

Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates

Sabee Grewal¹, Vishnu Iyer¹, William Kretschmer^{1,2}, and Daniel Liang^{1,3,4}

¹The University of Texas at Austin

²Simons Institute for the Theory of Computing, UC Berkeley

³Rice University

⁴Portland State University

We give a pair of algorithms that efficiently learn a quantum state prepared by Clifford gates and $O(\log n)$ non-Clifford gates. Specifically, for an n -qubit state $|\psi\rangle$ prepared with at most t non-Clifford gates, our algorithms use $\text{poly}(n, 2^t, 1/\varepsilon)$ time and copies of $|\psi\rangle$ to learn $|\psi\rangle$ to trace distance at most ε .

The first algorithm for this task is more efficient, but requires entangled measurements across two copies of $|\psi\rangle$. The second algorithm uses only single-copy measurements at the cost of polynomial factors in runtime and sample complexity. Our algorithms more generally learn any state with sufficiently large *stabilizer dimension*, where a quantum state has stabilizer dimension k if it is stabilized by an abelian group of 2^k Pauli operators. We also develop an efficient property testing algorithm for stabilizer dimension, which may be of independent interest.

Sabee Grewal: sabee@cs.utexas.edu, <https://sabeegrewal.com/>

Vishnu Iyer: vishnu.iyer@utexas.edu, <https://vishnuiyer.org/>

William Kretschmer: kretsch@cs.utexas.edu, <https://wkretschmer.github.io/>

Daniel Liang: daniel.liang@ll.mit.edu, <https://daniel-you-liang.github.io/>

Contents

1	Introduction	3
1.1	Our Contributions	4
1.2	Main Ideas	4
1.3	Related Work	7
2	Preliminaries	9
2.1	Stabilizer States and Clifford Circuits	11
2.2	Symplectic Vector Spaces	11
2.3	The Weyl Expansion and Bell Difference Sampling	13
2.4	Stabilizer Dimension	15
2.5	Clifford Action on $\mathbb{F}_2^{2n}$	16
3	Linear Algebra Subroutines	16
4	On Subspaces with Large p_ψ or q_ψ-mass	19
4.1	Subspaces With Large Mass Are Isotropic	19
4.2	Product State Structure	20
5	Property Testing Stabilizer Dimension	23
6	Tomography Reduces to Finding Heavy Subspaces	24
6.1	Pure State Tomography	24
6.2	Reduction Algorithm	24
7	Learning the Unsigned Stabilizer Group Using Bell Measurements	27
8	Computational Difference Sampling	29
9	Learning the Unsigned Stabilizer Group Using Single-Copy Measurements	31
9.1	The Output Has Large Dimension	32
9.2	The Output Has Large p_ψ -Mass	34
9.3	Putting It All Together	37
10	Open Problems	39
A	A Faster Algorithm From One Round of Adaptivity	45
B	Generalization to Mixed States	49
B.1	Bell Difference Sampling Mixed States	50
B.2	Reduction to Finding a Heavy Subspace	51
B.3	Efficient Tomography of Large-Stabilizer-Dimension Mixed States	53

1 Introduction

Quantum state tomography is the task of constructing a classical description of a quantum state, given copies of the state. This task—whose study dates back to the 1950s [Fan57]—is fundamentally important in quantum theory, and finds applications in the verification of quantum technologies and in experiments throughout physics, among other things. For a thorough history and motivation, we refer the reader to [MPS03, BCG13].

The optimal number of copies to perform quantum state tomography on a d -dimensional quantum *mixed* state is $\Theta(d^2)$ using entangled measurements [OW16, HHJ⁺17] and $\Theta(d^3)$ using single-copy measurements [KRT17, HHJ⁺17, CHL⁺23]. For a quantum *pure* state, $\Theta(d)$ copies are necessary and sufficient [BM99]. Alas, since the dimension d grows exponentially with the system size, the number of copies consumed by state tomography algorithms quickly becomes impractical, and, indeed, learning systems of even 10 qubits can require millions of measurements [SXL⁺17].

There have been several approaches to circumvent the exponential scaling of quantum state tomography, which we discuss further in Section 1.3. For example, one can try to recover less information about the state, or make additional assumptions about the state. While these results have drastically improved copy complexities relative to general quantum state tomography, many of them remain computationally inefficient.

In this work, we present a pure state tomography algorithm whose copy and time complexities scale in the complexity of a circuit that prepares the state. More specifically, we assume that the circuit is described by a gate set consisting of Clifford gates (i.e., Hadamard, phase, and CNOT gates) as well as single-qubit non-Clifford gates. Such gate sets are well-studied in quantum information because they are universal for quantum computation [Shi03], and have a number of desirable properties for quantum error correction and fault tolerance [Kni04, BK05, BBB⁺23], classical simulation of quantum circuits [BBC⁺19], and efficient implementation of approximate t -designs [HMMH⁺23].

Our main result is a tomography algorithm that scales *polynomially* in the number of qubits and *exponentially* in the number of non-Clifford gates needed to prepare the state.

Theorem 1.1 (Informal version of Corollaries 7.2 and 9.7). *Given $n, t \in \mathbb{N}$, there is an algorithm that uses $\text{poly}(n, 2^t, 1/\varepsilon)$ time and $\text{poly}(n, 1/\varepsilon)$ copies of an n -qubit state $|\psi\rangle$, and learns $|\psi\rangle$ to trace distance ε with high probability, promised that $|\psi\rangle$ can be produced by Clifford gates and at most t single-qubit non-Clifford gates.*

Hence, our algorithm learns in polynomial time any quantum state that can be prepared by Clifford gates and $O(\log n)$ single-qubit non-Clifford gates. This family of states is much more expressive than outputs of Clifford circuits alone (i.e., stabilizer states). For example, these states can form quantum k -designs for any constant k [HMMH⁺23], whereas stabilizer states are not a k -design for any $k > 3$. These states also arise naturally in quantum many-body physics as eigenstates of stabilizer Hamiltonians with a few non-commuting Pauli terms. We refer the reader to [GOL24] for more information regarding the contexts in which these states may appear.

Although our algorithm is no longer efficient when t exceeds $\omega(\log n)$, it still remains more efficient than standard pure state tomography as long as t is asymptotically smaller than n . While the limitation of $O(\log n)$ non-Clifford gates may seem restrictive, Theorem 1.1 is likely tight up to polynomial factors. This is because a polynomial dependence on n and $1/\varepsilon$ is clearly necessary, and an exponential dependence on t is necessary under a plausible cryptographic assumption. In particular, if there exist linear-time constructible pseudorandom quantum states [JLS18] with exponential security, then t -qubit states output by quantum circuits of size $O(t)$ cannot be learned in $2^{o(t)}$ time. Though this is a

strong assumption, an algorithm refuting it would be a major breakthrough in cryptography, because it would also rule out linear-time quantum-secure (weak) pseudorandom functions with exponential security, by the reduction of Brakerski and Shmueli [BS19].

1.1 Our Contributions

We, in fact, give *two* algorithms to solve the learning task [Theorem 1.1](#). The difference between the two algorithms is that one uses entangled measurements across two copies of $|\psi\rangle$, while the other measures only single copies of $|\psi\rangle$ at a time. The algorithm based on entangled measurements (given in [Corollary 7.2](#)) is faster and conceptually more elegant, but it requires the ability to perform a Bell measurement across two copies of $|\psi\rangle$. So, if the copies of $|\psi\rangle$ are provided in an online fashion, then the algorithm needs a coherent quantum memory to store copies between trials.

For the sake of experimental feasibility, it is often preferable to consider quantum algorithms that operate without an external quantum memory—i.e., that can only store classical information after each copy of $|\psi\rangle$ is measured [HKP21, CCHL22, HBC⁺22]. This is how our single-copy algorithm (given in [Corollary 9.7](#)) operates: it uses only unentangled measurements, at the cost of a polynomially worse runtime and sample complexity.

Our algorithms also learn a more general class of states, namely: quantum states with *stabilizer dimension* at least $n - t$ ([Definition 2.22](#)). Informally, a quantum state has stabilizer dimension $n - t$ if it is stabilized by an abelian group of 2^{n-t} Pauli operators.¹ We emphasize that this class of states is *much larger* than the class of states preparable with Clifford gates and t single-qubit non-Clifford gates. Indeed, there are many states with stabilizer dimension $n - t$ that require $2^{\Omega(t)}$ gates from *any* universal gate set to produce (e.g., a state of the form $|0^{n-t}\rangle \otimes |\psi\rangle$, where $|\psi\rangle$ is a Haar-random state on t qubits).

In [Appendix B](#), we generalize our algorithm that uses two-copy measurements to handle mixed states. A mixed state ρ has stabilizer dimension $n - t$ if $\text{tr}(P\rho) = 1$ for every P in an abelian group of 2^{n-t} Pauli operators. Surprisingly, essentially all of our results generalize to mixed states with no loss in parameters.

Finally, one might question whether the learning algorithms presented in this work follow directly from prior results on “near-stabilizer” states, such as those in [GIKL23, GIKL24]. We emphasize that this is not the case. Our algorithms are independent of the main technical results in [GIKL23]. While we do make use of certain results from [GIKL24] (namely, [Proposition 2.17](#), [Fact 2.18](#), [Lemma 2.23](#) in this work), these are not part of the main contributions of either their work or ours.

1.2 Main Ideas

Stabilizer dimension As remarked earlier, our algorithms learn states with stabilizer dimension at least $n - t$ ([Definition 2.22](#)). Quantum states prepared by Clifford gates and $t/2$ non-Clifford gates fall into this class because they have stabilizer dimension at least $n - t$ ([Lemma 2.23](#)).

Our first observation is that learning $|\psi\rangle$ reduces to learning the group G of 2^{n-t} stabilizing Pauli operators of $|\psi\rangle$. In particular, we show in [Lemma 3.2](#) that given a set of generators for G , we can efficiently construct a Clifford circuit C such that $C|\psi\rangle = |\varphi\rangle|x\rangle$, where $|x\rangle$ is a computational basis state on $n - t$ qubits and $|\varphi\rangle$ is a general state on t qubits. This construction builds on standard techniques for manipulating stabilizer tableaux, which

¹Recall that an operator U stabilizes a quantum state $|\psi\rangle$ when $U|\psi\rangle = |\psi\rangle$.

appeared e.g. in the Aaronson-Gottesman algorithm [AG04]. In some sense, this step “compresses” the non-Cliffordness of the state into the first t qubits.² Once we know C , we can easily learn $|x\rangle$ by measuring $C|\psi\rangle$, and can learn $|\varphi\rangle$ using a tomography algorithm on t qubits, which takes $2^{O(t)}$ time [BM99].

Unfortunately, learning G exactly is difficult in general, because there exist states that have no Pauli stabilizers, but that are arbitrarily close to states with a large stabilizer group (for example, consider perturbing a stabilizer state in a random direction). One of our key contributions is to make the above argument based around Lemma 3.2 robust: we show that to learn $|\psi\rangle$, it suffices to merely approximate G by a (possibly different) Pauli subgroup $\widehat{G}$, in a sense of approximation that we will make precise later. Roughly speaking, we require that $\widehat{G}$ be sufficiently large, and that a random element of G approximately stabilizes $|\psi\rangle$ (up to sign) on average:

$$\mathbf{E}_{P \sim G} [|\langle \psi | P | \psi \rangle|^2] \geq 1 - \varepsilon.$$

We show that given a group $\widehat{G}$ of size 2^{n-t} that satisfies this property, Algorithm 2 finds a Clifford circuit that approximately maps $|\psi\rangle$ to a state of the form $|\varphi\rangle|x\rangle$, which allows us to learn $|\psi\rangle$ using only single-copy measurements. To find such a $\widehat{G}$, we give two algorithms: one that uses entangled measurements, and another that uses only single-copy measurements.

Approximating G using Bell measurements Our first algorithm (Algorithm 3) for approximating G utilizes *Bell difference sampling*, a measurement primitive that has seen widespread use in learning algorithms relating to the stabilizer formalism [Mon17, GNW21, GIKL23, GIKL24]. We defer the details of Bell difference sampling to Section 2.3 but note that it involves measuring $|\psi\rangle^{\otimes 2}$ in the Bell basis, repeating twice, and interpreting the result as a Pauli operator. A key property of Bell difference sampling, which is not too hard to show (Fact 2.18), is that the sampled Pauli operators always commute with all elements of G . This suggests a natural approach to try to approximate G : Bell difference sample repeatedly, and then take $\widehat{G}$ to be the commutant of the sampled Pauli operators.

A priori, it is not at all clear why this strategy could work, because in general $\widehat{G}$ may be much larger than G . A key technical step in our proof amounts to showing that, after $\text{poly}(n, 1/\varepsilon)$ Bell difference samples, with high probability, $|\psi\rangle$ must be ε -close to a state that is stabilized by $\widehat{G}$. In other words, if after sufficiently many samples $\widehat{G}$ is larger than G , then this witnesses that $|\psi\rangle$ is close to a state with stabilizer dimension $n - \widehat{t}$ for some $\widehat{t} < t$. So, we can use the aforementioned reduction to perform tomography on $|\psi\rangle$ using $\widehat{G}$, because $\widehat{G}$ must approximately stabilize $|\psi\rangle$.

As a byproduct of this step in our proof, we obtain an algorithm for *property testing* stabilizer dimension, which may be of independent interest.

Theorem 1.2 (Informal version of Theorem 5.1). *There is an algorithm that, given $k \geq 1$ and copies of $|\psi\rangle$, distinguishes the following two cases: either (1) $|\psi\rangle$ has stabilizer dimension at least k , or (2) $|\psi\rangle$ has fidelity at most $1 - \varepsilon$ with all such states. The algorithm uses $O(n/\varepsilon)$ copies of $|\psi\rangle$ and $O(n^3/\varepsilon)$ time.*

Notably, this property testing algorithm is efficient for *all* choices of the stabilizer dimension k , unlike our learning algorithm. Theorem 1.2 thus greatly generalizes a re-

²We are not the first to utilize this type of decomposition; similar techniques of compressing non-Cliffordness have appeared in [ABNO22, LOLH24].

cent result of Grewal, Iyer, Kretschmer, and Liang [GIKL24] that Haar-random states are efficiently distinguishable from states with nonzero stabilizer dimension.

Approximating G using single-copy measurements Conceptually, our algorithm for approximating G via single-copy measurements provides a substantial strengthening and generalization of Aaronson and Gottesman’s algorithm [AG08] for learning stabilizer states from single-copy measurements.³ We start by briefly summarizing their algorithm. First, they select a uniformly random element C from the Clifford group, sample repeatedly from the measurement distribution of $C|\psi\rangle$ in the computational basis, and then iterate this process for many independently chosen C . Aaronson and Gottesman observe that because $C|\psi\rangle$ is a stabilizer state, its measurement distribution is supported uniformly over some affine subspace of $\mathbb{F}_2^n$. For most choices of C , this affine subspace will be all of $\mathbb{F}_2^n$, but with some constant probability, it will be a proper subspace of $\mathbb{F}_2^n$. In the latter case, learning the affine subspace⁴ reveals at least one nontrivial generator of the stabilizer group G of $|\psi\rangle$. For example, if the affine subspace is $0 \times \mathbb{F}_2^{n-1}$ (i.e., the first qubit is always measured to be $|0\rangle$), then this reveals that $C|\psi\rangle$ is stabilized by a Pauli- Z on the first qubit, and therefore $|\psi\rangle$ is stabilized by $C^\dagger(Z \otimes I^{\otimes n-1})C$. Aaronson and Gottesman then show that repeating this process for $O(n)$ randomly chosen C suffices to learn a complete set of generators for G , with high probability.

We find that a similar approach still works to approximately learn G when $|\psi\rangle$ is no longer a stabilizer state, but rather a state with stabilizer dimension $n - t$. Indeed, our algorithm (Algorithm 4) shares similarities with Aaronson and Gottesman’s, but the analysis is considerably more involved. One major difference is that the probability of choosing a Clifford circuit C that lets us learn a generator of G (i.e., for which the support of $C|\psi\rangle$ is a proper affine subspace) will be smaller, on the order of 2^{-t} (Lemma 9.1), and so we need to sample more random Clifford circuits to learn a complete set of generators for the stabilizer group. Furthermore, we require a significantly more careful error analysis, because the measurement distribution of $C|\psi\rangle$ need not be supported uniformly over an affine subspace for non-stabilizer states $|\psi\rangle$, and so it may be impossible to learn the full support of the distribution from a reasonable number of samples. We circumvent this issue in a similar fashion to the Bell difference sampling-based approach, by arguing that it suffices to merely *approximate* the support of the distribution to within some inverse-polynomial error.

There is another minor difference between our approach and that of Aaronson and Gottesman: we use a modified sampling technique that we call *computational difference sampling*, wherein we measure $C|\psi\rangle$ twice in the computational basis and output the difference (i.e., bitwise XOR) of the sampled strings. Taking this difference has the effect of canceling the shift on the affine subspace over which the measurement distribution of $C|\psi\rangle$ is supported, leaving a distribution over a (non-affine) subspace of $\mathbb{F}_2^n$.⁵ We then interpret the sampled string $x \in \mathbb{F}_2^n$ as corresponding to the Pauli- X -operator

$$\xi := \bigotimes_{i=1}^n X^{x_i}.$$

³Aaronson and Gottesman’s algorithm was presented in a short video lecture [AG08], but was never published. A brief sketch of the algorithm and a partial derivation of its workings is presented in Aaronson’s lecture notes [Aar22]. Part of our work’s contribution is to place their algorithm on more rigorous footing.

⁴I.e., computing from samples a basis that spans the space.

⁵Note that the Bell sampling approach also involves taking two samples and computing a difference. This is for a similar reason of canceling an unwanted affine shift [Mon17, GNW21, LC22].

We choose this convention so that the samples from computational difference sampling commute with all of the Pauli- Z stabilizers, as we prove in [Corollary 8.6](#).

Unsigned Pauli operators Throughout our proofs, we will actually find it more convenient to work with *unsigned* Pauli operators, meaning that we view elements of the Pauli group as operators modulo phase. In doing so, we can identify the n -qubit Pauli group under multiplication (modulo phase) with the vector space $\mathbb{F}_2^{2n}$ under addition, which simplifies many computations. In particular, this identification reduces many subroutines of our algorithms (e.g., computing the commutant of a Pauli subgroup) into straightforward linear-algebraic computations over $\mathbb{F}_2$, often by way of viewing $\mathbb{F}_2^{2n}$ as a symplectic vector space. See [Section 2.2](#) and [Section 2.3](#) for more details on the relation between Pauli operators and $\mathbb{F}_2^{2n}$.

Summary To recap, the steps in our learning algorithms are as follows: (1) use [Algorithm 3](#) or [Algorithm 4](#) to learn a Pauli subgroup $\widehat{G}$ that approximates the stabilizer group of $|\psi\rangle$, then (2) use [Algorithm 2](#) to compute from $\widehat{G}$ a Clifford circuit C such that $C|\psi\rangle \approx |\varphi\rangle|x\rangle$, and finally learn $|\varphi\rangle|x\rangle$. While some aspects of the analysis are technical, the algorithms themselves are quite simple and could be amenable to implementation on near-term devices. Indeed, the only quantum parts of the algorithms involve measuring pairs of qubits, applying Clifford circuits, measuring in the computational basis, and performing tomography on a t -qubit state. So, for example, the resource requirements of our algorithms are quite comparable to those of the classical shadows protocol [\[HKP20\]](#).

1.3 Related Work

There is a long line of work devoted to developing near-Clifford *simulation* algorithms [\[AG04, BG16, RLCK19, BBC⁺19, QPG21\]](#), which classically simulate quantum circuits dominated by Clifford gates. These algorithms scale polynomially in the number of qubits and Clifford gates and exponentially in the number of non-Clifford gates. The main contribution of this work is to complement these classical simulation algorithms with learning algorithms that scale comparably with respect to the number of non-Clifford gates.

There are a few other classes of quantum states for which time-efficient tomography algorithms are known. Among these are stabilizer states [\[AG08, Mon17\]](#), non-interacting fermion states [\[AG23\]](#), matrix product states [\[CPF⁺10\]](#), and certain classes of phase states [\[ABDY23\]](#). As a result of our work, the class of quantum states prepared by Clifford gates and $O(\log n)$ non-Clifford gates joins this list. We note that our results strictly generalize both Aaronson and Gottesman’s algorithm [\[AG08\]](#) and Montanaro’s algorithm [\[Mon17\]](#) for learning stabilizer states.

Lai and Cheng [\[LC22\]](#) gave an algorithm that learns a quantum state that is prepared via Clifford gates and a few T -gates, where the T -gate is the non-Clifford unitary $T = |0\rangle\langle 0| + e^{i\pi/4}|1\rangle\langle 1|$. However, their algorithm only works if the circuit U that prepares the input state meets two conditions. Firstly, U must be written as $C_1 T^v C_2$, where C_1 and C_2 are Clifford circuits, and $T^v = T^{v_1} \otimes \cdots \otimes T^{v_n}$ for a string $v \in \mathbb{F}_2^n$ of Hamming weight $O(\log n)$ (i.e., there is a single layer of $O(\log n)$ T -gates between two Clifford circuits). Secondly, the X -matrix of the stabilizer tableau of $C_1|0^n\rangle$ must be full rank, but only when restricted to the qubits that the T -gates act on. Suffice it to say, their algorithm works in a highly restricted setting. In contrast, our algorithms work for *any* quantum state prepared using Clifford gates and $O(\log n)$ arbitrary non-Clifford gates (not just the

T -gate), and with the non-Clifford gates allowed to be placed anywhere in the circuit. Our work is, therefore, a substantial generalization of Lai and Cheng’s algorithm.⁶

Besides Lai and Cheng [LC22], other authors have explored the complexity of some related learning problems that involve Clifford+ T circuits. For example, [HIN⁺23] observes that, given samples from the measurement distribution of a circuit comprised of Clifford gates and a single T -gate, learning this distribution can be as hard as the learning parities with noise (LPN) problem. Leone, Oliviero, Lloyd, and Hamma [LOLH24] give algorithms for learning the dynamics of a quantum circuit U comprised of Clifford gates and few T -gates, given oracle access to U . Both of these results are formally incomparable to our algorithm because the inputs and outputs of the respective learning tasks are different from ours. Nevertheless, the learning algorithm of [LOLH24] utilizes some similar techniques that involve compressing non-Cliffordness.

In another direction, one can reduce the computational complexity of learning by only estimating certain properties of quantum states, instead of producing an entire description of the state. For example, consider the *shadow tomography* problem [Aar20, AR19, BO21] where, given a list of known two-outcome observables and copies of an unknown quantum state, the goal is to estimate the expectation value of each observable with respect to the unknown state. Aaronson [Aar20] showed that shadow tomography requires a number of copies that scales polylogarithmically in both the number of two-outcome observables and the Hilbert space dimension, but the algorithm is not computationally efficient. More recently, Huang, Kueng, and Preskill [HKP20] introduced *classical shadows*, a shadow tomography algorithm that could be amenable to near-term quantum devices. Just like prior work, classical shadows uses exponentially fewer copies of the input state relative to state tomography, but, in general, is not computationally efficient.⁷

Finally, an important primitive in our work is Bell difference sampling, which was introduced (implicitly) by Montanaro [Mon17] and then further developed by Gross, Nezami, and Walter [GNW21]. In recent years, this technique has been at the core of several results involving stabilizer states and near-stabilizer states. For example, Bell difference sampling has been used to learn [Mon17] and property test [GNW21, GIKL24] stabilizer states. Bell difference sampling has also been used to prove computational lower bounds for constructing pseudorandom quantum states [GIKL23, GIKL24]. As mentioned earlier, our property testing algorithm (Theorem 1.2) subsumes the result from [GIKL24] that quantum states with nontrivial stabilizer dimension are not computationally pseudorandom.

Concurrent Work Shortly after the first version of this manuscript was posted, two works [LOH24, HG24] appeared that claim results similar to our Bell sampling-based learning algorithm.

Leone, Oliviero, and Hamma [LOH24] give an algorithm for learning output states of Clifford+ T circuits, whose time complexity scales polynomially in the number of qubits and exponentially in the number of T -gates. However, our result is stronger than [LOH24] for two reasons. Firstly, our algorithm allows for arbitrary non-Clifford gates (not just T -gates). In fact, our algorithm also efficiently learns any state that is stabilized by a sufficiently large Pauli subgroup, which is a more general class of states. By contrast, as

⁶We also believe there are some fixable errors in the most recent version of Lai and Cheng’s algorithm. For example, using notation that we introduce later in Section 2.3, [LC22] assumes that Bell difference sampling draws from the distribution that we call p_ψ . However, Bell difference sampling actually draws from a different distribution that we call q_ψ , which is the convolution of p_ψ with itself.

⁷There are certain settings where the classical shadows protocol is computationally efficient; see [HKP20] for more detail.

stated, [LOH24] relies on a specific algebraic property of the T -gate and does not readily extend in this fashion. Secondly, [LOH24]’s algorithm scales as $\text{poly}(n) \cdot \exp(t)$ where t is the number of T gates, whereas Corollary 7.2 scales as $\text{poly}(n) + \exp(t)$, which gives a polynomial in n savings for sufficiently large values of t .

Hangleiter and Gullans [HG24] exhibit a variety of protocols involving Bell sampling that have applications for learning and error detection. One of their results is an algorithm for learning outputs of Clifford+ T circuits that is similar to our Algorithms 2 and 3. While they only discuss T -gates, we believe that their algorithm, like ours, can also learn arbitrary states stabilized by a large group of Pauli matrices. However, [HG24] lacks a thorough time complexity analysis, and largely focuses on the sample complexity. We give a more comprehensive and careful analysis of the samples and time used by our algorithms. For example, Steps 2 and 3 of [HG24]’s algorithm corresponds to our Lemmas 3.1 and 3.2 respectively, where we give a detailed description of the computations involved and bound the runtime explicitly. Additionally, our bounds on sample complexity are tighter (e.g., compare the bound $O\left(\frac{n \log(1/\delta)}{\epsilon}\right)$ in [HG24, Lemma 4] and $O\left(\frac{n \log(n) \log(1/\delta)}{\epsilon}\right)$ in [HG24, Lemma 5] to our bound $O\left(\frac{n + \log(1/\delta)}{\epsilon}\right)$ in Lemma 2.3).

A later work by Chia, Lai, and Lin [CLL24] gives a single-copy learning algorithm analogous to our Algorithm 4, and appeared concurrently with the update to this manuscript to add Algorithm 4. It uses a similar technique of learning generators of the stabilizer group of $|\psi\rangle$ by repeatedly repeatedly in random Clifford bases. One difference between the two approaches is the use of adaptive measurements in [CLL24]’s algorithm (i.e., measurements that depend on the outcomes of measurements taken in a previous step). By contrast, our Algorithm 4 is non-adaptive and quartically faster in ϵ^{-1} dependence. However, it is slower by a factor of roughly n in both the sample and time complexity. Nevertheless, we show in Appendix A that a slight modification of Algorithm 4 with one additional round of adaptive measurements (inspired by the approach of [CLL24]) lets us remove this extra factor of n and get the best of both runtime dependencies.

2 Preliminaries

For a set of vectors $\{v_1, \dots, v_m\}$, we denote their span by $\langle v_1, \dots, v_m \rangle$. For a positive integer n , define $[n] := \{1, \dots, n\}$. We use $\log$ to denote the natural logarithm. For a probability distribution $\mathcal{D}$ on a set S , we denote drawing a sample $s \in S$ according to $\mathcal{D}$ by $s \sim \mathcal{D}$. For quantum pure states $|\psi\rangle, |\phi\rangle$, let $d_{\text{tr}}(|\psi\rangle, |\phi\rangle) = \sqrt{1 - |\langle\psi|\phi\rangle|^2}$ denote the trace distance and $F(|\psi\rangle, |\phi\rangle) = |\langle\psi|\phi\rangle|^2$ denote the fidelity. As a shorthand, we write the tensor product of pure states $|\psi\rangle \otimes |\varphi\rangle$ as $|\psi\rangle |\varphi\rangle$. We also use the following concentration inequalities.

Fact 2.1 (Chernoff bound). *Let $X_1, \dots, X_n$ be independent identically distributed random variables taking values in $\{0, 1\}$. Let X denote their sum and let $\mu = \mathbf{E}[X]$. Then for any $\delta > 0$,*

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\delta^2 \mu/2}.$$

Fact 2.2 (Hoeffding’s inequality). *Let $X_1, \dots, X_n$ be independent identically distributed random variables subject to $a_i \leq X_i \leq b_i$ for all i . Let $X = \sum_{i=1}^n X_i$ denote their sum and let $\mu = \mathbf{E}[X]$. Then for any $t \geq 0$,*

$$\Pr[X - \mu \geq t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right)$$

and

$$\Pr[|X - \mu| \geq t] \leq 2 \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

We require the following lemma, similar forms of which appeared in other works related to learning near-stabilizer states, including [GIKL24, Lemma 5.9] and [HG24, Lemma 5].

Lemma 2.3. *Let $\mathcal{D}$ be a distribution over $\mathbb{F}_2^d$, let $\varepsilon, \delta \in (0, 1)$, and suppose*

$$m \geq \frac{2 \log(1/\delta) + 2d}{\varepsilon}.$$

Let $A \subseteq \mathbb{F}_2^d$ be the subspace spanned by m independent samples drawn from $\mathcal{D}$. Then

$$\sum_{x \in A} \mathcal{D}(x) \geq 1 - \varepsilon$$

with probability at least $1 - \delta$.

Proof. For samples $x_1, \dots, x_m \in \mathbb{F}_2^d$ drawn from $\mathcal{D}$, let $A_i = \langle x_1, \dots, x_i \rangle$ be the subspace spanned by the first i samples for arbitrary $0 \leq i \leq m$, with the convention that A_0 is the trivial subspace. Define the indicator random variable X_i as

$$X_i = \begin{cases} 1 & \text{if } x_i \in \mathbb{F}_2^d \setminus A_{i-1} \text{ or } \sum_{x \in A_{i-1}} \mathcal{D}(x) \geq 1 - \varepsilon. \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively, $X_i = 1$ indicates that either the sample x_i has increased the span of A_i (i.e., $A_{i-1} \subset A_i$) or that A_{i-1} already accounts for a $1 - \varepsilon$ fraction of the mass of $\mathcal{D}$. If A_{i-1} does account for a $1 - \varepsilon$ fraction of the mass of $\mathcal{D}$, then $X_i = 1$ with probability 1. If not, then the probability mass on $\mathbb{F}_2^d \setminus A_{i-1}$ is at least ε , and therefore, $X_i = 1$ with probability at least ε . In both cases, $\Pr[X_i = 1] \geq \varepsilon$, and therefore $\mu := \mathbf{E}[\sum_i X_i] \geq m\varepsilon$.⁸

Once $\sum_{i=1}^m X_i \geq d$, the subspace A_m must account for a $1 - \varepsilon$ fraction of the mass of $\mathcal{D}$. This is because the dimension of $\mathbb{F}_2^d$ is d and so the span can only expand d times. Set $\gamma := 1 - \frac{d}{\mu}$. By a Chernoff bound (Fact 2.1),

$$\begin{aligned} \Pr\left[\sum_{i=1}^m X_i < d\right] &= \Pr\left[\sum_{i=1}^m X_i < (1 - \gamma)\mu\right] \\ &\leq \exp\left(-\frac{\mu}{2}\gamma^2\right) \\ &= \exp\left(-\frac{\mu}{2} - \frac{d^2}{\mu} + d\right) \\ &\leq \exp\left(-\frac{m\varepsilon}{2} + d\right). \end{aligned}$$

Hence, as long as

$$m \geq \frac{2 \log(1/\delta) + 2d}{\varepsilon},$$

A_m will account for a $1 - \varepsilon$ fraction of the $\mathcal{D}$ -mass with probability at least $1 - \delta$. □

⁸A careful observer will note that, since the X_i are dependent on previous samples, the Chernoff bound does not necessarily hold. However, we can replace each X_i with the random variable Y_i such that $\{Y_i\}$ consists of i.i.d Bernoulli random variables with $\Pr[Y_i = 1] = \varepsilon$. Since $\Pr[\sum_i Y_i < d] \geq \Pr[\sum_i X_i < d]$, we can simply apply the Chernoff bound to $\sum_i Y_i$ instead. Put another way, the dependence *only helps*.

2.1 Stabilizer States and Clifford Circuits

The n -qubit Pauli group is the set $\{\pm 1, \pm i\} \times \{I, X, Y, Z\}^{\otimes n}$. We refer to Hadamard, phase, and CNOT gates as Clifford gates, where $|0\rangle\langle 0| + i|1\rangle\langle 1|$ is the phase gate and CNOT is the controlled-not gate. The Clifford group is the group of unitary transformations comprised only of Clifford gates, and we refer to a unitary transformation in the Clifford group as a Clifford circuit. We note that a uniformly random element of the Clifford group can be sampled efficiently, by the works of [KS14, BM21, VDB21].

Proposition 2.4 ([VDB21]). *There is a classical algorithm that samples a uniformly random element of the n -qubit Clifford group and outputs a Clifford circuit implementation in time $O(n^2)$.*

A unitary transformation U stabilizes a state $|\psi\rangle$ when $U|\psi\rangle = |\psi\rangle$. A stabilizer state is a quantum state that is stabilized by an abelian group of 2^n Pauli operators, or equivalently, any quantum state reachable from $|0^n\rangle$ by applying a Clifford circuit. For more background on the stabilizer formalism, see, e.g., [Got97, NC02].

Any gate set comprised of Clifford gates and just one single-qubit non-Clifford gate is universal for quantum computation [Shi03]. We introduce the following definition for convenience, borrowing terminology from [LOLH24].

Definition 2.5 (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.

2.2 Symplectic Vector Spaces

Throughout this work, $\mathbb{F}_2^{2n}$ is equipped with a canonical symplectic bilinear form that we term the symplectic product.

Definition 2.6 (Symplectic product). For $x, y \in \mathbb{F}_2^{2n}$, we define the *symplectic product* as $[x, y] = x_1 \cdot y_{n+1} + x_2 \cdot y_{n+2} + \dots + x_n \cdot y_{2n} + x_{n+1} \cdot y_1 + x_{n+2} \cdot y_2 + \dots + x_{2n} \cdot y_n$, where all operations are performed over $\mathbb{F}_2$.

Similar to the orthogonal complement for the standard inner product, the symplectic product gives rise to a *symplectic complement*.

Definition 2.7 (Symplectic complement). Let $A \subseteq \mathbb{F}_2^{2n}$ be a subspace. The *symplectic complement* of A , denoted by $A^\perp$, is defined by

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

A subspace $A \subset \mathbb{F}_2^{2n}$ is *isotropic* when for all $x, y \in A$, $[x, y] = 0$. A subspace $A \subset \mathbb{F}_2^{2n}$ is *coisotropic* when $A^\perp$ is isotropic.

To aid intuition, we present the following facts about the symplectic complement, many of which are similar to the more familiar orthogonal complement.

Fact 2.8. *Let A and B be subspaces of $\mathbb{F}_2^{2n}$. Then:*

- $A^\perp$ is a subspace.
- $(A^\perp)^\perp = A$.
- $|A| \cdot |A^\perp| = 4^n$, or equivalently $\dim A + \dim A^\perp = 2n$.

- $A \subseteq B \iff B^\perp \subseteq A^\perp$.

Next, we introduce subspace addition and show how it behaves with respect to other operations on subspaces (namely, subspace intersection and the symplectic complement).

Definition 2.9 (Subspace addition). Given subspaces A and B over some vector space $\mathcal{V}$, let $A + B := \{a + b : a \in A \text{ and } b \in B\}$.

Fact 2.10. *Given subspaces $A \subseteq \mathcal{V}$ and $B \subseteq \mathcal{V}$ over some vector space $\mathcal{V}$, then the following are true:*

- $A + B = B + A$.
- $A + B$ is a subspace with dimension $\dim(A) + \dim(B) - \dim(A \cap B)$.
- $A \subseteq A + B$
- $B \subseteq A + B$

In the remainder of this section, we prove several useful identities that relate subspace addition to subspace intersection and the symplectic complement of a subspace.

Lemma 2.11. *Let A and B be subspaces of $\mathbb{F}_2^{2n}$. Then*

$$(A + B)^\perp = A^\perp \cap B^\perp.$$

Proof. $A^\perp \cap B^\perp \subseteq (A + B)^\perp$, since if $x \in \mathbb{F}_2^{2n}$ has $[x, y] = 0$ for every $y \in A + B$, then x must also have symplectic product zero with all subsets of $A + B$ (note that A and B are subsets of $A + B$ by [Fact 2.10](#)). Conversely, if $y \in \mathbb{F}_2^{2n}$ satisfies $[y, a] = 0$ and $[y, b] = 0$ for every $a \in A$ and $b \in B$ then, by linearity of the symplectic product, $[y, a + b] = 0$. Thus the two sets are equal. $\square$

Lemma 2.12. *Given subspaces A , B , and C over $\mathbb{F}_2^{2n}$ such that $C \subseteq A$ then,*

$$A \cap (B + C) = (A \cap B) + C$$

Proof. Observe that $A \cap (B + C) = \{b + c \in \mathbb{F}_2^{2n} : b \in B \wedge c \in C \wedge b + c \in A\}$. By assumption $C \subseteq A$, so $c \in C$ implies $c \in A$. Therefore, we can conclude that $b + c \in A$ if and only if $b \in A$. This is because for any $a \in A$, $a + b \in A$ if and only if $b \in A$, by the closure of subspaces under addition. Therefore $\{b + c \in \mathbb{F}_2^{2n} : b \in B \wedge c \in C \wedge b + c \in A\}$ can equivalently be written as $\{b + c \in \mathbb{F}_2^{2n} : b \in A \wedge b \in B \wedge c \in C\}$, which precisely the set $(A \cap B) + C$. $\square$

Corollary 2.13. *Let A and B be subspaces of $\mathbb{F}_2^{2n}$, where A is isotropic. Then*

$$(A^\perp \cap (A + B))^\perp = A^\perp \cap (A + B^\perp).$$

Proof. We note that if A is isotropic, then, by definition, $A \subseteq A^\perp$.

$$\begin{aligned} (A^\perp \cap (A + B))^\perp &= A + (A + B)^\perp && \text{(Lemma 2.11)} \\ &= A + (A^\perp \cap B^\perp) && \text{(Lemma 2.11)} \\ &= A^\perp \cap (A + B^\perp) && \text{(Lemma 2.12)} \end{aligned} \quad \square$$

Corollary 2.14. *Let A and B be subspaces of $\mathbb{F}_2^{2n}$, where A is isotropic and B is Lagrangian. Then*

$$A^\perp \cap (A + B)$$

is Lagrangian.

Proof. By [Corollary 2.13](#), we see that $A^\perp \cap (A + B)$ is equal to its symplectic complement if B is Lagrangian as well. $\square$

2.3 The Weyl Expansion and Bell Difference Sampling

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

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

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. We will often freely go back and forth between Weyl operators and $\mathbb{F}_2^{2n}$.

We denote by $\mathcal{X} := \mathbb{F}_2^n \times 0^n$ and $\mathcal{Z} := 0^n \times \mathbb{F}_2^n$ the Weyl operators corresponding to Pauli- X and Z strings, respectively.

We use the following notation for the stabilizer group of a quantum state, ignoring global phase.

Definition 2.15 (Unsigned stabilizer group). Let $\text{Weyl}(|\psi\rangle) := \{x \in \mathbb{F}_2^{2n} : W_x |\psi\rangle = \pm |\psi\rangle\}$ denote the unsigned stabilizer group of $|\psi\rangle$.

The symplectic product determines commutation relations between Weyl operators. Specifically, the Weyl operators W_x and W_y commute when $[x, y] = 0$, and anticommute when $[x, y] = 1$, where x and y are elements of $\mathbb{F}_2^{2n}$. If T is a subspace of $\mathbb{F}_2^{2n}$, then it is isotropic if and only if $\{W_x : x \in T\}$ is a set of mutually commuting Weyl operators. Furthermore, it is not difficult to show that $\text{Weyl}(|\psi\rangle)$ is always an isotropic subspace of $\mathbb{F}_2^{2n}$ (see [Fact 2.21](#)).

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

Definition 2.16 (Weyl expansion). Let $|\psi\rangle \in \mathbb{C}^{2^n}$ be an n -qubit quantum pure state. The Weyl expansion of $|\psi\rangle$ is

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

Using the Weyl expansion of $|\psi\rangle$ and the orthogonality of the Weyl operators, observe that

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

and therefore the squared coefficients of the Weyl expansion can be normalized to form a distribution over $\mathbb{F}_2^{2n}$. We denote this distribution by $p_\psi(x) := 2^{-n} \langle\psi|W_x|\psi\rangle^2$ and refer to it as the *characteristic distribution*. Gross, Nezami, and Walter [\[GNW21\]](#) introduced a measurement primitive called *Bell difference sampling*, which takes four copies of a quantum state as input and outputs a sample $x \in \mathbb{F}_2^{2n}$ according to the distribution

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

which is the convolution of p_ψ with itself.

Grewal, Iyer, Kretschmer, and Liang [GIKL24] proved additional properties of p_ψ and q_ψ that are needed in this work. First, they showed that p_ψ and q_ψ exhibit a strong duality property with respect to the commutation relations among Weyl operators. The proof involves symplectic Fourier analysis so the details are omitted.

Proposition 2.17 ([GIKL24, Theorem 3.1, 3.2]). *Let $A \subseteq \mathbb{F}_2^{2n}$ be a subspace, and let $|\psi\rangle$ be an n -qubit quantum pure state. Then*

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

and

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

They also prove that the support of these distributions is contained in the symplectic complement of the unsigned stabilizer group.

Fact 2.18 ([GIKL24, Lemma 4.3, 4.4]). *The support of q_ψ is contained in $\text{Weyl}(|\psi\rangle)^\perp$.*

Proof.

$$\begin{aligned} \sum_{x \in \text{Weyl}(|\psi\rangle)^\perp} q_\psi(x) &= |\text{Weyl}(|\psi\rangle)^\perp| \sum_{x \in \text{Weyl}(|\psi\rangle)} p_\psi(x)^2 && \text{(Proposition 2.17)} \\ &= |\text{Weyl}(|\psi\rangle)^\perp| \frac{|\text{Weyl}(|\psi\rangle)|}{4^n} && \text{(By definition of } \text{Weyl}(|\psi\rangle)) \\ &= 1. && \text{(Fact 2.8)} \quad \square \end{aligned}$$

We now use Proposition 2.17 to prove that the q_ψ -mass on a subspace is always bounded above by the p_ψ -mass.

Proposition 2.19. *Let $A \subseteq \mathbb{F}_2^{2n}$ be a subspace. Then*

$$\sum_{a \in A} q_\psi(a) \leq \sum_{a \in A} p_\psi(a).$$

Proof.

$$\begin{aligned} \sum_{a \in A} q_\psi(a) &= |A| \sum_{x \in A^\perp} p_\psi(x)^2 && \text{(Proposition 2.17)} \\ &\leq \frac{|A|}{2^n} \sum_{x \in A^\perp} p_\psi(x) && (p_\psi(x) \leq \frac{1}{2^n}) \\ &= \sum_{a \in A} p_\psi(a). && \text{(Proposition 2.17)} \quad \square \end{aligned}$$

Finally, consider the set $M = \{x \in \mathbb{F}_2^{2n} : 2^n p_\psi(x) > 1/2\}$. We show, via the Schrödinger uncertainty relation, that every pair of elements in this set commutes.⁹

⁹In 1927, Heisenberg observed a tradeoff between knowing a particle's position and momentum, which has since been generalized in several ways. This particular uncertainty relation was derived by Schrödinger in 1930.

Fact 2.20 (Schrödinger uncertainty relation [Sch30, AB08]). For a quantum state ρ and observables A and B ,

$$(\text{tr}(A^2\rho) - \text{tr}(A\rho)^2)(\text{tr}(B^2\rho) - \text{tr}(B\rho)^2) \geq \left| \frac{1}{2}\text{tr}((AB + BA)\rho) - \text{tr}(A\rho)\text{tr}(B\rho) \right|^2.$$

Fact 2.21. Let $M = \{x \in \mathbb{F}_2^{2n} : 2^n p_\psi(x) > \frac{1}{2}\}$. Then for all $x, y \in M$, $[x, y] = 0$.

Proof. Recall that Weyl operators are Hermitian. Let W_x and W_y be two Weyl operators in M . Simple calculations tell us that

$$(\langle \psi | W_x^2 | \psi \rangle - \langle \psi | W_x | \psi \rangle^2)(\langle \psi | W_y^2 | \psi \rangle - \langle \psi | W_y | \psi \rangle^2) = (1 - 2^n p_\psi(x))(1 - 2^n p_\psi(y)) < \frac{1}{4},$$

where we use the fact that $W_x^2 = I$ for all $x \in \mathbb{F}_2^{2n}$. Now, assume that $[x, y] = 1$ for the sake of contradiction. Then $W_x W_y + W_y W_x = 0$, and

$$\begin{aligned} \left| \frac{1}{2} \langle \psi | (W_x W_y + W_y W_x) | \psi \rangle - \langle \psi | W_x | \psi \rangle \langle \psi | W_y | \psi \rangle \right|^2 &= |\langle \psi | W_x | \psi \rangle \langle \psi | W_y | \psi \rangle|^2 \\ &= 4^n p_\psi(x) p_\psi(y) > \frac{1}{4}, \end{aligned}$$

which, by Fact 2.20, is a contradiction. Hence, $[x, y] = 0$. $\square$

We note that Gross, Nezami, and Walter [GNW21, Figure 1] previously gave a graphical proof of Fact 2.21.

2.4 Stabilizer Dimension

We use the following stabilizer complexity measure throughout this work.

Definition 2.22 (Stabilizer dimension). Let $|\psi\rangle$ be an n -qubit pure state. The *stabilizer dimension of $|\psi\rangle$* is the dimension of $\text{Weyl}(|\psi\rangle)$ as a subspace of $\mathbb{F}_2^{2n}$.

Put another way, the stabilizer dimension of a quantum state is k when there is an abelian group of 2^k Pauli matrices that stabilize the state. The stabilizer dimension of a stabilizer state is n , which is maximal, and, for most states, the stabilizer dimension is 0. We note that the stabilizer dimension is closely related to the stabilizer state nullity [BCHK20], and, in particular, for n -qubit states, the stabilizer dimension is n minus the stabilizer state nullity.

Roughly speaking, as one applies single-qubit non-Clifford gates to a stabilizer state, the state becomes farther from all stabilizer states. The stabilizer dimension quantifies this in the following sense.

Lemma 2.23 ([GIKL24, Lemma 4.2]). Let $|\psi\rangle$ be the output state of a t -doped Clifford circuit. Then the stabilizer dimension of $|\psi\rangle$ is at least $n - 2t$.

Proof sketch. First show that a single non-Clifford gate must commute with at least $1/4$ of the Weyl operators in $\text{Weyl}(|\psi\rangle)$, which decreases the stabilizer dimension by at most 2. The proof proceeds by induction on the non-Clifford gates in the circuit, using the fact that Clifford gates preserve stabilizer dimension. $\square$

2.5 Clifford Action on $\mathbb{F}_2^{2n}$

The Clifford group normalizes the Pauli group by conjugation, which induces an action on $\mathbb{F}_2^{2n}$ through the corresponding Weyl operators. To simplify notation, let $C(x)$ denote the action of a Clifford circuit C on $x \in \mathbb{F}_2^{2n}$, and define it to be the $y \in \mathbb{F}_2^{2n}$ such that $W_y = \pm CW_x C^\dagger$. Similarly, for a set $S \subseteq \mathbb{F}_2^{2n}$, define $C(S) := \{C(x) : x \in S\}$. As an example, $C(\mathbb{F}_2^{2n}) = \mathbb{F}_2^{2n}$, as all Clifford circuits normalize the Pauli group up to a ± 1 phase.

Two key properties of the Clifford action are linearity (i.e., $C(x) + C(y) = C(x + y)$), and preservation of the symplectic product (i.e., $[x, y] = [C(x), C(y)]$). In fact, these two properties can be taken as a *definition* of the Clifford group action, which is equivalent to the action of the symplectic group $\text{Sp}(2n, \mathbb{F}_2)$. These properties imply, for example, that the Clifford action preserves symplectic complements and inclusions among subspaces of $\mathbb{F}_2^{2n}$.

Below, we prove one further basic fact about the action of Clifford circuits on symplectic vector spaces over $\mathbb{F}_2$.

Fact 2.24. *Let $|\psi\rangle$ be an n -qubit quantum state, let C be a Clifford circuit, and define $|\phi\rangle := C|\psi\rangle$. Then*

$$p_\phi(x) = p_\psi(C^\dagger(x))$$

for all $x \in \mathbb{F}_2^{2n}$.

Proof.

$$2^n p_\phi(x) = \langle \psi | C^\dagger W_x C | \psi \rangle^2 = 2^n p_\psi(C^\dagger(x)). \quad \square$$

3 Linear Algebra Subroutines

Our algorithms use two linear-algebraic subroutines, which we describe below. First, we give an algorithm for computing the symplectic complement of a subspace.

Lemma 3.1. *Given a set of m vectors whose span is a subspace $A \subseteq \mathbb{F}_2^{2n}$, there is an algorithm that outputs a basis for $A^\perp$ in $O(mn \cdot \min(m, n))$ time.*

Proof. The algorithm works as follows. First, construct a $m \times 2n$ matrix whose rows are the m elements of A given as input. Then swap the left and right $m \times n$ block submatrices, and denote the resulting matrix by M . Observe that for a nonzero vector v , $Mv = 0$ only when the symplectic product between v and all vectors in A is 0. Hence, v is in $A^\perp$, and the nullspace of M is precisely $A^\perp$. Finding a basis for the nullspace of M can be done via Gaussian elimination, which takes $O(mn \cdot \min(m, n))$ time. $\square$

Next, we explain how to find a Clifford circuit whose action on $\mathbb{F}_2^{2n}$ maps an arbitrary d -dimensional isotropic subspace of $\mathbb{F}_2^{2n}$ to the subspace $0^{2n-d} \times \mathbb{F}_2^d$. We note that while the existence of such a Clifford circuit is not difficult to show (cf. [GIKL24, Lemma 5.1]), an explicit and efficient construction requires a bit more effort.

Lemma 3.2. *Given a set of m vectors whose span is a d -dimensional isotropic subspace $A \subset \mathbb{F}_2^{2n}$, there exists an efficient algorithm that outputs a Clifford circuit C such that $C(A) = 0^{2n-d} \times \mathbb{F}_2^d$. The algorithm runs in $O(mn \cdot \min(m, n))$ time, and the circuit size of C (i.e., the number of gates) is $O(nd)$.*

Proof. We will explain the algorithm and then prove its correctness. To begin, run Gaussian elimination on the set of m vectors to get a basis for H such that, when written as a $d \times 2n$ matrix $M = (m_{i,j})$, the matrix M is in row echelon form. This process takes $O(mn \cdot \min(m, n))$ time. The subspace spanned by the rows of M is precisely the subspace A .

The matrix M is essentially a *stabilizer tableau*, and therefore Clifford gates have the following effect on M (for additional detail see, e.g., [AG04]):

- Applying the Hadamard gate on the i th qubit corresponds to swapping the i th and $(n + i)$ th columns of M .
- Applying the phase gate on the i th qubit corresponds to adding the i th column of M to the $(n + i)$ th column of M .
- Applying the CNOT gate with control qubit i and target qubit j corresponds to adding the i th column of M to the j th column M and adding the $(n + j)$ th column of M to the $(n + i)$ th column of M .

Additionally, row operations do not change the subspace spanned by the rows of M and therefore can be done freely.

Our job now is to find a sequence of Hadamard, phase, and CNOT gates that maps M to a matrix whose rows span the subspace $0^{2n-d} \times \mathbb{F}_2^d$; in particular, a matrix with the following form

$$(0 \mid 0 \quad I), \quad (1)$$

where the first 0 is a $d \times n$ matrix of all 0's, the second is a $d \times (n - d)$ matrix of all 0's, and the last is a $d \times d$ identity matrix.

The remainder of the algorithm works as follows.

1. For each row $i \in [d]$ of M :
 - (a) For each $j \in \{i, \dots, n\}$, apply phase and Hadamard gates so that either $m_{i,j} = 1$ and $m_{i,n+j} = 0$ or both are 0.
 - (b) If $m_{i,i} = 0$, then find a $k \in \{i + 1, \dots, n\}$ for which $m_{i,k} = 1$.¹⁰ Apply a CNOT with control qubit i and target qubit k so that $m_{i,i} = 1$.
 - (c) For each $j \in \{i + 1, \dots, n\}$, if $m_{i,j} = 1$, apply CNOT with control qubit i and target qubit j .
 - (d) For $j \in \{i + 1, \dots, d\}$, set $m_{j,i} = 0$.¹¹
2. Apply a Hadamard gate to each of the first d qubits.
3. For $i \in \{0, \dots, d - 1\}$, apply a CNOT with control qubit $n - i$ and target qubit $d - i$. Then apply a CNOT with control qubit $d - i$ and target qubit $n - i$.

¹⁰At least one such k must exist, for if it didn't, then the i th row would be all 0's, which is impossible since the rows of M are linearly independent.

¹¹This corresponds to adding the i th row to the j th row, which, as mentioned earlier, does not change the subspace spanned by the rows of M .

Let C denote the Clifford circuit described by the above process. The algorithm concludes by outputting C . We apply $O(n)$ Clifford gates and do at most $O(d)$ row-sum operations per row. Thus, each iteration of Step 1 takes $O(nd)$ time. Since $d \leq \min(m, n)$, the overall running time is therefore $O(mn \cdot \min(m, n))$, due to the Gaussian elimination step at the beginning of the algorithm and the size of the circuit is at most $O(nd)$.

To prove correctness, we must argue that $C(A) = 0^{2n-d} \times \mathbb{F}_2^d$, or equivalently, that the algorithm above maps M to a matrix as in Eq. (1).

First, we show that Step 1 of the algorithm maps M to a matrix of the form

$$(I \quad 0 \mid 0),$$

where I is the $d \times d$ identity matrix. It is clear that after the first iteration of Step 1 completes, $m_{1,1} = 1$ and the remaining entries of the first row and column are 0's. By way of induction, assume that this is true after the first $i - 1$ iterations, so that the resulting matrix looks as follows:

$$\left(\begin{array}{c|c|c} \overbrace{\begin{array}{ccc} 1 & & \\ & \ddots & \\ & & 1 \end{array}}^{n \text{ columns}} & \begin{array}{c} 0 \dots 0 \\ \vdots \\ 0 \dots 0 \end{array} & \begin{array}{c} 0 \dots 0 \dots 0 \\ \vdots \\ 0 \dots 0 \dots 0 \end{array} \\ \hline \begin{array}{c} 0 \dots 0 \\ \vdots \\ 0 \dots 0 \end{array} & m_{i,i} & \begin{array}{c} 0 \dots 0 \\ \vdots \\ 0 \dots 0 \end{array} \end{array} \right). \quad (2)$$

$\underbrace{\hspace{10em}}_{i-1 \text{ columns}} \qquad \qquad \underbrace{\hspace{10em}}_{i-1 \text{ columns}}$

In the top row, from left to right, the first block is the $(i - 1) \times (i - 1)$ identity matrix, then an $(i - 1) \times (n - i + 1)$ block of all 0's, and finally an $(i - 1) \times 2n$ block of all 0's. In the bottom row, from left to right, the first block is a $(d - i + 1) \times (i - 1)$ matrix of all 0's, then a $(d - i + 1) \times (n - i + 1)$ block being processed by the algorithm, then a $(d - i + 1) \times (i - 1)$ block of all 0's, and finally a $(d - i + 1) \times (n - i + 1)$ block being processed by the algorithm.

We will argue that after the i th iteration of Step 1 finishes, the matrix will have the form of Eq. (2) but with $m_{i,i} = 1$ and the rest of the i th row and i th column cleared to 0. First, observe that the third block in the second row must be all 0's if the first block of the top row is the identity matrix because the subspace spanned by the rows is isotropic (and applying Clifford gates will not affect that). It is also clear that the operations performed in the i th iteration will set $m_{i,i} = 1$, set $m_{i,j} = 0$ for $j \in \{i + 1, \dots, 2n\}$, and set $m_{j,i} = 0$ for $j \in \{i + 1, \dots, d\}$. Therefore, we just need to argue that the i th iteration does not reintroduce 1's into the blocks of 0's or affect the $(i - 1) \times (i - 1)$ identity matrix in the first block of the first row. Observe that neither of these can happen as long as Hadamard gates and CNOT gates are not applied to the first $i - 1$ qubits in the i th iteration. Indeed, our algorithm does not apply any gates to the first $i - 1$ qubits, so the structure of the matrix in Eq. (2) is preserved. Therefore, once Step 1 terminates, the resulting matrix will be a $d \times d$ identity matrix in the first block and the remaining entries of the matrix will be 0's.

The layer of Hadamard gates in Step 2 maps the $d \times d$ identity matrix to the right block of the matrix, i.e.,

$$(0 \mid I \ 0).$$

Finally, the CNOT gates in Step 3 move the identity matrix to the rightmost side of the tableau, matching the goal shown in Eq. (1). This can be verified by explicit calculation. We note that the CNOT gates move the identity matrix by starting with the rightmost column and then proceeding leftward, which is critical for correctness when $d > n/2$. Hence, C performs the desired mapping. $\square$

4 On Subspaces with Large p_ψ or q_ψ -mass

We prove a series of lemmas about subspaces of $\mathbb{F}_2^{2n}$ that form the backbone of our property testing and tomography algorithms. Most of these lemmas presuppose that a given subspace S is “heavy,” in the sense that either p_ψ or q_ψ places a large probability mass on S . We begin by showing that heavy subspaces must be isotropic. Then, we show that if a sufficiently heavy subspace S is $(n - t)$ -dimensional, $|\psi\rangle$ is close in fidelity to a state of the form $C|\varphi\rangle|x\rangle$, where C is a Clifford circuit determined by S , $|x\rangle$ is an $(n - t)$ -qubit basis state, and $|\varphi\rangle$ is a t -qubit state. Furthermore, we explain how to use Lemma 3.2 to efficiently construct the Clifford circuit C given S . With these lemmas in hand, the remaining work in this paper will lie largely in showing how to find such heavy subspaces efficiently.

4.1 Subspaces With Large Mass Are Isotropic

First, we relate the isotropicity of a subspace to the q_ψ mass on its symplectic complement.

Lemma 4.1. *Let S be a subspace of $\mathbb{F}_2^{2n}$ such that*

$$\sum_{x \in S^\perp} q_\psi(x) > \frac{5}{8}.$$

Then S is isotropic.

Proof. Let $A := \{x \in S : 2^n p_\psi(x) > 1/2\}$. By Fact 2.21, every pair of elements in A has symplectic product zero. Furthermore, every pair of elements in $\langle A \rangle$ also has symplectic product zero by linearity of the symplectic product, showing that $\langle A \rangle$ is isotropic. Thus, if we can show that $\langle A \rangle = S$ then S must also be isotropic.

Suppose for a contradiction that $|A| \leq \frac{|S|}{2}$. Then:

$$\begin{aligned} \sum_{x \in S^\perp} q_\psi(x) &= |S^\perp| \sum_{x \in S} p_\psi(x)^2 && \text{(Proposition 2.17)} \\ &= |S^\perp| \left(\sum_{x \in A} p_\psi(x)^2 + \sum_{x \in S \setminus A} p_\psi(x)^2 \right) \\ &\leq |S^\perp| \left(\frac{|S|}{2} \cdot \frac{1}{4^n} + \frac{|S|}{2} \cdot \frac{1}{4 \cdot 4^n} \right) \\ &= \frac{|S| \cdot |S^\perp|}{4^n} \cdot \left(\frac{1}{2} + \frac{1}{8} \right) \\ &= \frac{5}{8}, \end{aligned}$$

which contradicts the assumption of the lemma. So, $|A| > \frac{|S|}{2}$. Since a proper subspace of S can have at most $\frac{|S|}{2}$ elements, it follows that $\langle A \rangle = S$ and S is isotropic. $\square$

Next, we prove an analogous statement with p_ψ in place of q_ψ . For the next lemma, note that by [Proposition 2.17](#), the assumption on S is equivalent to

$$\sum_{x \in S^\perp} p_\psi(x) > \frac{3}{4},$$

which is more directly comparable to [Lemma 4.1](#). The proof is also similar.

Lemma 4.2. *Let S be a subspace of $\mathbb{F}_2^{2n}$ such that*

$$\sum_{x \in S} p_\psi(x) > \frac{3}{4} \frac{|S|}{2^n}.$$

Then S is isotropic.

Proof. Let $A := \{x \in S : 2^n p_\psi(x) > 1/2\}$. By [Fact 2.21](#), every pair of elements in A has symplectic product zero. Furthermore, every pair of elements in $\langle A \rangle$ also has symplectic product zero by linearity of the symplectic product, showing that $\langle A \rangle$ is isotropic. Thus, if we can show that $\langle A \rangle = S$ then S must also be isotropic.

Suppose for a contradiction that $|A| \leq \frac{|S|}{2}$. Then:

$$\begin{aligned} \sum_{x \in S} p_\psi(x) &= \sum_{x \in A} p_\psi(x) + \sum_{x \in S \setminus A} p_\psi(x) \\ &\leq \frac{|S|}{2} \cdot \frac{1}{2^n} + \frac{|S|}{2} \cdot \frac{1}{2 \cdot 2^n} \\ &= \frac{|S|}{2^n} \cdot \left(\frac{1}{2} + \frac{1}{4} \right) \\ &= \frac{3}{4} \frac{|S|}{2^n}, \end{aligned}$$

which contradicts the assumption of the lemma. So, $|A| > \frac{|S|}{2}$. Since a proper subspace of S can have at most $\frac{|S|}{2}$ elements, it follows that $\langle A \rangle = S$ and therefore S is isotropic. $\square$

4.2 Product State Structure

We now establish the critical relation between p_ψ -mass on isotropic subspaces and stabilizer dimension. As noted earlier, we show that if an $(n-t)$ -dimensional isotropic subspace has large p_ψ -mass, then there is a Clifford circuit C (that can be constructed efficiently) that approximately maps $|\psi\rangle$ to a product state $|\varphi\rangle |x\rangle$, where $|x\rangle$ is an $(n-t)$ -qubit basis state and $|\varphi\rangle$ is an arbitrary t -qubit state.

Special cases of this relation have appeared implicitly before. In particular, if the p_ψ mass on the subspace is *maximal* (i.e., exactly $\frac{1}{2^t}$), then this can be viewed as a restatement of the fact that stabilizer codes can be encoded using Clifford circuits [\[Got06\]](#). This special case also essentially follows from [\[LOLH24, Theorem 2\]](#). The main contribution of this subsection is to show that this mapping is robust to error.

To begin, we prove a relation between sums over basis state projections and sums of Pauli- Z strings.

Proposition 4.3.

$$\sum_{x \in \mathbb{F}_2^k} |x\rangle\langle x| \otimes |x\rangle\langle x| = \frac{1}{2^k} \sum_{x \in \mathbb{F}_2^k} Z^x \otimes Z^x,$$

where $Z^x := Z^{x_1} \otimes \dots \otimes Z^{x_k}$.

Proof. We use the fact that for any $x \in \mathbb{F}_2^k$, $|x\rangle\langle x| = 2^{-k} \sum_{a \in \mathbb{F}_2^k} Z^a (-1)^{a \cdot x}$.¹²

$$\begin{aligned} \sum_{x \in \mathbb{F}_2^k} |x\rangle\langle x| \otimes |x\rangle\langle x| &= \sum_{x \in \mathbb{F}_2^k} \frac{1}{4^k} \left(\sum_{a \in \mathbb{F}_2^k} Z^a (-1)^{a \cdot x} \right) \otimes \left(\sum_{b \in \mathbb{F}_2^k} Z^b (-1)^{b \cdot x} \right) \\ &= \frac{1}{4^k} \sum_{a, b \in \mathbb{F}_2^k} Z^a \otimes Z^b \sum_{x \in \mathbb{F}_2^k} (-1)^{(a+b) \cdot x} \\ &= \frac{1}{2^k} \sum_{a \in \mathbb{F}_2^k} Z^a \otimes Z^a. \end{aligned} \quad \square$$

Next, we prove that the p_ψ -mass on an isotropic subspace has a nice operational interpretation.

Proposition 4.4. *Let $|\psi\rangle$ be an n -qubit state, and let $S = 0^{n+t} \times \mathbb{F}_2^{n-t}$. Upon measuring the last $n-t$ qubits in the computational basis on 2 copies of $|\psi\rangle$, the probability of observing the same string $x \in \mathbb{F}_2^{n-t}$ twice (i.e., the collision probability) is*

$$2^t \sum_{x \in S} p_\psi(x).$$

Proof. The probability of observing some $x \in \mathbb{F}_2^{n-t}$ twice is

$$\begin{aligned} \sum_{x \in \mathbb{F}_2^{n-t}} \langle \psi | (I^{\otimes t} \otimes |x\rangle\langle x|) | \psi \rangle^2 &= \langle \psi |^{\otimes 2} \left(\sum_{x \in \mathbb{F}_2^{n-t}} I^{\otimes t} \otimes |x\rangle\langle x| \otimes I^{\otimes t} \otimes |x\rangle\langle x| \right) | \psi \rangle^{\otimes 2} \\ &= \langle \psi |^{\otimes 2} \left(\frac{1}{2^{n-t}} \sum_{x \in \mathbb{F}_2^{n-t}} I^{\otimes t} \otimes Z^x \otimes I^{\otimes t} \otimes Z^x \right) | \psi \rangle^{\otimes 2} \\ &= \frac{1}{2^{n-t}} \sum_{x \in \mathbb{F}_2^{n-t}} \langle \psi | I^{\otimes t} \otimes Z^x | \psi \rangle^2 \\ &= 2^t \sum_{x \in S} p_\psi(x). \end{aligned}$$

The third step follows from Proposition 4.3 by treating the $I^{\otimes t}$ as constants. $\square$

Now we prove the main lemma of this subsection for the special case where the isotropic subspace is a subset of $\{I, Z\}^{\otimes n}$.

Lemma 4.5. *Let $S = 0^{n+t} \times \mathbb{F}_2^{n-t}$, and suppose that*

$$\sum_{x \in S} p_\psi(x) \geq \frac{1 - \varepsilon}{2^t}.$$

¹²This is essentially the Weyl expansion of $|x\rangle\langle x|$. Since the expression only involves I 's and Pauli Z 's, we can write the expansion in this simplified form.

Then there exists an $(n-t)$ -qubit computational basis state $|x\rangle$ and a t -qubit quantum state

$$|\varphi\rangle := \frac{(I \otimes \langle x|) |\psi\rangle}{\|(I \otimes \langle x|) |\psi\rangle\|_2},^{13}$$

such that the fidelity between $|\varphi\rangle |x\rangle$ and $|\psi\rangle$ is at least $1 - \varepsilon$.

Proof. We can always write $|\psi\rangle = \sum_{x \in \mathbb{F}_2^{n-t}} \alpha_x |\varphi_x\rangle |x\rangle$ where $\sum_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^2 = 1$. If we can show that

$$\max_{x \in \mathbb{F}_2^{n-t}} |\langle \psi | (|\varphi_x\rangle |x\rangle)|^2 \geq 1 - \varepsilon,$$

then we are done, by taking $|\varphi\rangle = |\varphi_x\rangle$. First,

$$\begin{aligned} \max_{x \in \mathbb{F}_2^{n-t}} |\langle \psi | (|\varphi_x\rangle |x\rangle)|^2 &= \max_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^2 \\ &= \max_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^2 \cdot \sum_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^2 \\ &\geq \sum_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^4. \end{aligned}$$

Observe that $\sum_x |\alpha_x|^4$ is precisely the collision probability when measuring the last $n-t$ qubits of $|\psi\rangle$ in the computational basis. Hence, by [Proposition 4.4](#),

$$\sum_{x \in \mathbb{F}_2^{n-t}} |\alpha_x|^4 = 2^t \sum_{x \in S} p_\psi(x) \geq 1 - \varepsilon. \quad \square$$

Finally, we generalize the previous lemma using the Clifford mapping algorithm from [Lemma 3.2](#).

Lemma 4.6. *Let S be an isotropic subspace of dimension $n-t$, and suppose that*

$$\sum_{x \in S} p_\psi(x) \geq \frac{1 - \varepsilon}{2^t}.$$

Then there exists a state $|\widehat{\psi}\rangle$ with $S \subseteq \text{Weyl}(|\widehat{\psi}\rangle)$ such that the fidelity between $|\widehat{\psi}\rangle$ and $|\psi\rangle$ is at least $1 - \varepsilon$.

In particular, $|\widehat{\psi}\rangle = C^\dagger |\varphi\rangle |x\rangle$, where $|x\rangle$ is an $(n-t)$ -qubit basis state,

$$|\varphi\rangle := \frac{(I \otimes \langle x|) C |\psi\rangle}{\|(I \otimes \langle x|) C |\psi\rangle\|_2}$$

is a t -qubit quantum state, and C is the Clifford circuit mapping S to $0^{n+t} \times \mathbb{F}_2^{n-t}$ described in [Lemma 3.2](#).

Proof. Define $|\phi\rangle := C |\psi\rangle$. Then, by [Fact 2.24](#),

$$\sum_{x \in S} p_\psi(x) = \sum_{x \in C(S)} p_\phi(x) \geq \frac{1 - \varepsilon}{2^t}.$$

Therefore, by [Lemma 4.5](#), $C |\psi\rangle$ is $(1 - \varepsilon)$ -close in fidelity to a state $|\varphi\rangle |x\rangle$, where $|x\rangle$ is an $(n-t)$ -qubit basis state and

$$|\varphi\rangle := \frac{(I \otimes \langle x|) C |\psi\rangle}{\|(I \otimes \langle x|) C |\psi\rangle\|_2}$$

is a t -qubit quantum pure state. Since fidelity is unitarily invariant, the fidelity between $|\psi\rangle$ and $|\widehat{\psi}\rangle = C^\dagger |\varphi\rangle |x\rangle$ is also at least $1 - \varepsilon$. Clearly, $C(S) \subseteq \text{Weyl}(|\varphi\rangle |x\rangle)$, and therefore, by [Fact 2.24](#), $S \subseteq \text{Weyl}(|\widehat{\psi}\rangle)$. $\square$

¹³This is to say that $|\varphi\rangle$ is obtained by postselecting on measuring the last $n-t$ qubits of $|\psi\rangle$ to be $|x\rangle$.

5 Property Testing Stabilizer Dimension

As a first application, we present an efficient algorithm for property testing stabilizer dimension. Recall that a property tester for a class $\mathcal{Q}$ of quantum states takes copies of a state $|\psi\rangle$ as input and determines whether $|\psi\rangle \in \mathcal{Q}$ or $|\psi\rangle$ is ε -far from all such states (according to some measure of distance), promised that one of these is the case. Our algorithm efficiently tests whether an input state has stabilizer dimension at least k or has fidelity less than $1 - \varepsilon$ with all such states.

Algorithm 1: Property Testing Stabilizer Dimension

Input: $\frac{16n+8\log(1/\delta)}{\varepsilon}$ copies of $|\psi\rangle$, $k \in [n]$, $\varepsilon \in (0, 3/8)$, and $\delta \in (0, 1]$

Promise: $|\psi\rangle$ has stabilizer dimension at least k or is ε -far in fidelity from all such states

Output: 1 if $|\psi\rangle$ has stabilizer dimension at least k , 0 otherwise, with probability at least $1 - \delta$

- 1 Perform Bell difference sampling to draw $\frac{4n+2\log(1/\delta)}{\varepsilon}$ samples from q_ψ . Denote the span of the samples by $S^\perp$.
 - 2 Let $\hat{k} = \dim S = 2n - \dim S^\perp$.
 - 3 **return** 1 if $\hat{k} \geq k$ and 0 otherwise.
-

Theorem 5.1. *Let $|\psi\rangle$ be an n -qubit quantum state. For $\varepsilon \in (0, 3/8)$, Algorithm 1 determines whether $|\psi\rangle$ has stabilizer dimension at least k or has fidelity at most $1 - \varepsilon$ with all such states, promised that one of these is the case. The algorithm uses $\frac{16n+8\log(1/\delta)}{\varepsilon}$ copies of $|\psi\rangle$ and $O\left(\frac{n^3+n^2\log(1/\delta)}{\varepsilon}\right)$ time, and succeeds with probability at least $1 - \delta$.*

Proof. First, suppose that $|\psi\rangle$ has stabilizer dimension at least k . By Fact 2.18, q_ψ is supported on a subspace of dimension at most $2n - k$. So, no matter what, $\dim S^\perp$ will never exceed $2n - k$, and therefore $\hat{k}$ will always be at least k . Hence, Algorithm 1 accepts with probability 1.

Now suppose that $|\psi\rangle$ has fidelity less than $1 - \varepsilon$ with every state of stabilizer dimension at least k . By Proposition 2.19 and Lemma 2.3, with probability at least $1 - \delta$, the subspace S satisfies

$$\sum_{x \in S^\perp} p_\psi(x) \geq \sum_{x \in S^\perp} q_\psi(x) \geq 1 - \varepsilon.$$

Assuming this occurs, applying Proposition 2.17 gives

$$\sum_{x \in S} p_\psi(x) = \frac{|S|}{2^n} \sum_{x \in S^\perp} p_\psi(x) \geq \frac{1 - \varepsilon}{2^{n-\hat{k}}}.$$

Since $\sum_{x \in S^\perp} q_\psi(x) \geq 1 - \varepsilon > 5/8$, S is isotropic by Lemma 4.1. As such, by Lemma 4.6, there exists a state of stabilizer dimension at least $\hat{k}$ that has fidelity at least $1 - \varepsilon$ with $|\psi\rangle$. By assumption, we must have $\hat{k} < k$, and thus Algorithm 1 rejects with probability at least $1 - \delta$.

Overall, we find that in either case, with probability at least $1 - \delta$, the algorithm succeeds. It remains to bound the runtime. Lemma 3.1 tells us that computing a basis for $S^\perp$ requires $O(mn^2)$ time, where $m = O\left(\frac{n+\log(1/\delta)}{\varepsilon}\right)$ is the number of Bell difference samples taken by the algorithm. This dominates the running time. $\square$

6 Tomography Reduces to Finding Heavy Subspaces

In this section, we describe a reduction from tomography on states $|\psi\rangle$ with large stabilizer dimension to finding any subspace $S \subseteq \mathbb{F}_2^{2n}$ that approximates $\text{Weyl}(|\psi\rangle)$ (in a sense that is made explicit below). This reduction uses only single-copy measurements. In the later sections, we will give algorithms to find such subspaces S using either entangled or single-copy measurements.

6.1 Pure State Tomography

Our reduction requires a general pure state tomography algorithm as a subroutine. The complexity of our reduction will be parameterized by the complexity of the tomography algorithm used. To describe these runtimes succinctly, it will be convenient to introduce the following notation.

Definition 6.1 (Single-copy pure state tomography copy and time complexities). Let $N_{n,\varepsilon,\delta}$ and $M_{n,\varepsilon,\delta}$ be the sample and time complexities, respectively, of pure state tomography on n qubits to trace distance ε with failure probability at most δ , using only single-copy measurements.

For context, it was shown by [FBaK21] that there exists a single-copy pure state tomography procedure that uses $O(2^n n \log(1/\delta)\varepsilon^{-4})$ copies and $O(4^n n^3 \log(1/\delta)\varepsilon^{-5})$ time to output a state $|\hat{\psi}\rangle$ that is ε -close to $|\psi\rangle$ in trace distance with probability at least $1 - \delta$. To our knowledge, this is the state-of-the-art when it comes to scaling in the dimension $d = 2^n$ of the system (though, it does not match the best-known lower bound of $\Omega(d)$).

Theorem 6.2 ([FBaK21]). *Given access to copies of an n -qubit pure state $|\psi\rangle$, there is a single-copy algorithm that uses $O(2^n n \log(1/\delta)\varepsilon^{-4})$ copies, $O(4^n n^3 \log(1/\delta)\varepsilon^{-5})$ time, and outputs a state $|\hat{\psi}\rangle$ that is ε -close to $|\psi\rangle$ in trace distance with probability at least $1 - \delta$. The algorithm only requires applying random Clifford circuits and classical post-processing.*

Alternatively, it was shown by [AG23] that there exists a single-copy pure state tomography procedure that uses $O(16^n \log(1/\delta)\varepsilon^{-2})$ copies and $O(32^n \log(1/\delta)\varepsilon^{-2})$ time, sacrificing dependence on n for a better dependence on ε .

Theorem 6.3 ([AG23, Section 5]). *Given access to copies of an n -qubit pure state $|\psi\rangle$, there is a single-copy algorithm that uses $O(16^n \log(1/\delta)\varepsilon^{-2})$ copies, $O(32^n \log(1/\delta)\varepsilon^{-2})$ time, and outputs a state $|\hat{\psi}\rangle$ that is ε -close to $|\psi\rangle$ in trace distance with probability at least $1 - \delta$. The algorithm only requires applying Clifford circuits and classical post-processing.*

In either case, we have the upper bound $M_{n,\varepsilon,\delta}, N_{n,\varepsilon,\delta} \leq \text{poly}(2^n, 1/\varepsilon, \log(1/\delta))$.

6.2 Reduction Algorithm

We are now ready to describe the reduction algorithm. Recall Lemma 4.6, which shows that given a large isotropic subspace that accounts for a large fraction of the p_ψ mass, there exists a Clifford circuit that approximately maps $|\psi\rangle$ to a product $|\varphi\rangle|x\rangle$ of a state $|\varphi\rangle$ on few qubits and a computational basis state $|x\rangle$. Our algorithm below amounts to computing this Clifford circuit C and then performing tomography on the states $|\varphi\rangle$ (via the general tomography algorithm) and $|x\rangle$ (by measuring in the computational basis).

Algorithm 2: Learning reduction using single-copy measurements

Input: Black-box access to copies of $|\psi\rangle$, $\varepsilon, \delta \in (0, 1)$, and $S \subseteq \mathbb{F}_2^{2n}$

Promise: S is a subspace such that $\sum_{x \in S} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n}$ and

$$\dim(S) \geq n - t$$

Output: A classical description of $|\hat{\psi}\rangle$ such that $d_{\text{tr}}(|\psi\rangle, |\hat{\psi}\rangle) \leq \varepsilon$ with probability at least $1 - \delta$

- 1 Let $\hat{t}$ be such that $\dim(S) = n - \hat{t}$ and $\hat{t} \leq t$.
 - 2 Use the algorithm in [Lemma 3.2](#) to construct a Clifford circuit C that maps S to $0^{n+\hat{t}} \times \mathbb{F}_2^{n-\hat{t}}$.
 - 3 Measure the last $n - \hat{t}$ qubits of $24 \log(3/\delta)$ copies of $C|\psi\rangle$. Let $\hat{x} \in \mathbb{F}_2^{n-\hat{t}}$ be the majority result.
 - 4 **for** $i \leftarrow 1$ **to** $\frac{4}{3}N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \frac{8}{9} \log(3/\delta)$ **do**
 - 5 Measure the last $n - \hat{t}$ qubits of a copy of $C|\psi\rangle$ and let $x_i \in \{0, 1\}^{n-\hat{t}}$ be the outcome.
 - 6 **if** $x_i = \hat{x}$ **then**
 - 7 Use the state on the first $\hat{t}$ qubits as the input to a pure state single-copy tomography algorithm with accuracy $\varepsilon/2$ in trace distance and failure probability at most $\delta/3$.
 - 8 **else**
 - 9 Discard the state.
 - 10 Let $|\hat{\varphi}\rangle$ be the classical description output by the pure state single-copy tomography algorithm.
 - 11 **return** $|\hat{\psi}\rangle = C^\dagger |\hat{\varphi}\rangle |\hat{x}\rangle$.
-

To aid intuition, note that if $|\psi\rangle$ has stabilizer dimension $n - t$ and $S = \text{Weyl}(|\psi\rangle)$, then the condition of the algorithm is satisfied with $\varepsilon = 0$. So, the algorithm may be understood as showing how to perform tomography on $|\psi\rangle$ given a sufficiently good approximation S to $\text{Weyl}(|\psi\rangle)$.

Theorem 6.4. *Let $|\psi\rangle$ be an n -qubit state, and suppose there exists a subspace $S \subseteq \mathbb{F}_2^{2n}$ that satisfies the following conditions:*

- (i) $\dim(S) \geq n - t$, and
- (ii) $\sum_{x \in S} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n}$.

Then, given S , [Algorithm 2](#) outputs a classical description of $|\hat{\psi}\rangle$ such that $d_{\text{tr}}(|\psi\rangle, |\hat{\psi}\rangle) \leq \varepsilon$ with probability at least $1 - \delta$ using $\frac{4}{3}N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \frac{224}{9} \log(3/\delta)$ copies of $|\psi\rangle$, and $M_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + O\left(n^2 \left(N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \log(1/\delta)\right)\right)$ time,¹⁴ while performing only single-copy measurements.

Proof. [Lemma 4.2](#) implies that a subspace S satisfying [Condition \(ii\)](#) must be isotropic, because $\varepsilon < 1$. The algorithm from [Lemma 3.2](#) produces a Clifford circuit C that maps S to $0^{n+\hat{t}} \times \mathbb{F}_2^{n-\hat{t}}$. By [Lemma 4.6](#), there exists an $(n - \hat{t})$ -qubit basis state $|x\rangle$ and a $\hat{t}$ -qubit state

$$|\varphi\rangle := \frac{(I \otimes \langle x|)C|\psi\rangle}{\|(I \otimes \langle x|)C|\psi\rangle\|_2},$$

¹⁴As seen by [Theorems 6.2](#) and [6.3](#), we can usually simplify this to just $O\left(M_{t, \frac{\varepsilon}{2}, \frac{\delta}{2}}\right)$ in practice.

such that the fidelity between $C|\psi\rangle$ and $|\varphi\rangle|x\rangle$ is at least $1 - \varepsilon^2/4$, implying that

$$d_{\text{tr}}(|\psi\rangle, C^\dagger |\varphi\rangle|x\rangle) \leq \frac{\varepsilon}{2}. \quad (3)$$

In [Line 3](#), for $1 \leq i \leq m = 24 \log(3/\delta)$, let x_i denote the i th $(n - \hat{t})$ -bit measurement outcome upon measuring the last $n - \hat{t}$ qubits of $C|\psi\rangle$ in the computational basis. Define the indicator random variable X_i as

$$X_i = \begin{cases} 1 & \text{if } x_i = x. \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Because the fidelity between $C|\psi\rangle$ and $|\varphi\rangle|x\rangle$ is at least $1 - \varepsilon^2/4$ and $\varepsilon \in (0, 1)$, $\Pr[X_i = 1] \geq 3/4$ and $\mu := \mathbf{E}[\sum_i X_i] \geq 3m/4$. By a Chernoff bound ([Fact 2.1](#)),

$$\begin{aligned} \Pr \left[\sum_{i=1}^m X_i \leq \frac{m}{2} \right] &\leq \Pr \left[\sum_{i=1}^m X_i \leq \left(1 - \frac{1}{3}\right) \mu \right] \\ &\leq \exp \left(-\frac{\mu}{18} \right) \\ &\leq \exp \left(-\frac{m}{24} \right) \\ &= \delta/3. \end{aligned}$$

Hence, over half of the m samples will be x with probability at least $1 - \delta/3$, which means that we will have $\hat{x} = x$ with probability $1 - \delta/3$.

Assume now that we have successfully learned $\hat{x} = x$. We then need to perform pure state tomography on $|\varphi\rangle$. Note that the definition of $|\varphi\rangle$ means that $|\varphi\rangle$ is the state on $\hat{t}$ qubits conditioned on measuring $|x\rangle$ on the last $n - \hat{t}$ qubits of $C|\psi\rangle$. In [Line 4](#), we need to repeat this postselection enough times until we have $N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}}$ copies of $|\varphi\rangle$ to feed into the tomography algorithm.

In [Line 5](#), for $1 \leq i \leq m = \frac{4}{3}N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \frac{8}{9} \log(3/\delta)$, define the random variable Y_i by

$$Y_i = \begin{cases} 1 & \text{if } x_i = x. \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

As with the X_i 's, because the fidelity between $C|\psi\rangle$ and $|\varphi\rangle|x\rangle$ is at least $1 - \varepsilon^2/4$ and $\varepsilon \in (0, 1)$, $\Pr[Y_i = 1] \geq 3/4$ and $\mu := \mathbf{E}[\sum_i Y_i] \geq 3m/4$. By Hoeffding's inequality ([Fact 2.2](#)),

$$\Pr \left[\sum_{i=1}^m Y_i < N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} \right] \leq \Pr \left[\sum_{i=1}^m Y_i - \mu \leq N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} - \frac{3}{4}m \right] \leq \exp \left(-\frac{2}{m} \left(N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} - \frac{3}{4}m \right)^2 \right).$$

Let $\gamma = \frac{8}{9} \log(3/\delta) / N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}}$ so that we may write $m = \left(\frac{4}{3} + \gamma\right) N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}}$. Then the above expression is upper bounded by

$$\exp \left(-\frac{9}{8} \frac{\gamma^2}{\frac{4}{3} + \gamma} N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} \right) < \exp \left(-\frac{9}{8} \gamma N_{\hat{t}, \frac{\varepsilon}{2}, \frac{\delta}{3}} \right) \leq \delta/3.$$

This is to say that with probability at least $1 - \delta/3$, the algorithm postselects enough copies of $|\varphi\rangle$ in [Line 7](#) to successfully run pure state tomography on $|\varphi\rangle$.

By assumption, the pure state tomography algorithm in [Line 10](#) fails with probability at most $\delta/3$. Assuming the tomography succeeds in outputting a state $|\hat{\varphi}\rangle$ satisfying $d_{\text{tr}}(|\varphi\rangle, |\hat{\varphi}\rangle) \leq \varepsilon/2$, the returned state $C^\dagger |\hat{\varphi}\rangle |\hat{x}\rangle$ satisfies

$$\begin{aligned} d_{\text{tr}}(|\psi\rangle, C^\dagger |\hat{\varphi}\rangle |\hat{x}\rangle) &\leq d_{\text{tr}}(|\psi\rangle, C^\dagger |\varphi\rangle |x\rangle) + d_{\text{tr}}(|\varphi\rangle |x\rangle, |\hat{\varphi}\rangle |\hat{x}\rangle) \\ &\leq \frac{\varepsilon}{2} + d_{\text{tr}}(|\varphi\rangle |x\rangle, |\hat{\varphi}\rangle |\hat{x}\rangle) \\ &\leq \varepsilon. \end{aligned}$$

Above, first step follows by the triangle inequality and the fact that trace distance is unitarily invariant. The second step follows from [Eq. \(3\)](#), and the final step follows assuming $\hat{x} = x$ and $d_{\text{tr}}(|\varphi\rangle, |\hat{\varphi}\rangle) \leq \varepsilon/2$.

By applying a union bound over the “bad” events (namely, [Lines 3, 7](#) and [10](#) of the algorithm), we have that the overall success probability of the algorithm is at least $1 - \delta$.

To conclude, we bound the copy and time complexities of the algorithm. Since $\hat{t} \leq t$, we use $O\left(N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \log(1/\delta)\right)$ copies in the for loop of [Line 4](#), which dominates the sample complexity. Producing these states requires time $O(n^2)$ per state because of the size of the Clifford circuit from [Lemma 3.2](#). The runtime of the pure state single-copy tomography algorithm in [Line 10](#) otherwise dominates the running time. The overall time complexity is therefore

$$O\left(n^2 \left(N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \log(1/\delta)\right)\right) + M_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}}. \quad \square$$

Note that the dependence on t comes entirely from the pure state tomography algorithm (e.g., [Theorem 6.2](#) or [Theorem 6.3](#)) in [Line 10](#) of [Algorithm 2](#). Hence, one can upgrade this part of the algorithm with improved pure state single-copy tomography algorithms when/if they are discovered.

We also note that if one is allowed a quantum memory, a constant factor improvement in sample complexity is possible by saving the first $\hat{t}$ qubits in [Line 3](#). Once x has been determined via majority, these copies of $|\varphi\rangle$ can be fed into the pure state tomography algorithm.

7 Learning the Unsigned Stabilizer Group Using Bell Measurements

In this section, we present the first of two algorithms for learning a subspace that satisfies the conditions of [Theorem 6.4](#). The algorithm in this section uses Bell difference sampling and is closely related to the property testing algorithm for stabilizer dimension ([Algorithm 1](#)).

Algorithm 3: Approximating Weyl($|\psi\rangle$) using Bell measurements

Input: Black-box access to copies of $|\psi\rangle$ and $\varepsilon, \delta \in (0, 1)$

Promise: $|\psi\rangle$ has stabilizer dimension at least $n - t$

Output: A subspace $S \subset \mathbb{F}_2^{2n}$ of dimension at least $n - t$ such that

$$\sum_{x \in S} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n} \text{ with probability at least } 1 - \delta$$

- 1 Perform Bell difference sampling to draw $\frac{8 \log(1/\delta) + 16n}{\varepsilon^2}$ samples from q_ψ . Denote the span of the samples by $S^\perp$.
 - 2 Use the algorithm in [Lemma 3.1](#) to compute the symplectic complement S of the subspace spanned by the samples.
 - 3 **return** S
-

Showing that [Algorithm 3](#) satisfies the conditions of [Theorem 6.4](#) proceeds similarly to the analysis of [Algorithm 1](#) in [Theorem 5.1](#).

Theorem 7.1. *Let $|\psi\rangle$ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of $|\psi\rangle$ as input, [Algorithm 3](#) succeeds at outputting a subspace S that satisfies the conditions of [Theorem 6.4](#) with probability at least $1 - \delta$. The algorithm uses*

$$\frac{32 \log(1/\delta) + 64n}{\varepsilon^2}$$

samples and

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

time.

Proof. By [Fact 2.18](#), the support of q_ψ is a subspace of dimension at most $n + t$, so $\dim S^\perp \leq n + t$ and therefore $\dim S \geq n - t$, by [Fact 2.8](#). This establishes [Condition \(i\)](#).

By [Lemma 2.3](#), except with probability at most δ , the Bell difference sampling phase of the algorithm in [Line 1](#) finds a subspace $S^\perp \subseteq \mathbb{F}_2^{2n}$ such that

$$\sum_{x \in S^\perp} q_\psi(x) \geq 1 - \frac{\varepsilon^2}{4}.$$

From [Proposition 2.19](#), we know that

$$\sum_{x \in S^\perp} p_\psi(x) \geq 1 - \frac{\varepsilon^2}{4}.$$

In the next step, the algorithm computes S . Using [Proposition 2.17](#), we have that

$$\sum_{x \in S} p_\psi(x) = \frac{|S|}{2^n} \sum_{x \in S^\perp} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n},$$

which establishes [Condition \(ii\)](#).

The upper bound on sample complexity is immediate from the fact that each Bell difference sample requires 4 copies of $|\psi\rangle$, so it only remains to bound the time complexity. The Bell difference sampling phase takes $O(n)$ time per sample, while [Lemma 3.1](#) implies that computing the symplectic complement takes time

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

which dominates the runtime. □

Combining with [Algorithm 2](#) gives us the following.

Corollary 7.2. *Let $|\psi\rangle$ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of $|\psi\rangle$ as input, the combination of [Algorithms 2](#) and [3](#) outputs a classical description of $|\psi\rangle$ such that $d_{\text{tr}}(|\psi\rangle, |\hat{\psi}\rangle) \leq \varepsilon$ with probability at least $1 - \delta$. The algorithm uses*

$$\frac{32 \log(2/\delta) + 64n}{\varepsilon^2} + \frac{4}{3} N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}} + \frac{224}{9} \log(6/\delta)$$

samples and

$$O\left(\frac{n^2 \log(1/\delta) + n^3}{\varepsilon^2} + n^2 \cdot N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}}\right) + M_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}},$$

time.

Proof. By [Theorems 6.4](#) and [7.1](#), each algorithm succeeds with probability at least $1 - \delta/2$. By the union bound, the total failure probability is at most δ . □

8 Computational Difference Sampling

In this section, we introduce and develop the theory of *computational difference sampling* and connect it to the characteristic distribution p_ψ (Section 2.3).

Definition 8.1 (Computational difference sampling). Computational difference sampling a quantum state $|\psi\rangle$ corresponds to the following quantum measurement. Measure $|\psi\rangle$ in the computational basis to get outcome $x \in \mathbb{F}_2^n$, measure an additional copy of $|\psi\rangle$ in the computational basis to get another outcome $y \in \mathbb{F}_2^n$, and output $(x + y, 0^n) \in \mathcal{X}$. Let $r_\psi(a)$ denote the probability of sampling $a \in \mathcal{X}$ from this process.

One can view computational difference sampling as a distribution over $\mathbb{F}_2^n$, but for our purposes, it will be most convenient to identify $\mathbb{F}_2^n$ with $\mathcal{X}$, which corresponds to the set of Pauli- X Weyl operators.

Computational difference sampling only uses single-copy measurements and is the core primitive used in our single-copy learning algorithm. Although it is weaker than the entangled sampling techniques used in Algorithm 3 and prior learning algorithms based on Bell sampling [Mon17, GNW21, LC22, GIKL24], computational difference sampling still reveals some information about the unsigned stabilizer group of a quantum state, as we will show in this section. We begin by proving some basic properties about r_ψ .

Fact 8.2. Let $|\psi\rangle = \sum_{x \in \mathbb{F}_2^n} \alpha_x |x\rangle$ be an n -qubit quantum state, and let $a \in \mathbb{F}_2^n$. Then

$$r_\psi(a, 0^n) = \sum_{x \in \mathbb{F}_2^n} |\alpha_x|^2 |\alpha_{x+a}^2|.$$

Proof. It is a basic fact in probability theory that the sum of two independent random variables is the convolution of their individual distributions. By the Born rule, we observe $x \in \mathbb{F}_2^n$ with probability $|\alpha_x|^2$ upon measuring $|\psi\rangle$ in the computational basis. Therefore, the probability of computational difference sampling $(a, 0^n) \in \mathcal{X}$ is precisely

$$r_\psi(a, 0^n) = \sum_{x \in \mathbb{F}_2^n} |\alpha_x|^2 |\alpha_{x+a}^2|,$$

which is the convolution of the Born distribution with itself. Equivalently, this is the probability of sampling two strings x and y whose difference is a . $\square$

Next, we relate r_ψ to p_ψ (defined in Section 2.3). To do so, we need the following proposition, which relates sums of basis state projections to sums of Pauli operators.

Proposition 8.3. For any $a \in \mathbb{F}_2^n$,

$$\sum_{x \in \mathbb{F}_2^n} |x\rangle\langle x| \otimes |x+a\rangle\langle x+a| = \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} (-1)^{a \cdot x} Z^x \otimes Z^x.$$

Proof.

$$\begin{aligned}
\sum_{x \in \mathbb{F}_2^n} |x\rangle\langle x| \otimes |x+a\rangle\langle x+a| &= \sum_{x \in \mathbb{F}_2^n} |x\rangle\langle x| \otimes X^a |x\rangle\langle x| X^a \\
&= (I \otimes X^a) \left(\sum_{x \in \mathbb{F}_2^n} |x\rangle\langle x| \otimes |x\rangle\langle x| \right) (I \otimes X^a) \\
&= \frac{1}{2^n} (I \otimes X^a) \left(\sum_{x \in \mathbb{F}_2^n} Z^x \otimes Z^x \right) (I \otimes X^a) \quad (\text{Proposition 4.3}) \\
&= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} Z^x \otimes X^a Z^x X^a \\
&= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} (-1)^{a \cdot x} Z^x \otimes Z^x. \quad \square
\end{aligned}$$

The formal relation between r_ψ and p_ψ is stated below. Intuitively, this shows that r_ψ is obtained from p_ψ by summing over all possible shifts by a Pauli-Z operator.

Proposition 8.4. *Let $|\psi\rangle$ be an n -qubit pure state, and let $a \in \mathbb{F}_2^n$. Then*

$$r_\psi(a, 0^n) = \sum_{w \in \mathbb{F}_2^n} p_\psi(a, w) = \sum_{x \in a \times \mathbb{F}_2^n} p_\psi(x).$$

Proof. We can always write $|\psi\rangle = \sum_{x \in \mathbb{F}_2^n} \alpha_x |x\rangle$.

$$\begin{aligned}
\sum_{x \in \mathbb{F}_2^n} p_\psi(a, x) &= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} \langle \psi | i^{a \cdot x} X^a Z^x | \psi \rangle^2 \\
&= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} (-1)^{a \cdot x} \langle \psi | X^a Z^x | \psi \rangle^2 \\
&= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} (-1)^{a \cdot x} \langle \psi |^{\otimes 2} (X^a Z^x \otimes X^a Z^x) | \psi \rangle^{\otimes 2} \\
&= \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} (-1)^{a \cdot x} \langle \psi |^{\otimes 2} (X^a \otimes X^a) (Z^x \otimes Z^x) | \psi \rangle^{\otimes 2} \\
&= \sum_{x \in \mathbb{F}_2^n} \langle \psi |^{\otimes 2} (X^a \otimes X^a) (|x\rangle\langle x| \otimes |x+a\rangle\langle x+a|) | \psi \rangle^{\otimes 2} \quad (\text{Proposition 8.3}) \\
&= \sum_{x \in \mathbb{F}_2^n} \langle \psi |^{\otimes 2} (|x+a\rangle\langle x| \otimes |x\rangle\langle x+a|) | \psi \rangle^{\otimes 2} \\
&= \sum_{x \in \mathbb{F}_2^n} |\alpha_x|^2 |\alpha_{x+a}|^2 \\
&= r_\psi(a, 0^n). \quad (\text{Fact 8.2}) \quad \square
\end{aligned}$$

The duality property of p_ψ with respect to the symplectic product, as captured by Proposition 2.17, extends to r_ψ as a consequence of Proposition 8.4.

Corollary 8.5. *Given a subspace $H \subseteq \mathcal{X}$, then*

$$\sum_{x \in H} r_\psi(x) = \frac{|H + \mathcal{Z}|}{2^n} \sum_{x \in (H + \mathcal{Z})^\perp} p_\psi(x).$$

Proof.

$$\begin{aligned}
\sum_{x \in H} r_\psi(x) &= \sum_{x \in H + \mathcal{Z}} p_\psi(x) && \text{(Proposition 8.4)} \\
&= \frac{|H + \mathcal{Z}|}{2^n} \sum_{x \in (H + \mathcal{Z})^\perp} p_\psi(x) && \text{(Proposition 2.17)} \quad \square
\end{aligned}$$

Next, we relate computational difference samples of $|\psi\rangle$ to the Pauli- Z strings that stabilize $|\psi\rangle$ modulo phase. In effect, [Corollary 8.6](#) tells us that learning the support of r_ψ allows us to learn the Pauli- Z operators that stabilize $|\psi\rangle$ (up to a ± 1 phase).

Corollary 8.6. *For $x \in \mathcal{X}$, $r_\psi(x) > 0$ implies $x \in (\text{Weyl}(|\psi\rangle) \cap \mathcal{Z})^\perp$.*

Proof. Set $H = (\text{Weyl}(|\psi\rangle) \cap \mathcal{Z})^\perp \cap \mathcal{X}$ in [Corollary 8.5](#). Then, by [Lemma 2.11](#),

$$\begin{aligned}
(H + \mathcal{Z})^\perp &= H^\perp \cap \mathcal{Z} \\
&= (\text{Weyl}(|\psi\rangle) \cap \mathcal{Z} + \mathcal{X}) \cap \mathcal{Z} \\
&= \text{Weyl}(|\psi\rangle) \cap \mathcal{Z} \\
&\subseteq \text{Weyl}(|\psi\rangle).
\end{aligned}$$

It follows that:

$$\begin{aligned}
\sum_{x \in H} r_\psi(x) &= \frac{|H + \mathcal{Z}|}{2^n} \sum_{x \in (H + \mathcal{Z})^\perp} p_\psi(x) && \text{(Corollary 8.5)} \\
&= \frac{|H + \mathcal{Z}|}{2^n} \cdot \frac{|(H + \mathcal{Z})^\perp|}{2^n} && ((H + \mathcal{Z})^\perp \subseteq \text{Weyl}(|\psi\rangle)) \\
&= 1 && \text{(Fact 2.8)}
\end{aligned}$$

We conclude that the probability of sampling something outside of H is zero. $\square$

9 Learning the Unsigned Stabilizer Group Using Single-Copy Measurements

In [Section 7](#), we gave an algorithm to learn a subspace S satisfying the conditions of [Theorem 6.4](#) using entangled measurements. Here, we present an algorithm that learns such a subspace S using only single-copy measurements. We begin by stating the algorithm.

Algorithm 4: Approximating $\text{Weyl}(|\psi\rangle)$ using single-copy measurements

Input: Black-box access to copies of $|\psi\rangle$ and $\varepsilon, \delta \in (0, 1)$

Promise: $|\psi\rangle$ has stabilizer dimension at least $n - t$

Output: A subspace $S \subset \mathbb{F}_2^{2n}$ of dimension at least $n - t$ such that

$$\sum_{x \in S} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n} \text{ with probability at least } 1 - \delta$$

- 1 Set $m_{\text{Clifford}} \leftarrow 2(2^{t+1} + 1)(n + \log(2/\delta))$.
 - 2 **for** $i \leftarrow 0$ **to** m_{Clifford} **do**
 - 3 Use the algorithm in [Proposition 2.4](#) to sample a random Clifford circuit C_i ,
and apply C_i to $\frac{32n}{\varepsilon^2}(n + \log(2m_{\text{Clifford}}/\delta))$ copies of $|\psi\rangle$. Define $|\psi_i\rangle := C_i |\psi\rangle$.
 - 4 Computational difference sample $|\psi_i\rangle$ to draw $\frac{16n}{\varepsilon^2}(n + \log(2m_{\text{Clifford}}/\delta))$
samples from r_{ψ_i} .
 - 5 Let $\widehat{H}_i \subseteq \mathcal{X}$ be the subspace spanned by the computational difference samples.
Use Gaussian elimination to compute generators of $\widehat{H}_i$.
 - 6 Using [Lemma 3.1](#), compute $S := \sum_{i=1}^{m_{\text{Clifford}}} C_i^\dagger \left((\widehat{H}_i + \mathcal{Z})^\perp \right)$ as the sum of subspaces
over $\mathbb{F}_2^{2n}$.¹⁵
 - 7 **return** S
-

Let us briefly explain the workings of the algorithm at an intuitive level. The idea is to use computational difference sampling to learn about $\text{Weyl}(|\psi\rangle)$ via the support of r_ψ . [Corollary 8.6](#) says that computational difference sampling allows us to learn the generators of $\text{Weyl}(|\psi\rangle)$ that are also contained in $\mathcal{Z}$. Of course, unless $\text{Weyl}(|\psi\rangle) \subseteq \mathcal{Z}$, we will not learn a complete set of generators for $\text{Weyl}(|\psi\rangle)$. However, if we apply a random Clifford circuit C to $|\psi\rangle$, then $\text{Weyl}(C|\psi\rangle) \cap \mathcal{Z}$ will be nontrivial with some probability, and therefore reveal additional generators of $\text{Weyl}(|\psi\rangle)$. We repeat this process many times until we are confident that we have learned a complete set of generators for $\text{Weyl}(|\psi\rangle)$. Doing so requires more samples when $\text{Weyl}(|\psi\rangle)$ is small, because then $\text{Weyl}(C|\psi\rangle) \cap \mathcal{Z}$ is more likely to be empty—this is why m_{Clifford} grows exponentially in t .

In practice, we will not be able to guarantee that we exactly learn the support of r_ψ . This also means that we cannot guarantee that we learn $\text{Weyl}(|\psi\rangle)$ exactly. Instead, we show that the output of [Algorithm 4](#) satisfies the conditions in [Theorem 6.4](#).

9.1 The Output Has Large Dimension

The goal of this section is to show that the output of [Algorithm 4](#) has dimension at least $n - t$. Since $\text{Weyl}(|\psi\rangle)$ naturally satisfies this property by assumption, if we can show that the output of [Algorithm 4](#) is a superset of $\text{Weyl}(|\psi\rangle)$, then this output must also have large dimension.

We start by showing that the probability of learning a new element of $\text{Weyl}(|\psi\rangle)$ (should one exist) is not *too* difficult when a uniformly random Clifford circuit is applied. In what follows, think of A as $\text{Weyl}(|\psi\rangle)$, T as the subspace of A learned so far, and B as the set of Pauli- Z operators.

¹⁵By cleverly using the regular orthogonal complement over $\mathbb{F}_2^n$, computing the symplectic complement of subspaces of the form $H + \mathcal{Z}$ will offer a computational savings by a multiplicative factor of about 4 over naively using [Lemma 3.1](#). This is because, if we let $\overline{H}$ be the projection of H onto $\mathbb{F}_2^n$ obtained by ignoring the last n coordinates, and let $\overline{G}$ be the orthogonal complement of $\overline{H}$, then $(H + \mathcal{Z})^\perp = (\overline{H} \times \mathbb{F}_2^n)^\perp = 0^n \times \overline{G}$.

Lemma 9.1. *Let A be an isotropic subspace of dimension $n - k$, let $T \subset A$ be a proper subspace of A , let B be a fixed Lagrangian subspace (e.g. $\mathcal{Z}$), and let C be a uniformly random Clifford circuit. Then*

$$\Pr[C(A \setminus T) \cap B \neq \emptyset] \geq \frac{1}{2^{k+1} + 1}.$$

Proof. Without loss of generality, assume that T has dimension $n - k - 1$, because the probability we want to bound can only be larger for smaller subspaces T . Write $A = \langle x \rangle + T$ for some generator $x \in T^\perp$, so that $A \setminus T$ is equivalent to the coset $x + T$. Observe that

$$\begin{aligned} \Pr[C(A \setminus T) \cap B \neq \emptyset] &= \Pr[C(x + T) \cap B \neq \emptyset] \\ &= \Pr[C(x) \in C(T) + B]. \end{aligned}$$

We complete the proof by computing this quantity. We can think of sampling $C(T)$ and $C(x)$ by first choosing $C(T)$, and then sampling $C(x)$ conditioned on $C(T)$. Observe that in doing so, conditioned on a choice of $C(T)$, $C(x)$ is uniformly distributed in $C(T^\perp \setminus T)$. As a result, it holds that

$$\begin{aligned} \Pr[C(x) \in C(T) + B \mid C(T)] &= \frac{|(C(T) + B) \cap C(T^\perp \setminus T)|}{|C(T^\perp \setminus T)|} \\ &= \frac{|(C(T) + B) \cap C(T^\perp)| - |C(T)|}{|C(T^\perp)| - |C(T)|} \\ &= \frac{|(C(T) + B) \cap C(T)^\perp| - |C(T)|}{|C(T^\perp)| - |C(T)|} \\ &= \frac{2^n - 2^{n-k-1}}{2^{n+k+1} - 2^{n-k-1}} \quad (\text{Corollary 2.14}) \\ &= \frac{1}{2^{k+1} + 1}. \end{aligned}$$

Because this holds for every choice of $C(T)$, the lemma follows. $\square$

Now that we have an upper bound on the difficulty of learning a new generator of $\text{Weyl}(|\psi\rangle)$, we can analyze the number of Clifford circuits we need in total to guarantee that the output of [Algorithm 4](#) is a superset of $\text{Weyl}(|\psi\rangle)$. The analysis is very similar to [Lemma 2.3](#), though we are no longer interested in approximating the support of a distribution.

Lemma 9.2. *Let $|\psi\rangle$ be a quantum state with stabilizer dimension $n - k$. Consider sampling m random Clifford circuits $C_1, \dots, C_m$, and let $S_i = \text{Weyl}(|\psi\rangle) \cap C_i^\dagger(\mathcal{Z})$. If*

$$m \geq 2(2^{k+1} + 1)(n + \log(1/\delta)),$$

then with probability at least $1 - \delta$, $\sum_{i=1}^m S_i = \text{Weyl}(|\psi\rangle)$.

Proof. Let $T_i = \sum_{j=1}^i S_j$, with the convention that we define T_0 to be the trivial subspace. Define the indicator random variable X_i as

$$X_i = \begin{cases} 1 & S_i \not\subseteq T_{i-1} \text{ or } T_{i-1} \supseteq \text{Weyl}(|\psi\rangle) \\ 0 & \text{otherwise.} \end{cases}$$

Informally, $X_i = 1$ indicates a step at which the algorithm has made progress towards sampling a complete set of generators for $\text{Weyl}(|\psi\rangle)$ as part of T_m . It suffices to show that

with probability $1 - \delta$, $\sum_{i=1}^m X_i \geq n$, as this guarantees that $T_m \supseteq \text{Weyl}(|\psi\rangle)$, because $\text{Weyl}(|\psi\rangle)$ has dimension at most n by assumption.

Let $c = 2^{k+1} + 1$. If $T_{i-1} \supseteq \text{Weyl}(|\psi\rangle)$ then $\mathbf{E}[X_i] = 1 \geq \frac{1}{c}$. Otherwise, T_{i-1} is missing some member of $\text{Weyl}(|\psi\rangle)$, such that $T_{i-1} \cap \text{Weyl}(|\psi\rangle)$ has dimension at most $n - (k+1)$. By [Lemma 9.1](#) with $A = \text{Weyl}(|\psi\rangle)$, $T = T_{i-1}$, and $B = \mathcal{Z}$,

$$\begin{aligned} \mathbf{E}[X_i \mid T_{i-1} \not\supseteq \text{Weyl}(|\psi\rangle)] &= \mathbf{Pr}[A \cap C^\dagger(B) \not\subseteq T \mid T \subset A] \\ &= \mathbf{Pr}[(A \setminus T) \cap C^\dagger(B) \neq \emptyset \mid T \subset A] \\ &= \mathbf{Pr}[C(A \setminus T) \cap B \neq \emptyset \mid T \subset A] \\ &\geq \frac{1}{1 + 2^{k+1}} \\ &= \frac{1}{c}. \end{aligned}$$

So, $\mathbf{E}[X_i] \geq \frac{1}{c}$ regardless of what T_i is.

Let $\gamma = 1 - \frac{n}{cm}$. Then, by the multiplicative Chernoff bound ([Fact 2.1](#)), we have

$$\begin{aligned} \mathbf{Pr}\left[\sum_{i=1}^m X_i < n\right] &= \mathbf{Pr}\left[\sum_{i=1}^m X_i < (1 - \gamma)cm\right] \\ &\leq \exp\left(-\gamma^2 \frac{cm}{2}\right) \\ &= \exp\left(-\left(1 - \frac{2n}{cm} + \frac{n^2}{c^2 m^2}\right) \frac{cm}{2}\right) \\ &\leq \exp\left(-\left(1 - \frac{2n}{cm}\right) \frac{cm}{2}\right) \\ &= \exp\left(n - \frac{cm}{2}\right) \end{aligned} \tag{6}$$

Hence, choosing

$$m \geq \frac{2}{c} \left(n + \log \frac{1}{\delta}\right) = 2(2^{k+1} + 1) \left(n + \log \frac{1}{\delta}\right)$$

suffices to guarantee that [Eq. \(6\)](#) is at most δ . $\square$

9.2 The Output Has Large p_ψ -Mass

In this subsection, we argue that the subspace output by [Algorithm 4](#) satisfies [Condition \(ii\)](#) of [Theorem 6.4](#) (i.e., it accounts for a near-maximal p_ψ -mass). In [Line 6](#), we take the span of subspaces, with the hope that each one learns the support of r_{ψ_i} exactly. This would in turn allow us to *only* learn elements of $\text{Weyl}(|\psi\rangle)$ via [Corollary 8.6](#). However, it is possible that we will fail to do so in practice. For example, if $|\psi_i\rangle$ is close to (but not exactly) a computational basis state, then the support of r_{ψ_i} is nontrivial, even though computational difference sampling will almost always sample the identity Weyl operator.

Nevertheless, we argue that even if we do not learn the support of r_{ψ_i} exactly, we can still succeed. In particular, [Corollary 8.5](#) also tells us that learning a large fraction of the r_{ψ_i} -mass corresponds to learning a subspace that has near maximal p_ψ -mass (i.e., even though we pick up elements outside of $\text{Weyl}(|\psi\rangle)$, they still have large p_ψ value on average). Therefore, for [Line 6](#) to work, we need to show that the sums of subspaces of $\mathbb{F}_2^{2n}$ with near-maximal p_ψ -mass form a (potentially) larger subspace that is still near-maximal. The next lemma establishes this.

Lemma 9.3. Let $S_1, \dots, S_k$ be subspaces of $\mathbb{F}_2^{2n}$ such that $\sum_{x \in S_i} p_\psi(x) \geq (1 - \varepsilon_i) \frac{|S_i|}{2^n}$ for all S_i . Then $S = \sum_{i=1}^k S_i$ satisfies

$$\sum_{x \in S} p_\psi(x) \geq \left(1 - \sum_{i=1}^k \varepsilon_i\right) \frac{|S|}{2^n}.$$

Proof. We proceed via induction on k . The base case $k = 1$ follows by assumption. Now, take $T = \sum_{i=1}^{k-1} S_i$ such that

$$\sum_{x \in T} p_\psi(x) \geq \left(1 - \sum_{i=1}^{k-1} \varepsilon_i\right) \frac{|T|}{2^n}$$

by the inductive hypothesis. Then,

$$\begin{aligned} \sum_{x \in S} p_\psi(x) &= \frac{|T + S_k|}{2^n} \sum_{x \in (T + S_k)^\perp} p_\psi(x) && \text{(Proposition 2.17)} \\ &= \frac{|T + S_k|}{2^n} \sum_{x \in T^\perp \cap S_k^\perp} p_\psi(x) && \text{(Lemma 2.11)} \\ &= \frac{|T + S_k|}{2^n} \left(\sum_{x \in T^\perp} p_\psi(x) + \sum_{x \in S_k^\perp} p_\psi(x) - \sum_{x \in T^\perp \cup S_k^\perp} p_\psi(x) \right) && \text{(Inclusion-Exclusion)} \\ &\geq \frac{|T + S_k|}{2^n} \left(\frac{2^n}{|S|} \sum_{x \in T} p_\psi(x) + \frac{2^n}{|S_k|} \sum_{x \in S_k} p_\psi(x) - 1 \right) \\ &= \left(1 - \varepsilon_k - \sum_{i=1}^{k-1} \varepsilon_i\right) \frac{|T + S_k|}{2^n} = \left(1 - \sum_{i=1}^k \varepsilon_i\right) \frac{|S|}{2^n}. \end{aligned}$$

□

Let m be the number of Clifford circuits needed from [Section 9.1](#). For [Algorithm 4](#) to fulfill [Condition \(ii\)](#) of [Theorem 6.4](#), we need each of the $\widehat{H}_i$ from [Line 5](#) to only capture $x \in \mathbb{F}_2^{2n}$ such that the average squared expectation, i.e.,

$$\frac{2^n}{|\widehat{H}_i|} \sum_{x \in \widehat{H}_i} p_\psi(x) = \mathbf{E}_{x \sim \widehat{H}_i} [\langle \psi | W_x | \psi \rangle^2],$$

is very close to 1. Using [Lemma 9.3](#), one might naïvely expect to require an average squared expectation of $1 - \frac{\varepsilon}{m}$ for each $\widehat{H}_i$. However, we can provide a slight optimization by noting that, since all of the subspaces lie in $\mathbb{F}_2^{2n}$, we only need to ever consider at most $2n$ subspaces when [Lemma 9.3](#) is applied. As a result, we only require that the average squared expectation is at least $1 - \frac{\varepsilon}{2n}$ for each $\widehat{H}_i$.

Lemma 9.4. Let $S_1, \dots, S_m$ denote m (potentially non-unique) subspaces of $\mathbb{F}_2^{2n}$ and let $S = \sum_{i=1}^m S_i$. Suppose that for some $\varepsilon > 0$ and all S_i ,

$$\sum_{x \in S_i} p_\psi(x) \geq \left(1 - \frac{\varepsilon}{2n}\right) \frac{|S_i|}{2^n}.$$

Then,

$$\sum_{x \in S} p_\psi(x) \geq (1 - \varepsilon) \frac{|S|}{2^n}.$$

Proof. If $m \leq 2n$ then we can directly use [Lemma 9.3](#) (with $k := m$) and we are done. Otherwise, also by [Lemma 9.3](#) (with $k := 2n$), for every subset $I \subseteq [m]$ of $2n$ indices (i.e., $|I| = 2n$),

$$\sum_{x \in S_I} p_\psi(x) \geq \left(1 - \sum_{i \in I} \frac{\varepsilon}{2n}\right) \frac{|S_I|}{2^n} = (1 - \varepsilon) \frac{|S_I|}{2^n},$$

where $S_I := \sum_{i \in I} S_i$. Because all of these subspaces lie in $\mathbb{F}_2^{2n}$, we note that $S = S_I$ for some $I \subseteq [m]$ of cardinality $2n$.¹⁶ Therefore, it must also be the case that

$$\sum_{x \in S} p_\psi(x) \geq (1 - \varepsilon) \frac{|S|}{2^n}. \quad \square$$

Using [Corollary 8.5](#) and [Fact 2.24](#), we can translate [Lemma 9.4](#) into something more closely related to analyzing [Algorithm 4](#).

Proposition 9.5. *Let $C_1, \dots, C_m$ denote m (potentially non-unique) n -qubit Clifford circuits and let $|\psi_i\rangle := C_i |\psi\rangle$ for some fixed n -qubit quantum state $|\psi\rangle$. Given subspaces $H_1, \dots, H_m \subseteq \mathcal{X}$, suppose that for some $\varepsilon > 0$ and for all H_i ,*

$$\sum_{x \in H_i} r_{\psi_i}(x) \geq 1 - \frac{\varepsilon}{2n}.$$

Define $S := \sum_{i=1}^m C_i^\dagger ((H_i + \mathcal{Z})^\perp)$. Then

$$\sum_{x \in S} p_\psi(x) \geq (1 - \varepsilon) \frac{|S|}{2^n}.$$

Proof. For convenience, define $S_i = C_i^\dagger ((H_i + \mathcal{Z})^\perp)$. For each S_i , we have

$$\begin{aligned} \sum_{x \in S_i} p_\psi(x) &= \sum_{x \in C_i^\dagger ((H_i + \mathcal{Z})^\perp)} p_\psi(x) \\ &= \sum_{x \in (H_i + \mathcal{Z})^\perp} p_{\psi_i}(x) && \text{(Fact 2.24)} \\ &= \frac{2^n}{|H_i + \mathcal{Z}|} \sum_{x \in H_i} r_{\psi_i}(x) && \text{(Corollary 8.5)} \\ &= \frac{|(H_i + \mathcal{Z})^\perp|}{2^n} \sum_{x \in H_i} r_{\psi_i}(x) && \text{(Fact 2.8)} \\ &\geq \left(1 - \frac{\varepsilon}{2n}\right) \frac{|(H_i + \mathcal{Z})^\perp|}{2^n} \\ &= \left(1 - \frac{\varepsilon}{2n}\right) \frac{|S_i|}{2^n}, \end{aligned}$$

using in the last line the fact that the Clifford action on subspaces preserves dimension. Finally, [Lemma 9.4](#) tells us that

$$\sum_{x \in S} p_\psi(x) \geq (1 - \varepsilon) \frac{|S|}{2^n}$$

as desired. $\square$

¹⁶A more careful analysis involving [Lemma 4.2](#) reveals that S must be isotropic for $\varepsilon < \frac{1}{4}$, and therefore $\dim(S) \leq n$. Thus I can be taken to have cardinality n rather than $2n$. Since ε is always less than $\frac{1}{4}$ in [Algorithm 4](#), this halves the number of samples needed in [Line 4](#) and, therefore, [Algorithm 4](#) as a whole. For the sake of clarity, we do not go through this analysis.

9.3 Putting It All Together

We now have all of the tools to show that [Algorithm 4](#) succeeds at learning a subspace that satisfies [Conditions \(i\) and \(ii\)](#) of [Theorem 6.4](#).

Theorem 9.6. *Let $|\psi\rangle$ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of $|\psi\rangle$ as input, [Algorithm 4](#) succeeds at outputting a subspace S that satisfies the conditions of [Theorem 6.4](#) with probability at least $1 - \delta$. The algorithm uses*

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

samples and

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

time.

Proof. We will first show that [Condition \(i\)](#) is met. By setting $m_{\text{Clifford}} = 2(2^{t+1} + 1)(n + \log(2/\delta))$, [Lemma 9.2](#) tells us that the subspaces $S_i = \text{Weyl}(|\psi\rangle) \cap C_i^\dagger(\mathcal{Z})$ span $\text{Weyl}(|\psi\rangle)$, except with failure probability at most $\delta/2$. Assume that we succeed. Then [Corollary 8.6](#) shows that we are guaranteed to only sample $y \sim r_{\psi_i}$ such that $y \in H_i$, where

$$H_i := (\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z})^\perp \cap \mathcal{X}$$

Hence, $\widehat{H}_i \subseteq H_i$, and therefore

$$\begin{aligned} (\widehat{H}_i + \mathcal{Z})^\perp &\supseteq (H_i + \mathcal{Z})^\perp \\ &= H_i^\perp \cap \mathcal{Z} && \text{(\a href="#">Lemma 2.11)} \\ &= ((\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z}) + \mathcal{X}) \cap \mathcal{Z} && \text{(\a href="#">Lemma 2.11)} \\ &= (\mathcal{X} \cap \mathcal{Z}) + (\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z}) && \text{(\a href="#">Lemma 2.12)} \\ &= \text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z} \\ &= C(S_i). \end{aligned}$$

By the end of the m_{Clifford} random Clifford circuits, it follows that

$$S = \sum_{i=1}^{m_{\text{Clifford}}} C_i^\dagger((\widehat{H}_i + \mathcal{Z})^\perp) \supseteq \sum_{i=1}^{m_{\text{Clifford}}} S_i = \text{Weyl}(|\psi\rangle)$$

and therefore $\dim(S) \geq \dim(\text{Weyl}(|\psi\rangle)) \geq n - t$.

Moving on to [Condition \(ii\)](#), [Lemma 2.3](#) shows that if we draw $\frac{16n}{\varepsilon^2}(n + \log(2m_{\text{Clifford}}/\delta))$ computational difference samples per $|\psi_i\rangle$ then

$$\sum_{x \in \widehat{H}_i} r_{\psi_i}(x) \geq 1 - \frac{\varepsilon^2}{8n} \tag{7}$$

for all $\widehat{H}_i$ with failure probability at most $\delta/2$. Again, assume that we succeed. By [Proposition 9.5](#), it follows that

$$\sum_{x \in S} p_\psi(x) \geq (1 - \varepsilon^2/4) \frac{|S|}{2^n},$$

thus satisfying [Condition \(ii\)](#).

By the union bound, this means we satisfy [Conditions \(i\) and \(ii\)](#) with probability at least $1 - \delta$.

We next turn to the sample complexity analysis. For simplicity, let m_{Comp} be the number of computational difference samples taken in each part of the inner loop. Due to the fact that $\log(m_{\text{Clifford}}) = O\left(t + \log n + \frac{\log(2/\delta)}{n}\right)$, we find that

$$m_{\text{Comp}} := \frac{16n}{\varepsilon^2} (n + \log(2m_{\text{Clifford}}/\delta)) = O\left(\frac{n(n + \log(1/\delta))}{\varepsilon^2}\right),$$

which will greatly simplify our asymptotic analysis. Since a computational difference sample involves 2 copies of $|\psi\rangle$, the total number of copies of $|\psi\rangle$ needed is

$$2m_{\text{Clifford}}m_{\text{Comp}} = O\left(\frac{n(n + \log(1/\delta))^2 2^t}{\varepsilon^2}\right)$$

To analyze the runtime, we will first analyze the inner loop of [Algorithm 4](#). By [Proposition 2.4](#), we require $O(n^2)$ time to sample a random Clifford circuit in the first part of [Line 3](#). Applying C_i to $2m_{\text{Comp}}$ copies of $|\psi\rangle$ in the second half [Line 3](#) requires $O(m_{\text{Comp}}n^2)$ time. In [Line 4](#), we require $O(m_{\text{Comp}}n)$ time to make $O(m_{\text{Comp}})$ many n -qubit measurements, then apply the difference to pairs of measurements. Finally, it costs $O(m_{\text{Comp}}n^2)$ time to run [Lemma 3.1](#) in [Line 5](#). Thus, the inner loop requires time $O(m_{\text{Comp}}n^2)$, which becomes

$$O(m_{\text{Clifford}}m_{\text{Comp}}n^2) = O\left(\frac{n^3(n + \log(1/\delta))^2 2^t}{\varepsilon^2}\right)$$

over all iterations of the inner loop. This will turn out to be the dominating term.

For completeness, we will analyze [Line 6](#). Given the generators of $\widehat{H}_i$, it takes $O(n^3)$ time to compute $(\widehat{H}_i + \mathcal{Z})^\perp$ using [Lemma 3.1](#). By the size of the Clifford circuit from [Proposition 2.4](#), we also need $O(n^3)$ time to apply $C_i^\dagger$,¹⁷ for a total time of $O(m_{\text{Clifford}} \cdot n^3)$ to compute generators of each $C_i^\dagger(\widehat{H}_i + \mathcal{Z})^\perp$. All that is left is to compute $\sum C_i^\dagger(\widehat{H}_i + \mathcal{Z})^\perp$. Since each $C_i^\dagger(\widehat{H}_i + \mathcal{Z})^\perp$ can be described by at most $2n$ generators, the generators of S can be computed by running Gaussian elimination on a $2nm_{\text{Clifford}} \times 2n$ matrix over $\mathbb{F}_2$. This also takes $O(m_{\text{Clifford}} \cdot n^3) = O(n^3 \cdot 2^t \cdot (n + \log(1/\delta)))$ time. $\square$

Combining with [Algorithm 2](#) gives us the following.

Corollary 9.7. *Let $|\psi\rangle$ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of $|\psi\rangle$ as input, the combination of [Algorithms 2 and 4](#) output a classical description of $|\psi\rangle$ such that $\text{d}_{\text{tr}}(|\psi\rangle, |\psi\rangle) \leq \varepsilon$ with probability at least $1 - \delta$. The algorithm uses*

$$O\left(\frac{n(n + \log(1/\delta))^2 2^t}{\varepsilon^2} + N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}}\right)$$

single-copy samples and

$$O\left(\frac{n^3(n + \log(1/\delta))^2 2^t}{\varepsilon^2} + n^2 \cdot N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}}\right) + M_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}},$$

time.

¹⁷Since we are now classically simulating the action of $C_i^\dagger$ on a subspace of $\mathbb{F}_2^{2n}$ as part of our classical post-processing, it will require an extra $O(n)$ time per gate to apply $C_i^\dagger$.

Proof. By [Theorems 6.4](#) and [9.6](#), each algorithm succeeds with probability at least $1 - \delta/2$. By the union bound, the total failure probability is at most δ . $\square$

10 Open Problems

Given $|\psi\rangle$, a state that can be prepared with Clifford gates and $O(\log n)$ non-Clifford gates, our tomography algorithms efficiently learn a classical description of $|\psi\rangle$. However, our algorithms are not *proper* learners—in other words, the circuit that the algorithms output to approximate the state does not necessarily decompose into few non-Clifford gates.

Question 10.1. *Given copies of an n -qubit state $|\psi\rangle$ that is the output of a Clifford circuit and $O(\log n)$ non-Clifford gates, is it possible to efficiently construct a circuit U comprised of Clifford and $O(\log n)$ non-Clifford gates such that $U|0^n\rangle$ is ε -close to $|\psi\rangle$?*

We note that an efficient proper learning algorithm for stabilizer states is known (namely, just run [Algorithm 2](#) with either [Algorithm 3](#) or [Algorithm 4](#) and $t = 0$). We also note that this task becomes trivial if polynomially many non-Clifford gates are allowed in U because one can use [Algorithm 2](#) to produce C , $|x\rangle$, and $|\varphi\rangle$, and then construct a circuit with at most $2^{O(\log n)} = \text{poly}(n)$ general gates that outputs $|\varphi\rangle$ [[SBM06](#)].

As a subroutine, [Algorithm 2](#) uses pure state tomography to recover a classical description of a pure state, and therefore the copy and time complexities of our algorithm can be improved if faster pure state tomography algorithms are developed. Developing a pure state tomography algorithm that achieves the optimal $O(2^n/\varepsilon^2)$ copy and time complexities is an interesting and important direction for future work.

Observe that [Algorithm 3](#) learns an approximation to $\text{Weyl}(|\psi\rangle)$ using only $O(n)$ two-copy measurements, independent of the size of $\text{Weyl}(|\psi\rangle)$. By contrast, [Algorithm 4](#) requires $\text{poly}(n)\exp(t)$ samples to do the same when $\text{Weyl}(|\psi\rangle)$ has dimension $n - t$. This raises the natural question of whether an exponential dependence on t is necessary using unentangled measurements:

Question 10.2. *Are $\exp(t)$ many single-copy measurements necessary to learn a subspace S that satisfies the conditions of [Theorem 6.4](#)?*

Note that analogous exponential separations between entangled and single-copy measurements have been shown in other learning contexts [[HKP21](#), [CCHL22](#)]. Moreover, these separations hold for closely-related learning tasks that may be viewed as identifying the large coefficients of the Weyl expansion ([Definition 2.16](#)), but for mixed states. For this reason, we are optimistic that the techniques of [[CCHL22](#)] could be adapted to answer [Question 10.2](#) in the affirmative.

Finally, stabilizer dimension is a rather rigid notion, and the property of high stabilizer dimension excludes states that could be considered “near-stabilizer.” Could one extend the techniques in this paper and [[GIKL24](#)] to states of high *approximate* stabilizer dimension: that is, states which have high expectation with a large subgroup of Pauli operators?

Acknowledgements

We thank Scott Aaronson, Richard Kueng, Marcel Hinsche, Nick Hunter-Jones, Luowen Qian, and John Wright for helpful conversations.

SG, VI, DL are supported by Scott Aaronson’s Vannevar Bush Fellowship from the US Department of Defense, NSF QLCI Award OMA-2016245, a Simons Investigator Award,

and the Simons “It from Qubit” collaboration. WK is supported by an NDSEG Fellowship and acknowledges support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. DL is also supported by NSF award FET-2243659. VI is also supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research in Quantum Computing and an NSF Graduate Research Fellowship.

This work was done in part while SG, VI, and DL were visiting the Simons Institute for the Theory of Computing.

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>.

References

- [Aar20] Scott Aaronson. Shadow Tomography of Quantum States. *SIAM Journal on Computing*, 49(5):STOC18–368–STOC18–394, 2020. doi:10.1145/3188745.3188802. [p. 8]
- [Aar22] Scott Aaronson. Introduction to Quantum Information Science II Lecture Notes, May 2022. URL: <https://www.scottaaronson.com/qisii.pdf>. [p. 6]
- [AB08] A. Angelow and M. C. Batoni. About Heisenberg Uncertainty Relation (by E. Schrodinger), 2008. [arXiv:quant-ph/9903100](https://arxiv.org/abs/quant-ph/9903100). [p. 15]
- [ABDY23] Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. Optimal Algorithms for Learning Quantum Phase States. In *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023)*, volume 266 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:24, 2023. doi:10.4230/LIPIcs.TQC.2023.3. [p. 7]
- [ABNO22] Srinivasan Arunachalam, Sergey Bravyi, Chinmay Nirkhe, and Bryan O’Gorman. The Parametrized Complexity of Quantum Verification. In *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, volume 232 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:18, 2022. doi:10.4230/LIPIcs.TQC.2022.3. [p. 5]
- [AG04] Scott Aaronson and Daniel Gottesman. Improved Simulation of Stabilizer Circuits. *Physical Review A*, 70(5), 2004. doi:10.1103/physreva.70.052328. [pp. 5, 7, 17]
- [AG08] Scott Aaronson and Daniel Gottesman. Identifying Stabilizer States, 2008. <https://pirsa.org/08080052>. [pp. 6, 7]

- [AG23] Scott Aaronson and Sabee Grewal. Efficient Tomography of Non-Interacting-Fermion States. In *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023)*, volume 266 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:18, 2023. doi:[10.4230/LIPIcs.TQC.2023.12](https://doi.org/10.4230/LIPIcs.TQC.2023.12). [pp. 7, 24, 51]
- [AR19] Scott Aaronson and Guy N. Rothblum. Gentle Measurement of Quantum States and Differential Privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 322–333, 2019. doi:[10.1145/3313276.3316378](https://doi.org/10.1145/3313276.3316378). [p. 8]
- [BBB⁺23] Sara Bartolucci, Patrick Birchall, Hector Bombin, Hugo Cable, Chris Dawson, Mercedes Gimeno-Segovia, Eric Johnston, Konrad Kieling, Naomi Nickerson, Mihir Pant, et al. Fusion-based quantum computation. *Nature Communications*, 14(1):912, 2023. doi:[10.1038/s41467-023-36493-1](https://doi.org/10.1038/s41467-023-36493-1). [p. 3]
- [BBC⁺19] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, 2019. doi:[10.22331/q-2019-09-02-181](https://doi.org/10.22331/q-2019-09-02-181). [pp. 3, 7]
- [BCG13] K. Banaszek, M. Cramer, and D. Gross. Focus on quantum tomography. *New Journal of Physics*, 15(12):125020, 2013. doi:[10.1088/1367-2630/15/12/125020](https://doi.org/10.1088/1367-2630/15/12/125020). [p. 3]
- [BCHK20] Michael Beverland, Earl Campbell, Mark Howard, and Vadym Kliuchnikov. Lower bounds on the non-Clifford resources for quantum computations. *Quantum Science and Technology*, 5(3):035009, 2020. doi:[10.1088/2058-9565/ab8963](https://doi.org/10.1088/2058-9565/ab8963). [p. 15]
- [BG16] Sergey Bravyi and David Gosset. Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates. *Phys. Rev. Lett.*, 116:250501, 2016. doi:[10.1103/PhysRevLett.116.250501](https://doi.org/10.1103/PhysRevLett.116.250501). [p. 7]
- [BK05] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A*, 71:022316, 2005. doi:[10.1103/PhysRevA.71.022316](https://doi.org/10.1103/PhysRevA.71.022316). [p. 3]
- [BM99] Dagmar Bruß and Chiara Macchiavello. Optimal state estimation for d -dimensional quantum systems. *Physics Letters A*, 253(5-6):249–251, 1999. doi:[10.1016/S0375-9601\(99\)00099-7](https://doi.org/10.1016/S0375-9601(99)00099-7). [pp. 3, 5]
- [BM21] Sergey Bravyi and Dmitri Maslov. Hadamard-Free Circuits Expose the Structure of the Clifford Group. *IEEE Transactions on Information Theory*, 67(7):4546–4563, 2021. doi:[10.1109/TIT.2021.3081415](https://doi.org/10.1109/TIT.2021.3081415). [p. 11]
- [BO21] Costin Bădescu and Ryan O’Donnell. Improved Quantum Data Analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 1398–1411, 2021. doi:[10.1145/3406325.3451109](https://doi.org/10.1145/3406325.3451109). [p. 8]
- [BS19] Zvika Brakerski and Omri Shmueli. (Pseudo) Random Quantum States with Binary Phase. In *Theory of Cryptography*, 2019. doi:[10.1007/978-3-030-36030-6_10](https://doi.org/10.1007/978-3-030-36030-6_10). [p. 4]
- [CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential Separations Between Learning With and Without Quantum Memory.

- In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 574–585. IEEE Computer Society, 2022. doi:[10.1109/FOCS52979.2021.00063](https://doi.org/10.1109/FOCS52979.2021.00063). [pp. 4, 39]
- [CHL⁺23] Sitan Chen, Brice Huang, Jerry Li, Allen Liu, and Mark Sellke. When Does Adaptivity Help for Quantum State Learning? . In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 391–404, Los Alamitos, CA, USA, November 2023. IEEE Computer Society. doi:[10.1109/FOCS57990.2023.00029](https://doi.org/10.1109/FOCS57990.2023.00029). [p. 3]
- [CLL24] Nai-Hui Chia, Ching-Yi Lai, and Han-Hsuan Lin. Efficient learning of t -doped stabilizer states with single-copy measurements. *Quantum*, 8:1250, February 2024. doi:[10.22331/q-2024-02-12-1250](https://doi.org/10.22331/q-2024-02-12-1250). [pp. 9, 45]
- [CPF⁺10] Marcus Cramer, Martin B. Plenio, Steven T. Flammia, Rolando Somma, David Gross, Stephen D. Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature Communications*, 1(1):1–7, 2010. doi:[10.1038/ncomms1147](https://doi.org/10.1038/ncomms1147). [p. 7]
- [Fan57] Ugo Fano. Description of States in Quantum Mechanics by Density Matrix and Operator Techniques. *Reviews of Modern Physics*, 29(1):74, 1957. doi:[10.1103/RevModPhys.29.74](https://doi.org/10.1103/RevModPhys.29.74). [p. 3]
- [FBaK21] Daniel Stilck França, Fernando G.S L. Brandão, and Richard Kueng. Fast and Robust Quantum State Tomography from Few Basis Measurements. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:13, 2021. doi:[10.4230/LIPIcs.TQC.2021.7](https://doi.org/10.4230/LIPIcs.TQC.2021.7). [pp. 24, 51]
- [FvdG99] C.A. Fuchs and J. van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999. doi:[10.1109/18.761271](https://doi.org/10.1109/18.761271). [p. 51]
- [GIKL23] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. Low-Stabilizer-Complexity Quantum States Are Not Pseudorandom. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 64:1–64:20, 2023. doi:[10.4230/LIPIcs.ITCS.2023.64](https://doi.org/10.4230/LIPIcs.ITCS.2023.64). [pp. 4, 5, 8]
- [GIKL24] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. Improved Stabilizer Estimation via Bell Difference Sampling. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, page 1352–1363, New York, NY, USA, 2024. Association for Computing Machinery. doi:[10.1145/3618260.3649738](https://doi.org/10.1145/3618260.3649738). [pp. 4, 5, 6, 8, 10, 14, 15, 16, 29, 39, 50]
- [GNW21] David Gross, Sepehr Nezami, and Michael Walter. 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):1325–1393, 2021. doi:[10.1007/s00220-021-04118-7](https://doi.org/10.1007/s00220-021-04118-7). [pp. 5, 6, 8, 13, 15, 29, 50]
- [GOL24] Andi Gu, Salvatore F. E. Oliviero, and Lorenzo Leone. Doped stabilizer states in many-body physics and where to find them. *Phys. Rev. A*, 110:062427, Dec 2024. doi:[10.1103/PhysRevA.110.062427](https://doi.org/10.1103/PhysRevA.110.062427). [p. 3]

- [Got97] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, May 1997. [doi:10.7907/rzr7-dt72](#). [p. 11]
- [Got06] Daniel Gottesman. Quantum Error Correction and Fault Tolerance. In Jean-Pierre Francoise, Gregory L. Naber, and Tsou Sheung Tsun, editors, *Encyclopedia of Mathematical Physics*, pages 196–201. Academic Press, Oxford, 2006. [doi:10.1016/B0-12-512666-2/00273-X](#). [p. 20]
- [HBC⁺22] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, and Jarrod R. McClean. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022. [doi:10.1126/science.abn7293](#). [p. 4]
- [HG24] Dominik Hangleiter and Michael J. Gullans. Bell Sampling from Quantum Circuits. *Phys. Rev. Lett.*, 133:020601, Jul 2024. [doi:10.1103/PhysRevLett.133.020601](#). [pp. 8, 9, 10]
- [HHJ⁺17] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-Optimal Tomography of Quantum States. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017. [doi:10.1109/TIT.2017.2719044](#). [p. 3]
- [HIN⁺23] Marcel Hinsche, Marios Ioannou, Alexander Nietner, Jonas Haferkamp, Yihui Quek, Dominik Hangleiter, Jean-Pierre Seifert, Jens Eisert, and Ryan Sweke. One T Gate Makes Distribution Learning Hard. *Phys. Rev. Lett.*, 130:240602, 2023. [doi:10.1103/PhysRevLett.130.240602](#). [p. 8]
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020. [doi:10.1038/s41567-020-0932-7](#). [pp. 7, 8]
- [HKP21] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Information-Theoretic Bounds on Quantum Advantage in Machine Learning. *Phys. Rev. Lett.*, 126:190505, 2021. [doi:10.1103/PhysRevLett.126.190505](#). [pp. 4, 39]
- [HMMH⁺23] Jonas Haferkamp, Felipe Montealegre-Mora, Markus Heinrich, Jens Eisert, David Gross, and Ingo Roth. Efficient Unitary Designs with a System-Size Independent Number of Non-Clifford Gates. *Communications in Mathematical Physics*, 397:995–1041, 2023. [doi:10.1007/s00220-022-04507-6](#). [p. 3]
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom Quantum States. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference*, pages 126–152. Springer, 2018. [doi:10.1007/978-3-319-96878-0_5](#). [p. 3]
- [Kni04] E. Knill. Fault-Tolerant Postselected Quantum Computation: Schemes, 2004. [arXiv:quant-ph/0402171](#). [p. 3]
- [KRT17] Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low Rank Matrix Recovery From Rank One Measurements. *Applied and Computational Harmonic Analysis*, 42(1):88–116, 2017. [doi:10.1016/j.acha.2015.07.007](#). [p. 3]

- [KS14] Robert Koenig and John A. Smolin. How to efficiently select an arbitrary Clifford group element. *Journal of Mathematical Physics*, 55(12):122202, 2014. doi:[10.1063/1.4903507](https://doi.org/10.1063/1.4903507). [p. [11](#)]
- [LC22] Ching-Yi Lai and Hao-Chung Cheng. Learning Quantum Circuits of Some T Gates. *IEEE Transactions on Information Theory*, 68(6):3951–3964, 2022. doi:[10.1109/TIT.2022.3151760](https://doi.org/10.1109/TIT.2022.3151760). [pp. [6](#), [7](#), [8](#), [29](#)]
- [LOH24] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. Learning t-doped stabilizer states. *Quantum*, 8:1361, May 2024. doi:[10.22331/q-2024-05-27-1361](https://doi.org/10.22331/q-2024-05-27-1361). [pp. [8](#), [9](#)]
- [LOLH24] Lorenzo Leone, Salvatore F. E. Oliviero, Seth Lloyd, and Alioscia Hamma. Learning efficient decoders for quasichaotic quantum scramblers. *Phys. Rev. A*, 109:022429, Feb 2024. doi:[10.1103/PhysRevA.109.022429](https://doi.org/10.1103/PhysRevA.109.022429). [pp. [5](#), [8](#), [11](#), [20](#)]
- [Mon17] Ashley Montanaro. Learning stabilizer states by Bell sampling, 2017. arXiv:[1707.04012](https://arxiv.org/abs/1707.04012). [pp. [5](#), [6](#), [7](#), [8](#), [29](#)]
- [MPS03] G. Mauro D’Ariano, Matteo G.A. Paris, and Massimiliano F. Sacchi. Quantum Tomography. *Advances in Imaging and Electron Physics*, 128:205–308, 2003. doi:[10.1016/S1076-5670\(03\)80065-4](https://doi.org/10.1016/S1076-5670(03)80065-4). [p. [3](#)]
- [NC02] Michael A. Nielsen and Isaac Chuang. Quantum Computation and Quantum Information, 2002. doi:[10.1017/CB09780511976667](https://doi.org/10.1017/CB09780511976667). [p. [11](#)]
- [OW16] Ryan O’Donnell and John Wright. Efficient Quantum Tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pages 899–912, 2016. doi:[10.1145/2897518.2897544](https://doi.org/10.1145/2897518.2897544). [p. [3](#)]
- [QPG21] Hammam Qassim, Hakop Pashayan, and David Gosset. Improved upper bounds on the stabilizer rank of magic states. *Quantum*, 5:606, 2021. doi:[10.22331/q-2021-12-20-606](https://doi.org/10.22331/q-2021-12-20-606). [p. [7](#)]
- [RLCK19] Patrick Rall, Daniel Liang, Jeremy Cook, and William Kretschmer. Simulation of qubit quantum circuits via Pauli propagation. *Phys. Rev. A*, 99:062337, 2019. doi:[10.1103/PhysRevA.99.062337](https://doi.org/10.1103/PhysRevA.99.062337). [p. [7](#)]
- [SBM06] Vivek V. Shende, Stephen S. Bullock, and Igor L. Markov. Synthesis of quantum-logic circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25(6):1000–1010, 2006. doi:[10.1109/TCAD.2005.855930](https://doi.org/10.1109/TCAD.2005.855930). [p. [39](#)]
- [Sch30] Erwin Schrödinger. The Uncertainty Relations in Quantum Mechanics. Zum Heisenbergschen Unschärfeprinzip. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse*, 14:296–303, 1930. [p. [15](#)]
- [Shi03] Yaoyun Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computing. *Quantum Info. Comput.*, 3(1):84–92, jan 2003. doi:<https://doi.org/10.26421/QIC3.1-7>. [pp. [3](#), [11](#)]
- [SXL⁺17] Chao Song, Kai Xu, Wuxin Liu, Chui-Ping Yang, Shi-Biao Zheng, Hui Deng, Qiwei Xie, Keqiang Huang, Qiujiang Guo, Libo Zhang, et al. 10-qubit entanglement and parallel logic operations with a superconducting circuit. *Physical Review Letters*, 119(18):180511, 2017. doi:[10.1103/PhysRevLett.119.180511](https://doi.org/10.1103/PhysRevLett.119.180511). [p. [3](#)]

- [VDB21] Ewout Van Den Berg. A simple method for sampling random Clifford operators. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 54–59, 2021. doi:10.1109/QCE52317.2021.00021. [p. 11]
- [Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. [p. 51]

A A Faster Algorithm From One Round of Adaptivity

Both Algorithms 3 and 4 are non-adaptive, in the sense that the measurements performed on $|\psi\rangle$ are independent of the outcomes of previous measurements, and so they can be measured in parallel. In this appendix, we explain how to improve both the sample and time complexities of Algorithm 4—the algorithm that only uses single-copy measurements—by using one round of adaptivity.¹⁸ Specifically, our modified algorithm will have two phases, where the second phase depends on the measurement outcomes of the first phase, while continuing to only use single-copy measurements. We credit the concurrent work by Chia, Lai, and Lin [CLL24] for the idea behind this modification.

The main idea of the algorithm is reasonably simple: suppose we knew an isotropic subspace A that contains $\text{Weyl}(|\psi\rangle)$. Then by Lemma 3.2, we could find a Clifford circuit C such that $\text{Weyl}(C|\psi\rangle) \subseteq \mathcal{Z}$, by mapping A to a subspace of $\mathcal{Z}$. But notice that if all of the stabilizers of $C|\psi\rangle$ are contained in $\mathcal{Z}$, then we can approximately learn $\text{Weyl}(C|\psi\rangle) = C(\text{Weyl}(|\psi\rangle))$ from computational difference samples of $C|\psi\rangle$, as a consequence of Corollary 8.5.

A priori, it is not clear what this approach buys us, because learning an isotropic subspace A that contains $\text{Weyl}(|\psi\rangle)$ could be roughly as hard as learning a subspace S that satisfies the conditions of Theorem 6.4. As we shall see, the advantage is that in Line 3 of Algorithm 4, we can get away with taking a number of computational difference samples that scales linearly in n , rather than quadratically in n/ε .

We are now ready to state the algorithm.

¹⁸A similar modification can be done to Algorithm 3. However, the modification will not yield any asymptotic improvements on the sample and time complexities, like it will for our modification to Algorithm 4.

Algorithm 5: Approximating $\text{Weyl}(|\psi\rangle)$ using adaptive single-copy measurements

Input: Black-box access to copies of $|\psi\rangle$ and $\varepsilon, \delta \in (0, 1)$

Promise: $|\psi\rangle$ has stabilizer dimension at least $n - t$

Output: A subspace $S \subset \mathbb{F}_2^{2n}$ of dimension at least $n - t$ such that

$$\sum_{x \in S} p_\psi(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n} \text{ with probability at least } 1 - \delta$$

- 1 Set $m_{\text{Clifford}} \leftarrow 2(2^{t+1} + 1)(n + \log(3/\delta))$.
 - 2 **for** $i \leftarrow 0$ **to** m_{Clifford} **do**
 - 3 Use the algorithm in [Proposition 2.4](#) to sample a random Clifford circuit C_i , and apply C_i to $16(n + \log(3m_{\text{Clifford}}/\delta)) + 2$ copies of $|\psi\rangle$. Define $|\psi_i\rangle := C_i |\psi\rangle$.
 - 4 Computational difference sample $|\psi_i\rangle$ to draw $8(n + \log(3m_{\text{Clifford}}/\delta)) + 1$ samples from r_{ψ_i} .
 - 5 Let $\widehat{H}_i \subseteq \mathcal{X}$ be the subspace spanned by the computational difference samples. Use Gaussian elimination to compute generators of $\widehat{H}_i$.
 - 6 Using [Lemma 3.1](#), compute $\widehat{A} := \sum_{i=1}^{m_{\text{Clifford}}} C_i^\dagger ((\widehat{H}_i + \mathcal{Z})^\perp)$.¹⁹
 - 7 Apply [Lemma 3.2](#) to find the Clifford circuit $C_{\mathcal{Z}}$ that maps $\widehat{A}$ to a subspace of $\mathcal{Z}$.
 - 8 Obtain $\frac{16(n + \log(3/\delta))}{\varepsilon^2}$ copies of $|\psi'\rangle := C_{\mathcal{Z}} |\psi\rangle$.
 - 9 Computational difference sample all copies of $|\psi'\rangle$ to obtain $\frac{8(n + \log(3/\delta))}{\varepsilon^2}$ samples from $r_{\psi'}$.
 - 10 Let $\widehat{H} \subseteq \mathcal{X}$ be the subspace spanned by the computational difference samples. Use [Lemma 3.1](#) to compute generators of $(\widehat{H} + \mathcal{Z})^\perp$.
 - 11 **return** $S := C_{\mathcal{Z}}^\dagger ((\widehat{H} + \mathcal{Z})^\perp)$.
-

Observe that [Lines 1 to 6](#) are the same as in [Algorithm 4](#), except that the number of computational difference samples taken in [Line 4](#) is smaller. The algorithm is adaptive only because the Clifford circuit obtained in [Line 7](#) depends on earlier measurement outcomes.

We now work to show the correctness of [Algorithm 5](#). We first require a generalization of [Lemma 4.2](#). This ensures that $\widehat{A}$ in [Line 6](#) will be isotropic, even when [Condition \(ii\)](#) is not met.

Lemma A.1. *Let S be a subspace of $\mathbb{F}_2^{2n}$ such that*

$$\sum_{x \in S} p_\psi(x) > \frac{3}{4} \frac{|S|}{2^n}.$$

Then $S \subseteq \langle M \rangle$, where $M := \{x \in \mathbb{F}_2^{2n} : 2^n p_\psi(x) > 1/2\}$.

Proof. In the proof of [Lemma 4.2](#), we showed that $S = \langle A \rangle$, where $A := \{x \in S : 2^n p_\psi(x) > 1/2\}$. Since $A \subseteq M$, we clearly have $S = \langle A \rangle \subseteq \langle M \rangle$, which proves the lemma. $\square$

Corollary A.2. *Let H be a subspace of $\mathcal{X}$ such that*

$$\sum_{x \in H} r_\psi(x) > \frac{3}{4}.$$

Then $(H + \mathcal{Z})^\perp \subseteq \langle M \rangle$, where $M := \{x \in \mathbb{F}_2^{2n} : 2^n p_\psi(x) > \frac{1}{2}\}$.

¹⁹Again, by smartly using the regular orthogonal complement over $\mathbb{F}_2^n$, computing the symplectic complement of subspaces of the form $H + \mathcal{Z}$ will offer a computational savings by a multiplicative factor of about 4 over naively using [Lemma 3.1](#).

Proof. By [Corollary 8.5](#) and [Fact 2.8](#),

$$\sum_{x \in (H + \mathcal{Z})^\perp} p_\psi(x) = \frac{2^n}{|H + \mathcal{Z}|} \sum_{x \in H} r_\psi(x) > \frac{3}{4} \frac{|(H + \mathcal{Z})^\perp|}{2^n}.$$

[Lemma A.1](#) then shows that $(H + \mathcal{Z})^\perp \subseteq \langle M \rangle$. $\square$

We now have all of the tools to show that [Algorithm 5](#) succeeds at learning a subspace that satisfies [Conditions \(i\) and \(ii\)](#) of [Theorem 6.4](#).

Theorem A.3. *Let $|\psi\rangle$ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of $|\psi\rangle$ as input, [Algorithm 5](#) succeeds at outputting a subspace S that satisfies all conditions of [Theorem 6.4](#) with probability at least $1 - \delta$. The algorithm uses*

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

samples and

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

time.

Proof. The algorithm runs in two stages: [Lines 1 to 6](#) constitute the first non-adaptive stage, while [Lines 7 to 11](#) make up the second stage based on $\hat{A}$.

We will first show that at the end of stage one, $\text{Weyl}(|\psi\rangle) \subseteq \hat{A}$ and $\hat{A}$ is isotropic, with high probability. By setting $m_{\text{Clifford}} = 2(2^{t+1} + 1)(n + \log(3/\delta))$, [Lemma 9.2](#) tells us that the subspaces $S_i = \text{Weyl}(|\psi\rangle) \cap C_i^\dagger(\mathcal{Z})$ span $\text{Weyl}(|\psi\rangle)$, except with failure probability at most $\delta/3$. Assume that we succeed. Then [Corollary 8.6](#) shows that we are guaranteed to only sample $y \sim r_{\psi_i}$ such that $y \in H_i$, where

$$H_i := (\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z})^\perp \cap \mathcal{X}$$

Hence, $\hat{H}_i \subseteq H_i$, and therefore

$$\begin{aligned} (\hat{H}_i + \mathcal{Z})^\perp &\supseteq (H_i + \mathcal{Z})^\perp \\ &= H_i^\perp \cap \mathcal{Z} && \text{(\a href="#">Lemma 2.11)} \\ &= ((\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z}) + \mathcal{X}) \cap \mathcal{Z} && \text{(\a href="#">Lemma 2.11)} \\ &= (\mathcal{X} \cap \mathcal{Z}) + (\text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z}) && \text{(\a href="#">Lemma 2.12)} \\ &= \text{Weyl}(|\psi_i\rangle) \cap \mathcal{Z} \\ &= C_i(S_i). \end{aligned}$$

By the end of the m_{Clifford} random Clifford circuits, it follows that

$$\hat{A} = \sum_{i=1}^{m_{\text{Clifford}}} C_i^\dagger \left((\hat{H}_i + \mathcal{Z})^\perp \right) \supseteq \sum_{i=1}^{m_{\text{Clifford}}} S_i = \text{Weyl}(|\psi\rangle).$$

Next, we argue that $\hat{A}$ is isotropic. [Lemma 2.3](#) shows that if we draw

$$8(n + \log(3m_{\text{Clifford}}/\delta)) + 1$$

computational difference samples per $|\psi_i\rangle$ then

$$\sum_{x \in \widehat{H}_i} r_{\psi_i} > \frac{3}{4} \quad (8)$$

for all $\widehat{H}_i$ with failure probability at most $\delta/3$. Again, assume that we succeed. By [Corollary A.2](#), it follows that each $C_i^\dagger \left((\widehat{H}_i + \mathcal{Z})^\perp \right) \subseteq \langle M \rangle$ where $M := \{x \in \mathbb{F}_2^{2n} : 2^n p_\psi(x) > \frac{1}{2}\}$. The addition of these subspaces $\sum_i C_i^\dagger \left((\widehat{H}_i + \mathcal{Z})^\perp \right)$ then must also lie in $\langle M \rangle$. By [Fact 2.21](#), $\langle M \rangle$ is isotropic, and so too must be $\widehat{A}$.

In stage two, we need to prove that [Conditions \(i\) and \(ii\)](#) are met by the output S . Since we assume that stage one succeeded, $\widehat{A}$ is isotropic. Therefore, we can apply [Lemma 3.2](#) to find a Clifford circuit such that $C_{\mathcal{Z}}(\widehat{A}) \subseteq \mathcal{Z}$, as is done in [Line 7](#). Since $\text{Weyl}(|\psi\rangle) \subseteq \widehat{A}$, it also follows that $\text{Weyl}(|\psi'\rangle) \subseteq \mathcal{Z}$.

We now start the second round of computational difference sampling in [Lines 8 and 9](#). By [Lemma 2.3](#), $\widehat{H} \subseteq \mathcal{X}$ from [Line 10](#) will be such that

$$\sum_{x \in \widehat{H}} r_{\psi'}(x) \geq 1 - \frac{\varepsilon^2}{4}$$

with probability at least $\delta/3$. Once again, assume that we succeed. It follows that

$$\begin{aligned} \sum_{x \in S} p_\psi(x) &= \sum_{x \in (\widehat{H} + \mathcal{Z})^\perp} p_{\psi'}(x) && \text{(Fact 2.24)} \\ &= \frac{2^n}{|\widehat{H} + \mathcal{Z}|} \sum_{x \in \widehat{H}} r_{\psi'}(x) && \text{(Corollary 8.5)} \\ &\geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{2^n}{|\widehat{H} + \mathcal{Z}|} \\ &= \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n}, && \text{(Fact 2.8)} \end{aligned}$$

thus satisfying [Condition \(ii\)](#).

Finally, [Corollary 8.6](#) guarantees that we only sample from

$$H := (\text{Weyl}(|\psi'\rangle) \cap \mathcal{Z})^\perp \cap \mathcal{X} = \text{Weyl}(|\psi'\rangle)^\perp \cap \mathcal{X} = (\text{Weyl}(|\psi'\rangle) + \mathcal{X})^\perp.$$

Hence, $\widehat{H} \subseteq H$, and therefore

$$\begin{aligned} (\widehat{H} + \mathcal{Z})^\perp &\supseteq (H + \mathcal{Z})^\perp \\ &= H^\perp \cap \mathcal{Z} && \text{(Lemma 2.11)} \\ &= (\text{Weyl}(|\psi'\rangle) + \mathcal{X}) \cap \mathcal{Z} \\ &= (\mathcal{X} \cap \mathcal{Z}) + \text{Weyl}(|\psi'\rangle) && \text{(Lemma 2.12)} \\ &= \text{Weyl}(|\psi'\rangle). \end{aligned}$$

It follows that $S \supseteq \text{Weyl}(|\psi\rangle)$, and since $\dim(\text{Weyl}(|\psi\rangle)) \geq n - t$, we conclude that S satisfies [Condition \(i\)](#).

By the union bound, this means we satisfy [Conditions \(i\) and \(ii\)](#) with probability at least $1 - \delta$.

We next turn to the sample complexity analysis. For simplicity, let m_{Comp} be the number of computational difference samples taken in each part of the inner loop. Because $\log(m_{\text{Clifford}}) = O\left(t + \log n + \frac{\log(3/\delta)}{n}\right)$, we find that

$$m_{\text{Comp}} := 8(n + \log(3m_{\text{Clifford}}/\delta)) + 1 = O(n + \log(1/\delta)),$$

which will greatly simplify our asymptotic analysis. Since a computational difference sample involves 2 copies of $|\psi\rangle$, the total number of copies of $|\psi\rangle$ needed in stage 1 is

$$2m_{\text{Clifford}}m_{\text{Comp}} = O\left((n + \log(1/\delta))^2 2^t\right).$$

In stage 2 we use $\frac{16(n+\log(3/\delta))}{\varepsilon^2}$ samples, so the total number of samples is

$$O\left((n + \log(1/\delta))^2 2^t + \frac{n + \log(1/\delta)}{\varepsilon^2}\right).$$

To analyze the runtime, we will first analyze the inner loop of [Algorithm 5](#). By [Proposition 2.4](#), we require $O(n^2)$ time to sample a random Clifford circuit in the first part of [Line 1](#). Applying C_i to $2m_{\text{Comp}}$ copies of $|\psi\rangle$ in the second half [Line 1](#) requires $O(m_{\text{Comp}}n^2)$ time. In [Line 2](#), we require $O(m_{\text{Comp}}n)$ time to make $O(m_{\text{Comp}})$ many n -qubit measurements, then apply the difference to pairs of measurements. Finally, it costs $O(m_{\text{Comp}}n^2)$ time to run Gaussian elimination in [Line 3](#) to find the generators of each $\widehat{H}_i$. Thus, the inner loop requires time $O(m_{\text{Comp}}n^2)$, which becomes

$$O(m_{\text{Clifford}}m_{\text{Comp}}n^2) = O\left(n^2(n + \log(1/\delta))^2 2^t\right)$$

over all iterations of the inner loop. This will be one of the dominating terms.

Analyzing from [Line 6](#) and onwards, it takes $O(m_{\text{Clifford}}n^3)$ time to apply [Lemma 3.1](#) to each of the $\widehat{H}_i + \mathcal{Z}$. It similarly requires $O(m_{\text{Clifford}}n^3)$ time to apply $C_i^\dagger$ to each $(\widehat{H}_i + \mathcal{Z})^\perp$. [Lemma 3.2](#) requires $O(n^2)$ time in [Line 7](#) to produce a circuit of size at most $O(n^2)$. As a result, the computational difference sampling in [Lines 8](#) and [9](#) requires $O\left(\frac{n^2(n+\log(1/\delta))}{\varepsilon^2}\right)$ time. Finally, it takes $O\left(\frac{n^2(n+\log(1/\delta))}{\varepsilon^2}\right)$ time to perform [Lemma 3.1](#) one last time in [Line 10](#). This is the second dominating term. $\square$

B Generalization to Mixed States

We show that [Algorithms 2](#) and [3](#) remain correct when run on mixed states. For a mixed state ρ , we define $\text{Weyl}(\rho) := \{x \in \mathbb{F}_2^{2n} : \text{tr}(W_x \rho) = \pm 1\}$. It is easy to see that $\text{Weyl}(\rho)$ is still an isotropic subspace, just as it is for pure states. The stabilizer dimension of ρ is then the dimension of the subspace $\text{Weyl}(\rho)$. We will prove the following theorem.

Theorem B.1. *Given $n, t \in \mathbb{N}$, there is an algorithm that uses $\text{poly}(n, 2^t, 1/\varepsilon)$ time and $\text{poly}(n, 1/\varepsilon)$ copies of an n -qubit mixed state ρ and learns ρ to trace distance ε with high probability, promised that ρ has stabilizer dimension at least $n - t$.*

We must recall the trace distance and fidelity between mixed quantum states. The trace distance between two mixed states ρ, σ is $d_{\text{tr}}(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1$, where, for a matrix A , $\|A\|_1$ is the sum of the absolute values of the singular values of A . The fidelity between ρ and σ is $F(\rho, \sigma) := \text{tr}(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$.

B.1 Bell Difference Sampling Mixed States

To prove [Theorem B.1](#), we must develop the theory of Bell difference sampling for mixed states. As we will show, many properties that hold for pure states remain true for mixed states. Indeed, we believe many of the results/algorithms in the literature that are based on Bell difference sampling can be generalized to mixed states if one desires to.

Recall from [\[GNW21, Equation 3.7\]](#) that one can describe Bell difference sampling as a projective measurement. In particular, upon Bell difference sampling a mixed state ρ , one observes $a \in \mathbb{F}_2^{2n}$ with probability $\text{tr}(\Pi_a \rho^{\otimes 4})$, where Π_a is defined as

$$\Pi_a := \frac{1}{4^n} \sum_{x \in \mathbb{F}_2^{2n}} (-1)^{[a,x]} W_x^{\otimes 4}. \quad (9)$$

Hence, we define $q_\rho(x) := \text{tr}(\Pi_x \rho^{\otimes 4})$, the distribution induced by the Bell difference sampling measurement.

We find it convenient to define $p_\rho(x) := 2^{-n} \text{tr}(W_x \rho)^2$ in analogy with p_ψ . However, many properties of p_ψ fail to carry over to p_ρ . For example, p_ρ is no longer a probability distribution, so one can no longer express q_ρ as the convolution of p_ρ with itself as a result. Note that for any $x \in \mathbb{F}_2^{2n}$, we still have that $0 \leq p_\rho(x) \leq 2^{-n}$.

A crucial property of q_ψ is that it exhibits a duality property ([Proposition 2.17](#)). We prove that this property holds for q_ρ as well.

Theorem B.2 (Generalization of [Proposition 2.17](#)). *Let $A \subseteq \mathbb{F}_2^{2n}$ be a subspace. Then*

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

Proof. We have

$$\begin{aligned} \sum_{a \in A} q_\rho(a) &= \sum_{a \in A} \text{tr}(\Pi_a \rho^{\otimes 4}) \\ &= \sum_{a \in A} \sum_{x \in \mathbb{F}_2^{2n}} (-1)^{[a,x]} p_\rho(x)^2 \\ &= \sum_{x \in \mathbb{F}_2^{2n}} p_\rho(x)^2 \sum_{a \in A} (-1)^{[a,x]} \\ &= |A| \sum_{x \in \mathbb{F}_2^{2n}} p_\rho(x)^2 1_{x \in A^\perp} \\ &= |A| \sum_{x \in A^\perp} p_\rho(x)^2. \end{aligned}$$

Above, the first three lines follow from [Eq. \(9\)](#), the linearity of the trace, and some basic arithmetic manipulations. The fourth line follows from the fact that for any subspace $T \subseteq \mathbb{F}_2^{2n}$ and a fixed $x \in \mathbb{F}_2^{2n}$,

$$\sum_{a \in T} (-1)^{[a,x]} = |T| \cdot 1_{x \in T^\perp}.$$

See, for example, [\[GIKL24, Lemma 2.11\]](#) for a proof of this fact. Readers familiar with Boolean Fourier analysis will recognize this as a basic fact about summing over (symplectic) Fourier characters. $\square$

The duality theorem for q_ρ allows us to prove the following two statements. The first states that the support of q_ρ is contained in $\text{Weyl}(\rho)^\perp$, generalizing [\[GIKL24, Corollary 4.4\]](#). The second bounds the q_ρ -mass on a subspace A in terms of p_ρ .

Lemma B.3. For any quantum mixed state ρ , the support of q_ρ is contained in $\text{Weyl}(\rho)^\perp$.

Proof.

$$\begin{aligned} \sum_{x \in \text{Weyl}(\rho)^\perp} q_\rho(x) &= |\text{Weyl}(\rho)^\perp| \sum_{x \in \text{Weyl}(\rho)} p_\rho(x)^2 && (\text{Theorem B.2}) \\ &= |\text{Weyl}(\rho)^\perp| \cdot \frac{|\text{Weyl}(\rho)|}{4^n} \\ &= 1. \end{aligned}$$

The second line follows from that fact that for $x \in \text{Weyl}(\rho)$, $p_\rho(x)^2 = 4^{-n}$. $\square$

Proposition B.4. Let $A \subseteq \mathbb{F}_2^{2n}$ be a subspace. Then

$$\sum_{a \in A} q_\rho(a) \leq \frac{|A|}{2^n} \sum_{a \in A^\perp} p_\rho(a).$$

Proof.

$$\begin{aligned} \sum_{a \in A} q_\rho(a) &= |A| \sum_{x \in A^\perp} p_\rho(x)^2 && (\text{Theorem B.2}) \\ &\leq \frac{|A|}{2^n} \sum_{x \in A^\perp} p_\rho(x). && (p_\rho(x) \leq \frac{1}{2^n}) \end{aligned} \quad \square$$

B.2 Reduction to Finding a Heavy Subspace

We argue in [Theorem 6.4](#) that efficient tomography of large-stabilizer-dimension states reduces to finding a subspace of $\mathbb{F}_2^{2n}$ that isn't too large and has sufficiently large p_ψ mass. In this section, we prove that this reduction holds for mixed states as well. The formal statement of this follows.

Theorem B.5 (Generalization of [Theorem 6.4](#)). Let ρ be an n -qubit mixed state, and suppose there exists a subspace $S \subseteq \mathbb{F}_2^{2n}$ that satisfies the following conditions:

- (i) $\dim(S) \geq n - t$, and
- (ii) $\sum_{x \in S} p_\rho(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n}$.

Then, given S and copies of ρ as input, [Algorithm 2](#) outputs a classical description of $\hat{\rho}$ such that $d_{\text{tr}}(\rho, \hat{\rho}) \leq \varepsilon$ with probability at least $1 - \delta$ using $\frac{4}{3}N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \frac{224}{9} \log(3/\delta)$ copies of ρ , and $M_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + O\left(n^2 \left(N_{t, \frac{\varepsilon}{2}, \frac{\delta}{3}} + \log(1/\delta)\right)\right)$ time, while performing only single-copy measurements.

Remark B.6. In this appendix, $M_{n, \varepsilon, \delta}$ and $N_{n, \varepsilon, \delta}$ above are the same as in [Definition 6.1](#) except for mixed state tomography. That is, we let $M_{n, \varepsilon, \delta}$ and $N_{n, \varepsilon, \delta}$ denote the copy and time complexities, respectively, of state tomography of an n -qubit mixed state to trace distance ε and success probability at least $1 - \delta$. As for pure states, one can use the state tomography algorithms in [\[FBaK21\]](#) or [\[AG23, Section 5\]](#) to instantiate [Theorem B.5](#).

Proof. The only aspects of the proof of [Theorem 6.4](#) that use the purity of the input are the [Lemmas 4.2](#) and [4.6](#), and the identity $d_{\text{tr}}(|\psi\rangle, |\phi\rangle) = \sqrt{1 - F(|\psi\rangle, |\phi\rangle)}$. We have generalized the first two statements below to mixed states in [Lemmas B.7](#) and [B.9](#), and, for two mixed states ρ, σ , $d_{\text{tr}}(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)}$ [\[FvdG99, Wat18\]](#). Hence, the proof goes through with no loss in parameters. $\square$

To complete the proof of [Theorem B.5](#), it remains to generalize [Lemmas 4.2](#) and [4.6](#) to mixed states. We prove these generalizations in the remainder of this section.

Lemma B.7 (Generalization of [Lemma 4.2](#)). *Let S be a subspace of $\mathbb{F}_2^{2n}$ such that*

$$\sum_{x \in S} p_\rho(x) > \frac{3}{4} \frac{|S|}{2^n}.$$

Then S is isotropic.

Proof. Let $A := \{x \in S : 2^n p_\rho(x) > 1/2\}$. Because [Fact 2.20](#) holds for mixed states, it follows that [Fact 2.21](#) holds as well. Therefore, the rest of the argument in [Lemma 4.2](#) follows, because, aside from the above statements, the proof only uses facts about symplectic vector spaces. $\square$

The final two lemmas of this subsection generalize [Lemma 4.6](#) to mixed states.

Lemma B.8 (Generalization of [Lemma 4.5](#)). *Let $S = 0^{n+t} \times \mathbb{F}_2^{n-t}$, and suppose that*

$$\sum_{x \in S} p_\rho(x) \geq \frac{1 - \varepsilon}{2^t}.$$

Then there exists an $(n-t)$ -qubit computational basis state $|x\rangle$ and a t -qubit quantum state

$$\sigma := \frac{\text{tr}_B((I \otimes |x\rangle\langle x|)\rho)}{\text{tr}((I \otimes |x\rangle\langle x|)\rho)}, \quad (10)$$

such that the fidelity between $\sigma \otimes |x\rangle\langle x|$ and ρ is at least $1 - \varepsilon$. The partial trace in [Eq. \(10\)](#) is over the last $n - t$ qubits.²⁰

Proof. We can always write:

$$\rho = \sum_{y,z \in \mathbb{F}_2^{n-t}} \alpha_{y,z} \cdot \sigma_{y,z} \otimes |y\rangle\langle z|.$$

Here, the $\alpha_{z,z}$ are the diagonal elements of $\text{tr}_A(\rho)$, where A corresponds to the first t qubits, so they are nonnegative real numbers that sum to 1. The $\sigma_{y,z}$ are t -qubit mixed states.

Define $x := \arg\max_{z \in \mathbb{F}_2^{n-t}} \alpha_{z,z}$. Observe that $\text{tr}_B((I \otimes |x\rangle\langle x|)\rho) = \alpha_{x,x} \cdot \sigma_{x,x}$. We will argue that the fidelity between $\sigma_{x,x} \otimes |x\rangle\langle x|$ and ρ is equal to $\alpha_{x,x}$. Then we will argue that $\alpha_{x,x} \geq 1 - \varepsilon$. Hence, taking $\sigma := \sigma_{x,x}$ gives the desired state claimed in [Eq. \(10\)](#).

First, we prove that $F(\rho, \sigma_{x,x} \otimes |x\rangle\langle x|) = \alpha_{x,x}$. We have

$$\begin{aligned} F(\rho, \sigma_{x,x} \otimes |x\rangle\langle x|) &= \text{tr} \left(\sqrt{\sqrt{\sigma_{x,x}} \otimes |x\rangle\langle x| \cdot \left(\sum_{y,z \in \mathbb{F}_2^{n-t}} \alpha_{y,z} \cdot \sigma_{y,z} \otimes |y\rangle\langle z| \right) \cdot \sqrt{\sigma_{x,x}} \otimes |x\rangle\langle x|}} \right)^2 \\ &= \text{tr} \left(\sqrt{\sum_{y,z \in \mathbb{F}_2^{n-t}} \alpha_{y,z} \cdot (\sqrt{\sigma_{x,x}} \sigma_{y,z} \sqrt{\sigma_{x,x}}) \otimes (|x\rangle\langle x| y \rangle \langle z| x \rangle \langle x|)} \right)^2 \\ &= \text{tr} \left(\sqrt{\alpha_{x,x} \cdot \sigma_{x,x}^2 \otimes |x\rangle\langle x|} \right)^2 \\ &= \alpha_{x,x} \cdot \text{tr}(\sigma_{x,x} \otimes |x\rangle\langle x|)^2 \\ &= \alpha_{x,x}. \end{aligned}$$

²⁰This is to say that σ is obtained by postselecting on measuring the last $n - t$ qubits of ρ to be $|x\rangle$.

Next, we prove that $\alpha_{x,x} \geq 1 - \varepsilon$.

$$\begin{aligned}
1 - \varepsilon &\leq 2^t \sum_{x \in S} p_\rho(x) \\
&= \frac{1}{2^{n-t}} \sum_{z \in \mathbb{F}_2^{n-t}} \text{tr} (I \otimes W_z \otimes I \otimes W_z \cdot \rho^{\otimes 2}) \\
&= \sum_{z \in \mathbb{F}_2^{n-t}} \text{tr} ((I \otimes |z\rangle\langle z|) \rho)^2 && \text{(Proposition 4.3)} \\
&\leq \left(\max_{z \in \mathbb{F}_2^{n-t}} \text{tr} ((I \otimes |z\rangle\langle z|) \rho) \right) \cdot \sum_{z \in \mathbb{F}_2^{n-t}} \text{tr} ((I \otimes |z\rangle\langle z|) \rho) \\
&= \alpha_{x,x} \cdot \sum_{z \in \mathbb{F}_2^{n-t}} \alpha_{z,z} \\
&= \alpha_{x,x}.
\end{aligned}$$

On the fourth line, we used the fact that $\text{tr}((I \otimes |z\rangle\langle z|)\rho) = \alpha_{z,z}$. $\square$

Lemma B.9 (Generalization of Lemma 4.6). *Let S be an isotropic subspace of dimension $n - t$, and suppose that*

$$\sum_{x \in S} p_\rho(x) \geq \frac{1 - \varepsilon}{2^t}.$$

Then there exists a state $\hat{\rho}$ with $S \subseteq \text{Weyl}(\hat{\rho})$ such that the fidelity between $\hat{\rho}$ and ρ is at least $1 - \varepsilon$.

In particular, $\hat{\rho} = C^\dagger(\sigma \otimes |x\rangle\langle x|)C$, where $|x\rangle$ is an $(n - t)$ -qubit basis state,

$$\sigma = \frac{\text{tr}_B((I \otimes |x\rangle\langle x|)C\rho C^\dagger)}{\text{tr}((I \otimes |x\rangle\langle x|)C\rho C^\dagger)} \quad (11)$$

is a t -qubit quantum state, C is the Clifford circuit mapping S to $0^{n+t} \times \mathbb{F}_2^{n-t}$ described in Lemma 3.2, and the partial trace in Eq. (11) is over the last $n - t$ qubits.

Proof. Following the proof of Lemma 4.6, the Clifford circuit that maps S to $0^{n+t} \times \mathbb{F}_2^{n-t}$ permutes the p_ρ mass in the same way that it permutes the p_ψ mass for a pure state as per Fact 2.24. That is, for $\rho' := C\rho C^\dagger$,

$$p_{\rho'}(x) = \frac{1}{2^n} \text{tr} (W_x C \rho C^\dagger) = p_\rho(C^\dagger(x)).$$

The only remaining step of the proof that relies on a pure state is when Lemma 4.5 is used. Since we generalized Lemma 4.5 to hold for mixed states, via Lemma B.8, the proof goes through. $\square$

B.3 Efficient Tomography of Large-Stabilizer-Dimension Mixed States

We give the remaining details of our efficient tomography algorithm. The final missing piece is to show that one can learn a subspace satisfying the two conditions in Theorem B.5 by repeatedly Bell difference sampling copies of a mixed state.

Theorem B.10 (Generalization of Theorem 7.1). *Let ρ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of ρ , Algorithm 3 succeeds at outputting a*

subspace S that satisfies the conditions of [Theorem B.5](#) with probability at least $1 - \delta$ with no loss in parameters. In particular, the algorithm uses

$$\frac{32 \log(1/\delta) + 64n}{\varepsilon^2}$$

samples and

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

time.

Proof. By [Lemma B.3](#), the support of q_ρ is a subspace of dimension at most $n + t$, so $\dim S^\perp \leq n + t$ and therefore $\dim S \geq n - t$, by [Fact 2.8](#). This establishes [Condition \(i\)](#).

By [Lemma 2.3](#), except with probability at most δ , the Bell difference sampling phase of the algorithm in [Line 1](#) finds a subspace $S^\perp \subseteq \mathbb{F}_2^{2n}$ such that

$$\sum_{x \in S^\perp} q_\rho(x) \geq 1 - \frac{\varepsilon^2}{4}.$$

From [Proposition B.4](#), we know that

$$\frac{|S^\perp|}{2^n} \sum_{x \in S} p_\rho(x) \geq 1 - \frac{\varepsilon^2}{4},$$

and because $|S| \cdot |S^\perp| = 4^n$, we have

$$\sum_{x \in S} p_\rho(x) \geq \left(1 - \frac{\varepsilon^2}{4}\right) \frac{|S|}{2^n},$$

which establishes [Condition \(ii\)](#).

The running time analysis is the same as the analysis given in the proof of [Theorem 7.1](#). $\square$

We can now combine [Theorems B.5](#) and [B.10](#) to prove the main theorem of this appendix. We note that the generalization to mixed states goes through with no loss in parameters (except potentially for the parameters referenced in [Remark B.6](#)).

Corollary B.11 (Formal version of [Theorem B.1](#)). *Let ρ be an n -qubit quantum state with stabilizer dimension at least $n - t$. Given copies of ρ as input, the combination of [Algorithms 2](#) and [3](#) outputs a classical description of $\hat{\rho}$ such that $d_{\text{tr}}(\rho, \hat{\rho}) \leq \varepsilon$ with probability at least $1 - \delta$. The algorithm uses*

$$\frac{32 \log(2/\delta) + 64n}{\varepsilon^2} + \frac{4}{3} N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}} + \frac{224}{9} \log(6/\delta)$$

samples and

$$O\left(\frac{n^2 \log(1/\delta) + n^3}{\varepsilon^2} + n^2 \cdot N_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}}\right) + M_{t, \frac{\varepsilon}{2}, \frac{\delta}{6}},$$

time.

Proof. By [Theorems B.5](#) and [B.10](#), each algorithm succeeds with probability at least $1 - \delta/2$. By the union bound, the total failure probability is at most δ . $\square$

As a final remark, we note that [Corollary B.11](#) learns the input state in *fidelity*, which, for mixed states, is a stronger result than trace distance tomography. This is apparent in the Fuchs-van de Graaf inequality $d_{\text{tr}}(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)}$, which says that if ρ and σ are close in fidelity, then they are also close in trace distance. Hence, if one chooses a fidelity state tomography algorithm in [Theorem B.5](#), the final output of [Corollary B.11](#) will be close to the input in fidelity.