gaussian elimination
    if |S| = 2^n then
        foreach stabilizer state |φ⟩ with Weyl(|φ⟩) = S do
            Let ô_φ be the estimator of |⟨ψ|φ⟩|^2 computed using the algorithm in Theorem 4.6
return whichever |φ⟩ maximizes ô_φ
```

suffices to guarantee that with probability at least $1-\delta$, Algorithm 2 outputs a state $|\phi\rangle$ satisfying $|\langle\phi|\psi\rangle|^2 \geq F_{\mathcal{S}}(|\psi\rangle) - \varepsilon$ and it uses $O\left(\frac{n+\log(1/\delta)}{\varepsilon^2\tau^4}\right)$ samples and $\exp\left(O\left(\frac{n+\log(1/\delta)}{\tau^4}\right)\right) \cdot \frac{1}{\varepsilon^2}$ time.

5 BOUNDED-DISTANCE STABILIZER APPROXIMATION

We present an efficient algorithm that, given copies of a quantum state $|\psi\rangle$ that has fidelity at least $\cos^2(\pi/8)+\gamma$ with a stabilizer state $|\phi\rangle$, outputs $|\phi\rangle$ with high probability. We note that this algorithm solves the same task as Algorithm 2, but with the difference that it only works for $\tau > \cos^2(\pi/8)$.

We start by bounding the squared expectation of Weyl operators in the unsigned stabilizer group of $|\phi\rangle$.

Proposition 5.1. *Let $|\psi\rangle$ be an n-qubit quantum state that has fidelity τ with a stabilizer state $|\phi\rangle$, where $\tau \geq 1/2$. If $x \in \mathrm{Weyl}(|\phi\rangle)$, then $\langle\psi|W_x|\psi\rangle^2 \geq (2\tau-1)^2$.*

If an operator is not in $|\phi\rangle$'s unsigned stabilizer group, then it must anticommute with at least half of the Pauli operators in that group. The uncertainty principle states that the expectation of these operators must be small, since the expectation of the operators in $\mathrm{Weyl}(|\phi\rangle)$ is large. To show this formally, we use the Schrödinger uncertainty relation.

Fact 5.2 (Schrödinger uncertainty relation [7, 59]). *For a quantum state $|\psi\rangle$ and observables A and B,*

$$\left(\langle\psi|A^2|\psi\rangle - \langle\psi|A|\psi\rangle^2\right)\left(\langle\psi|B^2|\psi\rangle - \langle\psi|B|\psi\rangle^2\right) \geq \left|\frac{1}{2}\langle\psi|(AB+BA)|\psi\rangle - \langle\psi|A|\psi\rangle\langle\psi|B|\psi\rangle\right|^2.$$

Proposition 5.3. *Let $|\psi\rangle$ be an n-qubit quantum state that has fidelity τ with a stabilizer state $|\phi\rangle$, where $\tau \geq 1/2$. If $y \notin \mathrm{Weyl}(|\phi\rangle)$, then*

$$\langle\psi|W_y|\psi\rangle^2 \leq 4\tau(1-\tau).$$

Proposition 5.1 and Proposition 5.3 suggest that we can determine whether a given Pauli operator is in the unsigned stabilizer group only from its squared expectation as long as, for all $y \notin \mathrm{Weyl}(|\phi\rangle)$ and for all $x \in \mathrm{Weyl}(|\phi\rangle)$,

$$\langle\psi|W_y|\psi\rangle^2 < \langle\psi|W_x|\psi\rangle^2,$$

which happens only when $4\tau(1-\tau) < (2\tau-1)^2 \iff \cos^2(\pi/8) < \tau$. However, we must also take into account the fact that we cannot know the squared expectations exactly. Rather, we can only recover them to some $\pm O(\gamma)$ accuracy, which in turn implies that τ must be at least $\cos^2(\pi/8)+\gamma$ for some $\gamma > 0$. We formalize this in the following corollary.

Corollary 5.4. *Let $|\psi\rangle$ be an n-qubit quantum state that has fidelity $\cos^2(\pi/8)+\gamma$ with a stabilizer state $|\phi\rangle$ for some $\gamma > 0$. Then for all $y \notin \mathrm{Weyl}(|\phi\rangle)$ and all $x \in \mathrm{Weyl}(|\phi\rangle)$,*

$$\langle\psi|W_x|\psi\rangle^2 \geq \frac{1}{2} + 4\gamma^2 + 2\sqrt{2}\gamma > \frac{1}{2} + 2\sqrt{2}\gamma,$$

and

$$\langle\psi|W_y|\psi\rangle^2 \leq \frac{1}{2} - 4\gamma^2 - 2\sqrt{2}\gamma < \frac{1}{2} - 2\sqrt{2}\gamma.$$

A noteworthy consequence of Corollary 5.4 is that the state $|\phi\rangle$ must be unique:

Corollary 5.5. *If $|\psi\rangle$ has fidelity at least $\cos^2(\pi/8)+\gamma$ with a stabilizer state $|\phi\rangle$ for some $\gamma > 0$, then $|\phi\rangle$ must be unique.*

Observe that the threshold $\cos^2(\pi/8)$ in Corollary 5.5 is tight, because $\cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle$ has fidelity $\cos^2(\pi/8)$ with both $|0\rangle$ and $|+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$.

5.1 The Algorithm

We now state and analyze our algorithm. Corollary 5.4—which is the starting point of our algorithm—implies that, based only on the squared expectation of a Weyl operator, we can decide whether nor not it is in $\mathrm{Weyl}(|\phi\rangle)$, where $|\phi\rangle$ is the stabilizer state that has fidelity at least $\cos^2(\pi/8)+\gamma$ with the input state $|\psi\rangle$. At a high level, there are two missing pieces to complete our algorithm.

First, we need to find a polynomial-size list of Weyl operators that is guaranteed to contain a list of generators of $\mathrm{Weyl}(|\phi\rangle)$. By Lemma 4.8, we can achieve this by Bell difference sampling repeatedly from $\mathrm{Weyl}(|\phi\rangle)$. Second, we must estimate the squared expectations of the Weyl operators we sample. One way to do so is by naïvely measuring each Weyl operator repeatedly, one at a time. However, an algorithm due to Huang, Kueng, and Preskill [40] achieves a better runtime, letting us estimate the squared expectation value of many Weyl operators with only a logarithmic sample complexity.

Theorem 5.6 ([40, Proof of Theorem 4]). *Given any m Weyl operators $P_1, \ldots, P_m$ and copies of an unknown pure state $|\psi\rangle$, there is an algorithm that estimates $\langle\psi|P_i|\psi\rangle^2$ to $\pm\varepsilon$ accuracy with probability at least $1-\delta$ by performing Bell measurements on*

$$\frac{4\log(2m/\delta)}{\varepsilon^2}$$

copies of the unknown state $|\psi\rangle$. The time it takes is

$$O\left(n \cdot m \cdot \frac{\log(m/\delta)}{\varepsilon^2}\right).$$

Putting the pieces together, our algorithm works as follows.

Algorithm 3: Bounded-Distance Stabilizer Approximation

Input: $O\left(n + \frac{\log(n/\delta)}{\gamma^2}\right)$ copies of $|\psi\rangle$

Promise: $|\psi\rangle$ has fidelity at least $\cos^2(\pi/8)+\gamma$ with a stabilizer state $|\phi\rangle$

Output: $|\phi\rangle$ with probability at least $1-\delta$

1. Let $m := \frac{8+4\sqrt{3}}{\cos^8(\pi/8)}(n + \log(3/\delta))$.
2. Bell difference sample m times to obtain $x_1, \ldots, x_m \in \mathbb{F}_2^{2n}$.
3. Using the algorithm in Theorem 5.6, estimate $\langle\psi|W_{x_i}|\psi\rangle^2$ for each i to accuracy $\pm 2\sqrt{2}\gamma$ with failure probability at most $\delta/3$.
4. Discard any x_i's for which the estimate of $\langle\psi|W_{x_i}|\psi\rangle^2$ is less than $\frac{1}{2}$. Let S be the subspace spanned by the remaining samples. If S is not Lagrangian, then output "FAILURE". Otherwise, find a Clifford circuit that measures in the stabilizer basis induced by S.
5. Measure $4\log(3/\delta)$ copies of $|\psi\rangle$ in the stabilizer basis induced by S and output the majority result.

Theorem 5.7. *Let $|\psi\rangle$ be an n-qubit state with fidelity at least $\cos^2(\pi/8)+\gamma$ with a stabilizer state $|\phi\rangle$ for $\gamma > 0$. Given*

$$O\left(n + \frac{\log(n/\delta)}{\gamma^2}\right)$$

copies of $|\psi\rangle$ and

$$O\left(n^3 + \frac{n^2\log(n/\delta) + n\log^2(1/\delta)}{\gamma^2}\right)$$

time, Algorithm 3 outputs $|\phi\rangle$ with probability at least $1-\delta$.

6 TOLERANT PROPERTY TESTING OF STABILIZER STATES

We collect (and give alternative proofs of) a few results related to the property testing algorithm for stabilizer states due to Gross, Nezami, and Walter [33] (hereafter, the "GNW algorithm"). We combine these results with the prior work of [27] to give a *tolerant property testing* algorithm for stabilizer states. By tolerant property testing, we mean that the tester must accept inputs that are ε_1-close to having some property and reject inputs that are ε_2-far from having the same property. This is more general than the standard setting where ε_1 is set to 0.

The following two remarks are important for understanding the extent of our contribution. First, our algorithm is similar to the

algorithm given in [30][13] that distinguishes Haar-random states from quantum states with at least $1/\mathrm{poly}(n)$ stabilizer fidelity. Second, we note that our algorithm only works in certain parameter regimes, not for all sensible settings of ε_1 and ε_2. This is discussed further in Section 6.2.

To explain our property testing model in more detail, we are testing whether or not a quantum state is close to a stabilizer state, where distance is measured with fidelity. Specifically, we are given copies of an n-qubit quantum pure state $|\psi\rangle$ as input, and we must decide whether $F_{\mathcal{S}}(|\psi\rangle) \geq 1-\varepsilon_1$ or $F_{\mathcal{S}}(|\psi\rangle) \leq 1-\varepsilon_2$, promised that one of them is the case.

6.1 The Algorithm

Prior work [30, 33] has established that for all quantum states $|\psi\rangle$,

$$\frac{4\eta-1}{3} \leq F_{\mathcal{S}}(|\psi\rangle) \leq \eta^{1/6}.$$

To simplify notation, let $\alpha_1 := 1-\varepsilon_1$ and $\alpha_2 := 1-\varepsilon_2$. Observe that if $F_{\mathcal{S}}(|\psi\rangle) \geq \alpha_1$ then $\eta \geq \alpha_1^6$, and if $F_{\mathcal{S}}(|\psi\rangle) \leq \alpha_2$ then $\eta \leq \frac{3\alpha_2+1}{4}$. This is the basis of our testing algorithm. Specifically, as long as

$$\alpha_1^6 - \frac{3\alpha_2+1}{4} \geq \frac{1}{\mathrm{poly}(n)},$$

then we can efficiently distinguish the two cases simply by estimating η. For the remainder of this section, define

$$\gamma := \alpha_1^6 - \frac{3\alpha_2+1}{4}.$$

Our algorithm is stated in Algorithm 4.

Algorithm 4: Tolerant Property Testing of Stabilizer States

Input: $48\log(2/\delta)/\gamma^2$ copies of $|\psi\rangle$

Promise: Either case (i): $F_{\mathcal{S}}(|\psi\rangle) \geq \alpha_1$ or case (ii): $F_{\mathcal{S}}(|\psi\rangle) \leq \alpha_2$, for $\alpha_1, \alpha_2 \in [0,1]$ such that $\gamma > 0$

Output: 1 if case (i) holds and 0 if case (ii) holds, with probability at least $1-\delta$

Let $m = \frac{8\log(2/\delta)}{\gamma^2}$.

repeat m **times**

  Perform Bell difference sampling to obtain $W_x \sim q_\psi$.

  Perform the measurement $W_x^{\otimes 2}$ on $|\psi\rangle^{\otimes 2}$. Let $X_i \in \{\pm 1\}$ denote the measurement outcome.

Set $\widehat{\eta} = \frac{1}{m}\sum_i X_i$. Output 1 if $\widehat{\eta} > \alpha_1^6 - \frac{\gamma}{2}$ and 0 otherwise.

Theorem 6.1. *For $\gamma > 0$, Algorithm 4 is correct and it uses $48\log(2/\delta)/\gamma^2$ copies of the input state, $O(n\log(1/\delta)/\gamma^2)$ time, and succeeds with probability at least $1-\delta$.*

6.2 Parameter Regime Discussion

We conclude this section by studying the regime in which our tolerant testing algorithm works in comparison to prior work. The GNW algorithm already implicitly functions as a tolerant property tester: because it uses 6 copies of $|\psi\rangle$ and accepts any stabilizer state with probability 1, if the trace distance between $|\psi\rangle$ and some stabilizer state is at most ε, then the test accepts $|\psi\rangle$ with probability at least $1-6\varepsilon$. We can use this observation to establish the values of α_1 and α_2 in which repeated applications of the GNW algorithm works, given only the soundness analysis of [33] but not applying the completeness bound of [30].

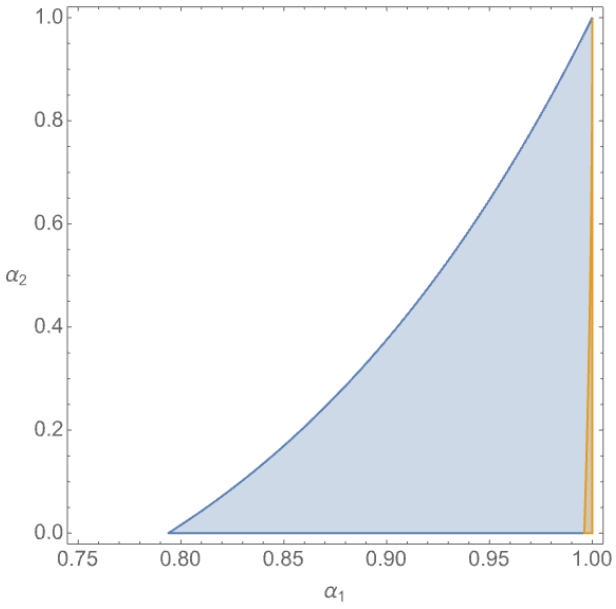

Figure 1: The shaded regions indicate the parameter regimes of α_1 and α_2 that are permissible by the analysis of the GNW algorithm in Section 6.2 (orange) and Algorithm 4 (blue). Thus, the difference between the orange and blue regions illustrates the improvement due to the upper bound $F_{\mathcal{S}}(|\psi\rangle) \leq \eta^{1/6}$ from [30].

Let p_{accept} denote the acceptance probability of the GNW algorithm. It is easy to show that $\eta = 2p_{\mathrm{accept}} - 1$ (see [33, Page 19]). As mentioned above, [33] proved that for any quantum state $|\psi\rangle$, $\frac{4\eta-1}{3} \leq F_{\mathcal{S}}(|\psi\rangle)$. Additionally, since the GNW algorithm uses 6 copies of the input state and accepts stabilizer states with probability 1, it follows that $1-6\sqrt{1-F_{\mathcal{S}}(|\psi\rangle)} \leq p_{\mathrm{accept}}$, where we are using the fact that the trace distance between $|\psi\rangle$ and the stabilizer state maximizing fidelity is $\sqrt{1-F_{\mathcal{S}}(|\psi\rangle)}$. Finally, using the fact that $\eta = 2p_{\mathrm{accept}} - 1$, we get $F_{\mathcal{S}}(|\psi\rangle) \leq \frac{1}{144}(2\eta - \eta^2 + 143)$.

In a "yes" instance, where we are promised that $F_{\mathcal{S}}(|\psi\rangle) \geq \alpha_1$, we have the inequality $\frac{1}{144}(2\eta - \eta^2 + 143) \geq \alpha_1$. Solving for η gives $\eta \geq 1 - 12\sqrt{1-\alpha_1}$. Similarly, in a "no" instance, where we are promised that $F_{\mathcal{S}}(|\psi\rangle) \leq \alpha_2$, we have $\eta \leq (3\alpha_2+1)/4$. Hence, following the same argument as in Section 6.1, the GNW algorithm tolerantly tests stabilizer states as long as:

$$1 - 12\sqrt{1-\alpha_1} > \frac{3\alpha_2+1}{4},$$

whereas, as shown earlier, our algorithm works as long as

$$\alpha_1^6 > \frac{3\alpha_2+1}{4}.$$

This is a significant improvement, which is shown visually in Fig. 1.

[13] Which is itself a repeated application of the base GNW algorithm, for the purposes of error amplification.

7 DISCUSSION AND OPEN PROBLEMS

A natural direction for future work is to improve the performance of our algorithms or to prove (conditional or unconditional) lower bounds. In particular, can the exponential running time of Algorithm 2 be improved upon, or is stabilizer state approximation computationally hard for general parameter regimes? We are optimistic that the exponential factors in our runtime analysis could be made much smaller in practice, because our bound on the sample complexity of finding a complete set of generators is probably far from optimal.

We also remark that, at least superficially, our problem of finding the nearest stabilizer state resembles the closest vector problem (CVP): given a lattice L and a target vector, find the nearest lattice point to the target vector. In our problem, we are given a target vector, and we want to find the nearest stabilizer state to the target vector. While not a lattice, the stabilizer states are "evenly spread" across the complex unit sphere due to their 3-design property [48, 65, 66]. CVP is known to be NP-hard to solve approximately to within any constant and some almost-polynomial factors [9, 22, 63]. Is there a formal connection between these two problems?

Can tighter bounds between η and stabilizer fidelity be proven? In [30, Appendix B], the authors prove that one can hope for at most a roughly quadratic improvement in the bound $F_S(|\psi\rangle)^6 \leq \eta$. In addition to η, are there other statistics related to stabilizer fidelity (or any other stabilizer complexity measure) that can be estimated efficiently? Progress in this direction would extend the parameter regimes for which our property testing algorithm works (see Fig. 1).

One can view the output of Algorithm 2 as an approximation of the input state by a nearby stabilizer state. Following this theme, a natural objective is to design similar approximation algorithms relative to other classes of quantum states such as product states or matchgate states. We note that it is even open to design a time-efficient algorithm that, given copies of an n-qubit quantum state, outputs the nearest state from the set $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n}$, which is a subset of stabilizer states. In addition to potentially improving Clifford+T simulation algorithms (as discussed in Section 1.1), are there other applications for these types of state approximation algorithms?

ACKNOWLEDGMENTS

We thank David Gosset, Srinivasan Arunachalam, Sepehr Nezami, and Arkopal Dutt for helpful conversations. SG, VI, DL are supported via Scott Aaronson by a Vannevar Bush Fellowship from the US Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons "It from Qubit" collaboration. WK is supported by an NDSEG Fellowship, and also acknowledges support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. VI is supported by an NSF Graduate Research Fellowship. DL is also supported by NSF award FET-2243659.

REFERENCES

[1] Scott Aaronson. 2007. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 463, 2088 (2007), 3089–3114. https://doi.org/10.1098/rspa.2007.0113

[2] Scott Aaronson. 2018. Shadow Tomography of Quantum States. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*. Association for Computing Machinery, 325–338. https://doi.org/10.1145/3188745.3188802

[3] Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. 2022. Quantum Pseudoentanglement. arXiv:2211.00747 [quant-ph]

[4] Scott Aaronson and Daniel Gottesman. 2008. Identifying Stabilizer States. https://pirsa.org/08080052.

[5] Scott Aaronson and Sabee Grewal. 2023. Efficient Tomography of Non-Interacting-Fermion States. In *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 266)*. 12:1–12:18. https://doi.org/10.4230/LIPIcs.TQC.2023.12

[6] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. 2022. Cryptography from Pseudorandom Quantum States. In *Advances in Cryptology – CRYPTO 2022*. Springer Nature Switzerland, 208–236. https://doi.org/10.1007/978-3-031-15802-5_8

[7] A. Angelow and M. C. Batoni. 2008. About Heisenberg Uncertainty Relation (by E. Schrodinger). arXiv:quant-ph/9903100 [quant-ph]

[8] Anurag Anshu and Srinivasan Arunachalam. 2023. A survey on the complexity of learning quantum states. arXiv:2305.20069 [quant-ph]

[9] Sanjeev Arora, László Babai, Jacques Stern, and Z Sweedyk. 1997. The Hardness of Approximate Optima in Lattices, Codes, and Systems of Linear Equations. *J. Comput. System Sci.* 54, 2 (1997), 317–331. https://doi.org/10.1006/jcss.1997.1472

[10] Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. 2023. Optimal Algorithms for Learning Quantum Phase States. In *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 266)*. 3:1–3:24. https://doi.org/10.4230/LIPIcs.TQC.2023.3

[11] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. 2021. One-Way Functions Imply Secure Computation in a Quantum World. In *Advances in Cryptology – CRYPTO 2021*, Tal Malkin and Chris Peikert (Eds.). Springer International Publishing, Cham, 467–496. https://doi.org/10.1007/978-3-030-84242-0_17

[12] Michael Beverland, Earl Campbell, Mark Howard, and Vadym Kliuchnikov. 2020. Lower bounds on the non-Clifford resources for quantum computations. *Quantum Science and Technology* 5, 3 (2020), 035009. https://doi.org/10.1088/2058-9565/ab8963

[13] Adam Bouland, Bill Fefferman, and Umesh Vazirani. 2020. Computational Pseudorandomness, the Wormhole Growth Paradox, and Constraints on the AdS/CFT Duality (Abstract). In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 151)*. 63:1–63:2. https://doi.org/10.4230/LIPIcs.ITCS.2020.63

[14] Zvika Brakerski. 2023. Black-Hole Radiation Decoding Is Quantum Cryptography. In *CRYPTO 2023*. Springer Nature Switzerland, 37–65. https://doi.org/10.1007/978-3-031-38554-4_2

[15] Zvika Brakerski and Omri Shmueli. 2019. (Pseudo) Random Quantum States with Binary Phase. In *Theory of Cryptography*. https://doi.org/10.1007/978-3-030-36030-6_10

[16] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. 2019. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum* 3 (2019), 181. https://doi.org/10.22331/q-2019-09-02-181

[17] Sergey Bravyi and Alexei Kitaev. 2005. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* 71 (2005), 022316. Issue 2. https://doi.org/10.1103/PhysRevA.71.022316

[18] Sergey Bravyi, Graeme Smith, and John A. Smolin. 2016. Trading Classical and Quantum Computational Resources. *Physical Review X* 6, 2 (2016), 021043. https://doi.org/10.1103/PhysRevX.6.021043

[19] Costin Bădescu and Ryan O'Donnell. 2021. Improved Quantum Data Analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*. Association for Computing Machinery, 1398–1411. https://doi.org/10.1145/3406325.3451109

[20] A. R. Calderbank and Peter W. Shor. 1996. Good quantum error-correcting codes exist. *Phys. Rev. A* 54 (1996), 1098–1105. Issue 2. https://doi.org/10.1103/PhysRevA.54.1098

[21] Marcus Cramer, Martin B. Plenio, Steven T. Flammia, Rolando Somma, David Gross, Stephen D. Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. 2010. Efficient quantum state tomography. *Nature Communications* 1, 1 (2010), 1–7. https://doi.org/10.1038/ncomms1147

[22] Irit Dinur, Guy Kindler, and Shmuel Safra. 1998. Approximating-CVP to within almost-polynomial factors is NP-hard. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (FOCS '98)*. IEEE, 99–109. https://doi.org/10.1109/SFCS.1998.743433

[23] Zhiyuan Fan, Jiatu Li, and Tianqi Yang. 2022. The Exact Complexity of Pseudorandom Functions and the Black-Box Natural Proof Barrier for Bootstrapping Results in Computational Complexity. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*. 962–975. https://doi.org/10.1145/3519935.3520010

[24] Oded Goldreich, Shari Goldwasser, and Dana Ron. 1998. Property Testing and Its Connection to Learning and Approximation. *Journal of the ACM (JACM)* 45, 4 (1998), 653–750. https://doi.org/10.1145/285055.285060

[25] Aravind Gollakota and Daniel Liang. 2022. On the Hardness of PAC-learning Stabilizer States with Noise. *Quantum* 6 (2022), 640. https://doi.org/10.22331/q-2022-02-02-640

[26] Daniel Gottesman. 1997. *Stabilizer Codes and Quantum Error Correction.* Ph. D. Dissertation. California Institute of Technology. https://doi.org/10.7907/rzr7-dt72

[27] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. 2023. Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates. arXiv:2305.13409 [quant-ph]

[28] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. 2023. Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates II: Single-Copy Measurements. arXiv:2308.07175 [quant-ph]

[29] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. 2023. Improved Stabilizer Estimation via Bell Difference Sampling. arXiv:2304.13915 [quant-ph]

[30] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. 2023. Low-Stabilizer-Complexity Quantum States Are Not Pseudorandom. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 251).* 64:1–64:20. https://doi.org/10.4230/LIPIcs.ITCS.2023.64

[31] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. 2021. Oblivious Transfer Is in MiniQCrypt. In *Advances in Cryptology – EUROCRYPT 2021*, Anne Canteaut and François-Xavier Standaert (Eds.). Springer International Publishing, Cham, 531–561. https://doi.org/10.1007/978-3-030-77886-6_18

[32] David Gross. 2006. Hudson's theorem for finite-dimensional quantum systems. *J. Math. Phys.* 47, 12 (2006), 122107. https://doi.org/10.1063/1.2393152

[33] David Gross, Sepehr Nezami, and Michael Walter. 2021. Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations. *Communications in Mathematical Physics* 385, 3 (2021), 1325–1393. https://doi.org/10.1007/s00220-021-04118-7

[34] Andi Gu, Lorenzo Leone, Soumik Ghosh, Jens Eisert, Susanne Yelin, and Yihui Quek. 2023. A little magic means a lot. arXiv:2308.16228 [quant-ph]

[35] Dominik Hangleiter and Michael J. Gullans. 2023. Bell sampling from quantum circuits. arXiv:2306.00083v1 [quant-ph]

[36] Tobias Haug and M. S. Kim. 2023. Scalable Measures of Magic Resource for Quantum Computers. *PRX Quantum* 4, 1 (2023), 010301. https://doi.org/10.1103/PRXQuantum.4.010301

[37] Minki Hhan, Tomoyuki Morimae, and Takashi Yamakawa. 2023. From the Hardness of Detecting Superpositions to Cryptography: Quantum Public Key Encryption and Commitments. In *EUROCRYPT 2023*. Springer Nature Switzerland, 639–667. https://doi.org/10.1007/978-3-031-30545-0_22

[38] Marcel Hinsche, Marios Ioannou, Alexander Nietner, Jonas Haferkamp, Yihui Quek, Dominik Hangleiter, Jean-Pierre Seifert, Jens Eisert, and Ryan Sweke. 2023. One T Gate Makes Distribution Learning Hard. *Phys. Rev. Lett.* 130 (2023), 240602. Issue 24. https://doi.org/10.1103/PhysRevLett.130.240602

[39] Hsin-Yuan Huang, Richard Kueng, and John Preskill. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics* 16, 10 (2020), 1050–1057. https://doi.org/10.1038/s41567-020-0932-7

[40] Hsin-Yuan Huang, Richard Kueng, and John Preskill. 2021. Information-Theoretic Bounds on Quantum Advantage in Machine Learning. *Physical Review Letters* 126, 19 (2021), 190505. https://doi.org/10.1103/PhysRevLett.126.190505

[41] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. 2008. Cryptography with Constant Computational Overhead. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing (STOC '08).* 433–442. https://doi.org/10.1145/1374376.1374438

[42] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. 2018. Pseudorandom Quantum States. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference*. Springer, 126–152. https://doi.org/10.1007/978-3-319-96878-0_5

[43] Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. 1992. Toward efficient agnostic learning. In *Proceedings of the fifth annual workshop on Computational learning theory.* 341–352. https://doi.org/10.1007/BF00993468

[44] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. 2008. Randomized benchmarking of quantum gates. *Physical Review A* 77, 1 (2008). https://doi.org/10.1103/physreva.77.012307

[45] Lucas Kocia. 2022. Improved Strong Simulation of Universal Quantum Circuits. arXiv:2012.11739 [quant-ph]

[46] William Kretschmer. 2021. Quantum Pseudorandomness and Classical Complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 197).* 2:1–2:20. https://doi.org/10.4230/LIPIcs.TQC.2021.2

[47] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. 2023. Quantum Cryptography in Algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (Orlando, FL, USA) *(STOC 2023).* Association for Computing Machinery, New York, NY, USA, 1589–1602. https://doi.org/10.1145/3564246.3585225

[48] Richard Kueng and David Gross. 2015. Qubit stabilizer states are complex projective 3-designs. arXiv:1510.02767 [quant-ph]

[49] Ching-Yi Lai and Hao-Chung Cheng. 2022. Learning Quantum Circuits of Some T Gates. *IEEE Transactions on Information Theory* 68, 6 (2022), 3951–3964. https://doi.org/10.1109/TIT.2022.3151760

[50] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. 2023. Learning t-doped stabilizer states. arXiv:2305.15398v3 [quant-ph]

[51] Lorenzo Leone, Salvatore F. E. Oliviero, Seth Lloyd, and Alioscia Hamma. 2022. Learning efficient decoders for quasi-chaotic quantum scramblers. arXiv:2212.11338 [quant-ph]

[52] Stéphane Mallat and Zhifeng Zhang. 1994. Matching Pursuit with Time-Frequency Dictionaries. *IEEE Transactions on Signal Processing* 41 (1994), 3397–3415. https://doi.org/10.1109/78.258082

[53] Ashley Montanaro. 2017. Learning stabilizer states by Bell sampling. arXiv:1707.04012 [quant-ph]

[54] Tomoyuki Morimae and Takashi Yamakawa. 2022. Quantum Commitments And Signatures Without One-Way Functions. In *Advances in Cryptology – CRYPTO 2022: 42nd Annual International Cryptology Conference*. 269–295. https://doi.org/10.1007/978-3-031-15802-5_10

[55] Ryan O'Donnell. 2014. *Analysis of Boolean Functions.* Cambridge University Press. https://doi.org/10.1017/CBO9781139814782

[56] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. 2006. Tolerant property testing and distance approximation. *J. Comput. System Sci.* 72, 6 (2006), 1012–1042. https://doi.org/10.1016/j.jcss.2006.03.002

[57] Robert Raussendorf and Hans J. Briegel. 2000. Quantum computing via measurements only. https://doi.org/10.48550/arxiv.quant-ph/0010033

[58] Ronitt Rubinfeld and Madhu Sudan. 1996. Robust Characterizations of Polynomials with Applications to Program Testing. *SIAM J. Comput.* 25, 2 (1996), 252–271. https://doi.org/10.1137/S0097539793255151

[59] Erwin Schrödinger. 1930. The Uncertainty Relations in Quantum Mechanics. Zum Heisenbergschen Unschärfeprinzip. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse* 14 (1930), 296–303.

[60] Peter W. Shor. 1995. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* 52 (1995), R2493–R2496. Issue 4. https://doi.org/10.1103/PhysRevA.52.R2493

[61] Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. 2006. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science* 363, 1 (2006), 28–42. https://doi.org/10.1016/j.tcs.2006.06.015

[62] Leslie G. Valiant. 1984. A Theory of the Learnable. *Commun. ACM* 27, 11 (1984), 1134–1142. https://doi.org/10.1145/1968.1972

[63] Peter van Emde Boas. 1981. Another NP-complete partition problem and the complexity of computing short vectors in lattices. *TR* (1981).

[64] Victor Veitch, S A Hamed Mousavian, Daniel Gottesman, and Joseph Emerson. 2014. The resource theory of stabilizer quantum computation. *New Journal of Physics* 16, 1 (2014), 013009. https://doi.org/10.1088/1367-2630/16/1/013009

[65] Zak Webb. 2016. The Clifford Group Forms a Unitary 3-Design. *Quantum Info. Comput.* 16, 15–16 (2016), 1379–1400. https://doi.org/10.26421/qic16.15-16-8

[66] Huangjun Zhu. 2017. Multiqubit Clifford groups are unitary 3-designs. *Physical Review A* 96, 6 (2017). https://doi.org/10.1103/PhysRevA.96.062336

Received 09-NOV-2023; accepted 2024-02-11