

Triply efficient shadow tomography

Robbie King^{*†} David Gosset^{*‡§} Robin Kothari^{*} Ryan Babbush^{*}

May 1, 2024

Abstract

Given copies of a quantum state ρ , a shadow tomography protocol aims to learn all expectation values from a fixed set of observables, to within a given precision ϵ . We say that a shadow tomography protocol is *triply efficient* if it is sample- and time-efficient, and only employs measurements that entangle a constant number of copies of ρ at a time. The classical shadows protocol based on random single-copy measurements is triply efficient for the set of local Pauli observables. This and other protocols based on random single-copy Clifford measurements can be understood as arising from fractional colorings of a graph G that encodes the commutation structure of the set of observables. Here we describe a framework for two-copy shadow tomography that uses an initial round of Bell measurements to reduce to a fractional coloring problem in an induced subgraph of G with bounded clique number. This coloring problem can be addressed using techniques from graph theory known as *chi-boundedness*. Using this framework we give the first triply efficient shadow tomography scheme for the set of local fermionic observables, which arise in a broad class of interacting fermionic systems in physics and chemistry. We also give a triply efficient scheme for the set of all n -qubit Pauli observables. Our protocols for these tasks use two-copy measurements, which is necessary: sample-efficient schemes are provably impossible using only single-copy measurements. Finally, we give a shadow tomography protocol that compresses an n -qubit quantum state into a $\text{poly}(n)$ -sized classical representation, from which one can extract the expected value of any of the 4^n Pauli observables in $\text{poly}(n)$ time, up to a small constant error.

1 Introduction

In this paper we discuss protocols for shadow tomography, as introduced by Aaronson [Aar18], specialized to the case of Pauli observables. Let

$$\mathcal{P}^{(n)} = \{\mathbb{1}, X, Y, Z\}^{\otimes n} \tag{1}$$

be the set of n -qubit Pauli operators and consider a subset $S \subseteq \mathcal{P}^{(n)}$. Let ρ be an unknown n -qubit quantum state. Given copies of ρ , we would like to learn the expectation values $\text{Tr}(P\rho)$ to precision ϵ for every $P \in S$.

^{*}Google Quantum AI, Venice, CA, USA

[†]California Institute of Technology, Pasadena, CA, USA

[‡]IQC and Department of Combinatorics and Optimization, University of Waterloo, Canada

[§]Perimeter Institute for Theoretical Physics, Waterloo, Canada

Definition 1 (Pauli shadow tomography). *The shadow tomography task for a set $S \subseteq \mathcal{P}^{(n)}$ is as follows. We are given copies of an unknown n -qubit state ρ , and our goal is to output estimates y_P such that, with high probability¹ we have $|y_P - \text{Tr}(P\rho)| \leq \epsilon$ for all $P \in S$.*

One can use a very naive tomography protocol to perform this task. For a given Pauli $P \in S$, if we measure its value $O((\log |S|)/\epsilon^2)$ times (using one copy of ρ for each measurement), then we can ensure that the sample mean is within ϵ of $\text{Tr}(\rho P)$ with probability at least $1 - 0.01/|S|$. If we follow this procedure for each of the $|S|$ Paulis in S , the union bound guarantees that, with high probability they will all be ϵ -close to their true values. This algorithm uses $O((|S| \log |S|)/\epsilon^2)$ copies of the unknown state ρ and is computationally efficient.

Remarkably, Aaronson described a protocol for shadow tomography of any set of bounded observables (such as Pauli observables) that uses exponentially fewer copies of the unknown state than the naive algorithm [Aar18, AR19, BKL⁺19]. The best known scaling of the number of samples for learning m general observables is $O(n(\log^2 m)/\epsilon^4)$ [BO21]. However, the general shadow tomography schemes suffer from two major caveats: they are explicitly exponential in computational runtime (even when the number of observables m scales polynomially with n), and they require entangled measurements on many copies of the unknown state ρ at a time.

These caveats can be avoided for certain restricted sets of observables such as low-weight Pauli operators. For k -local Paulis with $k = O(1)$, there are simple and computationally efficient protocols to learn m observables with $O((\log m)/\epsilon^2)$ single-copy measurements [CW20, EHF19, BMBO20, JKMN20, HKP20]. The classical shadows framework [HKP20] provides a broader family of learning protocols that can also handle other sets of non-Pauli observables with single-copy measurements, notably including rank-1 observables which are relevant to fidelity estimation. However, classical shadows and other single-copy learning strategies become inefficient for higher weight Pauli operators.

To go beyond low-weight Paulis one can use a shadow tomography protocol developed in [HKP21] which learns *any* subset S of Pauli operators using $O((\log |S|)/\epsilon^4)$ copies of ρ , and $\text{poly}(|S|, n, 1/\epsilon)$ runtime. The protocol proceeds in two stages. In the first stage—*learning magnitudes*—one computes estimates of the magnitudes $|\text{Tr}(\rho P)|$ to within $\epsilon/4$ error (say), for all Paulis $P \in S$. Remarkably, this can be achieved efficiently using only $O((\log |S|)/\epsilon^4)$ two-copy measurements using the well-known Bell sampling procedure [Mon17]. It is based on measuring copies of $\rho \otimes \rho$ in the basis which simultaneously diagonalizes the operators $P \otimes P$ for all $P \in \mathcal{P}^{(n)}$. In the second stage—*learning signs*—one computes the signs of all Paulis $P \in S$ that were estimated to have nonnegligible magnitude in the first stage. The learning signs protocol from Ref. [HKP21] proceeds by a sequence of “gentle measurements” on $O((\log |S|)/\epsilon^2)$ copies of ρ .

The requirement to perform joint entangled measurements on many copies of ρ is a significant drawback. If the state ρ is prepared by a real-world quantum computer, it may be prohibitive to require a quantum memory whose size is many times larger than the quantum system of interest. It is natural to ask if we can achieve sample- and computational efficiency while using measurements on only a few copies of ρ at a time. We call a shadow tomography protocol with these features *triply efficient*.

¹Throughout this paper, we use “with high probability” to mean with probability at least 99%, say.

Triply efficient shadow tomography

1. *Sample efficiency*: The number of samples scales as $\text{poly}(\log |S|, 1/\epsilon)$.
2. *Computational efficiency*: The classical and quantum computation is $\text{poly}(|S|, n, 1/\epsilon)$.
3. *Few-copy measurements*: The algorithm uses joint measurements on a constant number of copies of ρ (ideally 1 or 2).

In addition to the above, it would be reasonable to ask that the total quantum memory used by the shadow tomography algorithm is $O(n)$; all of the protocols we propose in this work satisfy this stronger requirement.

Is there a triply efficient shadow tomography protocol? While this question is well-posed for arbitrary subsets of observables $S \subseteq \mathcal{P}^{(n)}$, we restrict our attention to three subsets that are practically motivated and representative of the complexity of the problem:

1. $\mathcal{P}_k^{(n)}$: The set of k -local Pauli operators on n qubits, where $k = O(1)$.
2. $\mathcal{F}_k^{(n)}$: The set of k -body fermionic operators on n fermionic modes, where $k = O(1)$.
3. $\mathcal{P}^{(n)}$: The set of all Pauli operators on n qubits.

Note that a triply efficient shadow tomography protocol for the set of all Paulis does not directly give one for the set of k -local Paulis or k -body fermionic operators, even though they are subsets of the set of all Paulis. This is because the complexity requirements are functions of $|S|$. Thus a triply efficient protocol for the set of all Paulis can use more samples and time than we allow for smaller subsets of Paulis.

As we describe below, k -local Pauli operators and k -body fermionic operators arise in a variety of applications in many-body physics and quantum chemistry and are one of the most common algorithmic applications of shadow tomography. The set of all Pauli operators is of exponential size in n , and perhaps not as practically relevant, but seems important to study nonetheless since it represents a general Pauli learning task. Surprising tomography algorithms are possible in this case, including an algorithm we describe that, for any constant error ϵ , compresses the output (of size 4^n) into a polynomial-sized description from which an ϵ -estimate of the expected value of any Pauli observable can be extracted efficiently.

This brings us to the main question that guided this work:

Do there exist triply efficient shadow tomography protocols for these observables?

In fact, the schemes based on random single-copy measurements described in Refs. [HKP20, CW20, BMBO20, JKMN20, EHF19] already achieve triply efficient shadow tomography for the set $\mathcal{P}_k^{(n)}$ of k -local Pauli operators with $k = O(1)$. In this paper we present triply efficient shadow tomography algorithms for the set $\mathcal{F}_k^{(n)}$ of k -body fermionic operators for $k = O(1)$, and the set $\mathcal{P}^{(n)}$ of all n -qubit Pauli operators. Furthermore, our algorithms only use Clifford measurements on 2 copies of ρ at a time.

We will see that it is impossible to perform sample-efficient shadow tomography using only single-copy measurements for the set of k -body fermionic operators ([Theorem 3](#)), as

Observables	Triply efficient shadow tomography?
k -local Pauli operators	Yes [HKP20 , CW20 , BMBO20 , JKMN20 , EHF19]
k -body fermionic operators	Yes (Theorem 10)
All Pauli operators	Yes (Theorem 7)

Table 1: A qualitative summary of triply efficient shadow tomography. Cells in blue are new results in this paper.

well as for the set of all Paulis [[CCHL22](#)]. Taken together, our protocols and the single-copy lower bounds demonstrate that two-copy measurements are necessary and sufficient for Pauli and fermionic shadow tomography. Further, we see a striking difference between local Paulis, for which single-copy measurements suffice, and local fermionic operators where entangled measurements are necessary.

1.1 Local observables

In order to describe our results, let us now define the sets of local observables that are relevant to qubit and fermionic systems. Let $|P|$ denote the Pauli-weight of an operator $P \in \mathcal{P}^{(n)}$, i.e., the number of qubits on which it acts nontrivially. For example, $|X \otimes \mathbb{1} \otimes Y \otimes \mathbb{1} \otimes Z| = 3$.

A broad class of quantum many-body systems that arise in condensed matter physics are described by systems of spins with $k = O(1)$ particle interactions. Such systems are described by a Hamiltonian operator which can be expressed as a sum of operators from the set

$$\mathcal{P}_k^{(n)} = \{P \in \mathcal{P}^{(n)} : |P| = k\} \quad (k\text{-local Pauli operators}) \quad (2)$$

of all weight- k Pauli observables. Note that $|\mathcal{P}_k^{(n)}| = 3^k \binom{n}{k}$ and $\log |\mathcal{P}_k^{(n)}| = O(k \log n)$. Learning all Paulis in the set $\mathcal{P}_k^{(n)}$ is quite useful—it allows one to reconstruct all the k -qubit reduced density matrices of the state ρ and compute for example the expected value of any k -local Hamiltonian operator. So shadow tomography with the set $\mathcal{P}_k^{(n)}$ is particularly relevant for characterizing ground states of quantum spin systems with few-body interactions.

A different subset of Pauli operators describes few-body interactions between fermionic particles, such as the electronic structure of molecules. Just as before, there is a fermionic locality parameter k but it is fundamentally different from the one defined by Pauli weight. To describe it, one fixes any subset of $2n$ anticommuting n -qubit Pauli operators:

$$c_1, c_2, \dots, c_{2n} \in \mathcal{P}^{(n)} \quad \forall a, b : c_a c_b + c_b c_a = 2\delta_{ab} \mathbb{1}, \quad (3)$$

where δ_{ab} is 1 if $a = b$ and 0 otherwise. Note that since $c_a \in \mathcal{P}^{(n)}$ we also have $c_a^\dagger = c_a$ for each $1 \leq i \leq 2n$. These are known as the Majorana fermion operators associated with a fermionic system with n modes (a fermionic mode is a state that can either be occupied or unoccupied by a fermionic particle). Note that there is a freedom here—a particular choice of operators in [Eq. \(3\)](#) is a “fermion-to-qubit mapping” that describes how we associate the degrees of freedom of the n -mode fermionic system with those of the n -qubit Hilbert space. Such mappings have a long history and there are several choices that are used in practice to design algorithms for fermionic systems on a quantum computer (see, e.g.,

Refs [JW28, BK02, SRL12, JKMN20, DKBC21], and Section 4). However, our discussion and the results described below apply to any choice of fermion-to-qubit mapping.

With our Majoranas (Eq. (3)) in hand, let us now define the Majorana monomials as

$$\Gamma(x) = i^{|x| \cdot (|x|-1)/2} c_1^{x_1} c_2^{x_2} \dots c_{2n}^{x_{2n}} \quad \forall x \in \{0, 1\}^{2n}. \quad (4)$$

The overall phase factor $i^{|x| \cdot (|x|-1)/2}$ ensures that $\Gamma(x)$ is Hermitian for all $x \in \{0, 1\}^{2n}$. In fact, the 4^n operators

$$\{\Gamma(x) : x \in \{0, 1\}^{2n}\} \quad (5)$$

coincide with the 4^n n -qubit Pauli operators in $\mathcal{P}^{(n)}$, up to (efficiently computable) signs. Now let us define the k -body fermionic operators

$$\mathcal{F}_k^{(n)} = \{\Gamma(x) : |x| = 2k\}, \quad (\text{\textit{k}-body fermionic operators}) \quad (6)$$

which should be compared with Eq. (2). Note that $|\mathcal{F}_k^{(n)}| = \binom{2n}{2k}$ and $\log |\mathcal{F}_k^{(n)}| = O(k \log n)$. A system of n fermionic modes with k -particle interactions is described by a Hamiltonian operator that is a sum of terms from $\mathcal{F}_k^{(n)}$. Note that $\mathcal{F}_k^{(n)}$ consists of the Majorana monomials in Eq. (5) of degree $2k$; typically only these even-degree monomials are relevant to physics and chemistry due to conservation of fermionic parity.

The notion of k -locality for fermions is fundamentally more expressive than of k -locality for spin systems; the former subsumes the latter in the sense that compact “qubit-to-fermion mappings” exist which embed the k -local n -qubit Pauli operators within a subset of the k -body fermionic operators on $O(n)$ fermionic modes (see for example Ref. [BGKT19]), whereas fermion-to-qubit mappings necessarily represent the Majorana fermion operators Eq. (3) using Pauli operators of average weight at least $\Omega(\log(n))$ [JKMN20].

The expectation values of the fermionic operators $\mathcal{F}_k^{(n)}$ comprise the matrix elements of what physicists and chemists refer to as the k -body reduced density matrix, or k -RDM. One can efficiently compute most physically relevant local observables of a fermionic system (e.g., dipole moment, charge density, and importantly—energy) from just the 1- and 2-RDMs as a consequence of fermions being identical particles that interact pairwise. A large body of literature exists on methods for computing the fermionic 2-RDM matrix elements as a means of estimating the energy of chemical systems during the course of a quantum variational algorithm (see e.g., [VYI20, PZC23, BMBO20, ZRM21]).

A number of important methods for post-processing the output of quantum simulations actually *require* measuring k -RDMs. For example, subspace expansion techniques for approximating excited states from ground states via linear response typically require the 4-RDM [MKCdJ17, YHM⁺22]. Perturbation theory [GWH⁺16, SKGA17] and multi-reference configuration interaction methods [TRJ⁺20] for relaxing ground state calculations in small basis sets towards their continuum (large basis) limit often require the 4-RDM but converge even faster given access to higher order RDMs. Finally, there are popular impurity model schemes for extrapolating finite simulations of condensed phase fermionic systems towards their thermodynamic limits (e.g., density matrix embedding theory [WJHSC16]) and hybrid quantum-classical schemes for quantum Monte Carlo [HOR⁺22], that require the full 1-RDM.

1.2 Results

Single-copy measurements. It would be very practical if we could achieve triply efficient Pauli shadow tomography using only measurements on one copy of ρ at a time—and for k -local Paulis, it can be done. The shadow tomography task for $\mathcal{P}_k^{(n)}$ can be performed, using a time-efficient algorithm, using only *single-copy* measurements on $O(3^k(\log |\mathcal{P}_k^{(n)}|)/\epsilon^2) = O(3^k(k \log n)/\epsilon^2)$ copies of ρ [HKP20, CW20, BMBO20, JKMN20, EHF19].

One might hope that we can similarly achieve triply efficient Pauli shadow tomography for $\mathcal{F}_k^{(n)}$ and $\mathcal{P}^{(n)}$ as well. Unfortunately, this is not possible in either case, even if we only care about sample efficiency and allow unbounded computation time. It was shown in Ref [CCHL22] that sample-efficient shadow tomography with single-copy measurements is impossible for the set of all Paulis.

Theorem 2 ([CCHL22]). *There is no sample-efficient shadow tomography protocol with single-copy measurements for the set $S = \mathcal{P}^{(n)}$ of all Paulis. In particular, any protocol based on single-copy measurements must consume $\Omega(2^n/\epsilon^2)$ copies of ρ .*

For k -body fermionic operators, one can also establish a lower bound, see [Section 4.1](#).²

Theorem 3. *There is no sample-efficient single-copy shadow tomography protocol for the set $\mathcal{F}_k^{(n)}$ of k -body fermionic operators. In particular, for $k = O(1)$, any protocol based on single-copy measurements must consume $\Omega(n^k/\epsilon^2)$ copies of ρ .*

Several efficient shadow tomography algorithms based on single-copy measurements are known which achieve the $\Omega(n^k/\epsilon^2)$ lower bound, up to a log factor [BMBO20, JKMN20, WHLB23, HWM⁺22].

The lower bound in [Theorem 3](#) demonstrates a significant difference between learning local fermionic observables and local Pauli observables; for local Paulis there are sample-efficient single-copy protocols, whereas this is impossible for local fermionic observables.

In order to achieve triply efficient shadow tomography for $\mathcal{F}_k^{(n)}$ and $\mathcal{P}^{(n)}$, we will need to measure 2 or more copies of ρ at a time. Before jumping into 2-copy measurements, we review the 1-copy algorithm for k -local Pauli operators and offer a new interpretation in terms of fractional graph colorings that will be used in our algorithms.

The classical shadows single-copy protocol for k -local Paulis is very simple: it is based on measuring each qubit of ρ (in each copy of ρ) in a random single-qubit Pauli basis X, Y or Z uniformly at random [HKP20]. The postprocessing of the measurement data to compute expected values is equally simple.

The high-level format of this protocol is as follows: one selects a Clifford basis at random according to some probability distribution p , and then measures in that basis. The distribution has the property that each of the Paulis P in the set S of observables of interest (in the above, $S = \mathcal{P}_k^{(n)}$) has a high chance (at least 3^{-k}) of being diagonal in a basis sampled from p . Every time we pick a basis in which a Pauli is diagonal, we learn some information about its expected value and hence the sample complexity of the protocol is inversely related to the probability of being diagonal in a randomly sampled basis.

²Theorem 1 of [BMBO20] contains a lower bound for single-copy non-adaptive Clifford measurements; on the other hand, [Theorem 3](#) applies to arbitrary and even adaptive single-copy measurements.

We reinterpret this measurement strategy, and other protocols based on random single-copy Clifford measurements, as arising from fractional colorings of the commutation graph of the observables S , defined as follows:

Definition 4. *The commutation graph $G(S)$ of a set $S \subseteq \mathcal{P}^{(n)}$ of Pauli operators is the graph with vertex set S and an edge between every pair of anticommuting operators.*

An independent set in $G(S)$ corresponds to a set of commuting observables that can be measured simultaneously via a Clifford measurement. Similarly, a coloring of this graph with χ colors describes a learning strategy with deterministic single-copy Clifford measurements, based on measuring χ disjoint sets of commuting Pauli observables. Such deterministic graph coloring strategies for learning Pauli observables have been explored previously, see for example Ref. [JGM19, VYI20]. But the protocol for local Pauli observables described above is not based on a coloring of the commutation graph $G(S)$: the measurement bases are associated with *overlapping* sets of commuting Pauli observables, and are chosen randomly rather than deterministically. As we will see, a probabilistic Clifford measurement strategy can be viewed as defining a *fractional coloring* of $G(S)$. A fractional coloring is a well-studied relaxation of the notion of graph coloring, see Section 2.2 for details. We show that the sample complexity of single-copy learning with Clifford measurements is upper bounded by the fractional chromatic number of $G(S)$, which is the size of the smallest fractional coloring of $G(S)$. In Section 2.2, we prove the following theorem.

Theorem 5. *Let $S \subseteq \mathcal{P}^{(n)}$. Suppose the commutation graph $G(S)$ admits a fractional coloring of size χ that can be sampled by a classical randomized algorithm with runtime T . Then there is an algorithm using only single-copy Clifford measurements of ρ which can estimate $\text{Tr}(P\rho)$ within error ϵ for all $P \in S$ with high probability using*

$$O(\chi(\log |S|)/\epsilon^2) \tag{7}$$

copies of ρ . The runtime of the algorithm is $O((T + n^3) \cdot \chi(\log |S|)/\epsilon^2)$.

The single-copy measurement strategies based on fractional colorings from Theorem 5 have the special feature that they only use Clifford measurements, which have efficient classical descriptions. Such protocols learn a *compressed classical representation* of ρ that can be used to compute Pauli observables, see Section 2.2 for details.

The power of two copies. In light of Theorems 2 and 3 we see that there exist sets of Pauli observables for which sample-efficient shadow tomography cannot be achieved with one-copy measurements. Are two-copy measurements enough?

Our starting point here is a new algorithm that shows that sample-efficient shadow tomography is indeed possible in the general case with two-copy measurements. The protocol has three steps.

The first step is the learning magnitudes subroutine from Ref. [HKP21], which we now review. In this step we perform Bell sampling to measure copies of $\rho \otimes \rho$ in the Clifford basis that diagonalizes the commuting Pauli observables $P \otimes P$ for all $P \in \mathcal{P}^{(n)}$. Since these operators commute, we can learn the all observables $\text{Tr}((P \otimes P)(\rho \otimes \rho)) = \text{Tr}(\rho P)^2$ for

$P \in S$ to error δ using $O((\log |S|)/\delta^2)$ measurements. By choosing $\delta = \Theta(\epsilon^2)$, we see that $O((\log |S|)/\epsilon^4)$ two-copy measurements suffices to compute estimates $\{u_P\}_{P \in S}$ such that

$$|u_P - |\text{Tr}(\rho P)|| \leq \epsilon/4 \quad \text{for all } P \in S, \quad (8)$$

with high probability. After we have learned the magnitudes in this way, we may find that some of our estimates are negligible; if our estimate u_P from the first stage is less than $3\epsilon/4$ then 0 is an ϵ -approximation to the expected value $\text{Tr}(\rho P)$ and we can forget about this Pauli P going forward. So in the second stage of the algorithm we are only concerned with observables in the set

$$S_\epsilon = \{P \in S : |u_P| \geq 3\epsilon/4\}. \quad (9)$$

This set S_ϵ is a random variable determined by the output of Bell sampling, but the condition in Eq. (8) implies that with high probability we have

$$|\text{Tr}(\rho P)| \geq \epsilon/2 \quad \text{for all } P \in S_\epsilon. \quad (10)$$

To complete the learning task it suffices to then compute the sign of $\text{Tr}(\rho P)$ for all Paulis $P \in S_\epsilon$. To this end, in the second step of the protocol we compute a classical description of a *mimicking state* σ such that

$$|\text{Tr}(\sigma P)| \geq \epsilon/4 \quad \text{for all } P \in S_\epsilon. \quad (11)$$

A key observation is that, if Eq. (10) holds (which occurs with high probability), then (A) there always exists at least one mimicking state σ (for instance $\sigma = \rho$ is a valid mimicking state), and (B) given the set S_ϵ , a mimicking state can be found without using any additional copies of ρ , by brute-force search. So this second step of the protocol, while computationally inefficient, can be performed without using any additional copies of ρ . In the final, third step of the protocol, we now perform Bell sampling on the tensor product $\rho \otimes \sigma$. This will require us to repeatedly prepare the mimicking state σ on our quantum computer.

The resulting samples are used to estimate mean values

$$\text{Tr}(P \otimes P(\rho \otimes \sigma)) = \text{Tr}(P\rho)\text{Tr}(P\sigma) \quad P \in S_\epsilon. \quad (12)$$

Since σ is a mimicking state and satisfies Eq. (11), each of these mean values has magnitude at least $\epsilon^2/8$ and we can compute all of them up to $\epsilon^2/16$ error using only $O((\log |S|)/\epsilon^4)$ Bell samples. Since σ is known to us, we can compute $\text{Tr}(\sigma P)$ exactly (including the sign), and using this knowledge and our estimates of the mean values in Eq. (12) we can infer the sign of each mean value $\text{Tr}(\rho P)$ for $P \in S_\epsilon$.

In summary, we have described a sample-efficient shadow tomography protocol for any subset of n -qubit Pauli observables, that only uses two-copy Clifford measurements.

Theorem 6. *There exists a shadow tomography protocol for any subset $S \subseteq \mathcal{P}^{(n)}$ of Pauli observables, that uses only $O((\log |S|)/\epsilon^4)$ two-copy Clifford measurements.*

Note that this matches the sample complexity of the protocol from Ref. [HKP21], while only using two-copy measurements. Although the above protocol is sample efficient, it is extremely inefficient in terms of runtime as it seems to require a brute-force search over the set of all n -qubit quantum states. Here we show that the outrageous runtime can be tamed – the protocol can be modified so that its runtime is upper bounded as $\text{poly}(2^n)$. In this way we obtain a sample- and time-efficient two-copy protocol for the set of all Paulis.

Theorem 7. *There exists a triply efficient shadow tomography protocol for the set $S = \mathcal{P}^{(n)}$ of all n -qubit Paulis that uses only two-copy Clifford measurements. In particular, it has sample complexity $O(n \log(n/\epsilon)/\epsilon^4)$ and time complexity $\text{poly}(2^n, 1/\epsilon)$.*

In Section 3.1 we complete the proof of Theorem 7 by showing that a suitable mimicking state σ satisfying Eq. (11) can be computed using $O(\text{poly}(2^n))$ time and $O(n \log(n/\epsilon)/\epsilon^4)$ additional single-copy measurements of ρ . The algorithm uses the matrix multiplicative weights technique [AK07, ACH⁺18]. Once a classical description of a mimicking state has been computed, we can create the corresponding quantum state using $O(4^n)$ gates [SBM06].

Our idea to use a mimicking state computed via the matrix multiplicative weights algorithm has already found an application: Ref. [KWM24] applies this to perform shadow tomography on bosonic displacement operators.

Improved algorithms using graph theory. The remainder of our results employ a general framework for two-copy shadow tomography which is based on fractional graph coloring. Our framework can be viewed as an extension of a heuristic learning algorithm proposed in Appendix E.2.d of Ref. [HKP21]; in this work we use it to obtain algorithms with rigorous performance guarantees.

We have already seen that single-copy tomography for any set of observables S reduces to fractional graph coloring for the commutation graph $G(S)$. Likewise, via Bell sampling, two-copy tomography reduces to fractional graph coloring for the commutation graph $G(S_\epsilon)$. That is, we propose to use the single-copy algorithm to learn all observables in S_ϵ , once we have already determined S_ϵ using an initial stage of Bell sampling. But how do the two-copy measurements help us?

A key insight is that the Paulis in S_ϵ cannot be very anticommuting. Intuition from the Heisenberg uncertainty principle tells us that anticommuting (traceless) observables cannot simultaneously be large on a quantum state, since the quantum state cannot simultaneously be an eigenvector of anticommuting observables. But the high probability event in Eq. (10) implies that the Paulis in S_ϵ are simultaneously large on the state ρ , and thus cannot anticommute with each other too often. This can be formalized as an upper bound on the clique number of their commutation graph. In Section 2.3 we show the following:

Lemma 8. *The largest clique in the commutation graph $G(S_\epsilon)$ has size at most $4/\epsilon^2$ with high probability.*

Going forward, our aim is to exploit this upper bound on the clique number to compute good (fractional) colorings.

Unfortunately, it is well known that the chromatic (or fractional chromatic) number is not upper bounded by any function of the clique number in general³. However, such upper bounds can be established for certain families of graphs, a research direction pioneered by Gyárfás [Gyá87]. A family of graphs for which this is possible is called *chi-bounded* and the upper bound on chromatic number is said to be expressed in terms of a chi-binding function (see Refs. [SR19, SS20] for recent surveys). Our shadow tomography learning task for a set

³For example, there exists a family of triangle-free graphs with chromatic number $\Omega(\sqrt{m/\log m})$ where m is the number of vertices [Kim95]

of observables S thus reduces to establishing a suitable chi-binding function for the family of induced subgraphs of the commutation graph $G(S)$, see [Section 2.3](#) for details.

We show that the family of induced subgraphs of the commutation graph of k -body fermionic observables admits a polynomial chi-binding function (that does not depend on n).

Lemma 9. *Let $k \geq 1$, and let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_k^{(n)})$ of k -body fermionic observables, and let ω be the size of the largest clique in G' . Then the fractional chromatic number of G' satisfies*

$$\chi_f(G') \leq p_k(\omega). \quad (13)$$

where p_k is a polynomial. Moreover, for any $k = O(1)$ we can sample from a fractional coloring of G' with size $p_k(\omega)$ using a classical algorithm with runtime $\text{poly}(n)$. The polynomials for $k = 1, 2$ are $p_1(\omega) = \omega + 1$ and $p_2(\omega) = O(\omega^8)$.

The proof of [Lemma 9](#) is provided in [Section 4.3](#). As discussed above, the commutation graph $G(S_\epsilon)$ is an induced subgraph of $G(S)$ with clique number at most $O(1/\epsilon^2)$. For k -body fermionic observables $S = \mathcal{F}_k^{(n)}$, [Lemma 9](#) tells us there is an efficiently computable fractional coloring of $G(S_\epsilon)$ with at most $\text{poly}(1/\epsilon^2)$ colors. We can then use the single-copy learning protocol from [Theorem 5](#) to learn all observables in S_ϵ . This reduction, which describes how to convert [Lemma 9](#) into a two-copy learning protocol, is formalized in [Lemma 20](#). Putting it all together gives the following theorem.

Theorem 10. *Let $k = O(1)$. There exists a triply efficient shadow tomography protocol for the set $S = \mathcal{F}_k^{(n)}$ of k -body fermionic observables that uses only two-copy Clifford measurements.*

This triply efficient protocol has sample complexity

$$O\left(\frac{\log |\mathcal{F}_k^{(n)}|}{\epsilon^4} + \frac{p_k(4/\epsilon^2) \log |\mathcal{F}_k^{(n)}|}{\epsilon^2}\right) = O((k \log n) p_k(4/\epsilon^2)/\epsilon^2), \quad (14)$$

where p_k is the polynomial from [Lemma 9](#) that depends on the locality k , and we also used the fact that $p_k(\omega) = \Omega(\omega)$ for all $k \geq 1$. For each $k \geq 1$ we obtain an exponential improvement over single-copy learning protocols in terms of the sample complexity as a function of system size n . For $k = 1$ our learning algorithm has sample complexity $O((\log n)/\epsilon^4)$, and the measurements and postprocessing are simple to implement. We anticipate that this learning algorithm could find applications in quantum simulations of chemistry and fermionic physics. With our current analysis, the degree of the polynomial p_k increases very rapidly as a function of k rendering the scheme less practical for $k \geq 2$. We hope this could be improved in future work. An upper bound on the ϵ -dependence of the sample complexity is $\sim \epsilon^{-O((2k)^{k+1})}$. For $k = 2, 3$ the sample complexity is $\sim \epsilon^{-18}$ and $\sim \epsilon^{-110}$ respectively.

It is natural to ask how far we can push this two-copy framework based on Bell sampling and fractional coloring. Below, we show that it provides a nontrivial shadow tomography protocol for *any* subset of Pauli observables $S \subseteq \mathcal{P}^{(n)}$. This gives hope that our framework could lead to triply efficient shadow tomography in the general case.

We shall exploit the fact that the longest induced path in the commutation graph $G(\mathcal{P}^{(n)})$ contains at most $2n + 1$ vertices, see [Section 3.2](#) for details. The following upper bound on chromatic number then follows from a seminal result in chi-boundedness due to Gyárfás [[Gyá87](#)]; see [Section 3.2](#).

Lemma 11. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{P}^{(n)})$, and let ω be the size of the largest clique in G' . The chromatic number of G' is upper bounded as*

$$\chi(G') \leq (2n + 1)^{\omega-1}. \quad (15)$$

Moreover, a coloring with this many colors can be computed by a classical algorithm with runtime $\text{poly}(|G'|, n^\omega)$.

To get a shadow tomography algorithm for any set of Pauli observables $S \subseteq P^{(n)}$, we follow the strategy outlined above and formalized in [Lemma 20](#). That is, we apply [Lemma 11](#) to the subgraph $G' = G(S_\epsilon)$ induced by the set S_ϵ computed via Bell sampling. From [Lemma 8](#) we have that with high probability the largest clique in G' has size $\omega = O(1/\epsilon^2)$. So we get an coloring of $G(S_\epsilon)$ with at most $n^{O(1/\epsilon^2)}$ colors, that can be computed with runtime $\text{poly}(|S|, n^{1/\epsilon^2})$. When $\epsilon = \Omega(1)$ is a small constant, this protocol is time-efficient, has sample complexity $\text{poly}(n)$, and only uses two-copy measurements, for any subset of Pauli observables S . This gives a sample-efficient protocol only when $|S|$ is exponentially large as a function of n . At a technical level this is a consequence of the factor of n appearing in [Eq. \(15\)](#), and we do not know if this can be avoided.

This leaves open the question of triply efficient shadow tomography for arbitrary subsets of Pauli observables. However, we will see that it provides insight into a related question concerning compressed classical representations of quantum states.

Rapid-retrieval Pauli compression. Can we compress an n -qubit quantum state into a small amount of classical information, so that the compressed classical description is sufficient to extract the expectation values of any bounded observable to within a small constant error? This question has been studied using tools from communication complexity and it is known that an exponential classical description size is necessary if one wishes to recover bounded observables in the general case; a representation size $\tilde{\Theta}(\sqrt{2^n})$ is necessary and sufficient for n -qubit pure states [[Raz99](#), [GKK⁺07](#), [GS19](#)].

On the other hand, if we restrict our attention to the set of n -qubit Pauli observables (or other sets of observables with only singly exponential size), a classical description of size $\text{poly}(n)$ exists and can be computed using the matrix multiplicative weights algorithm [[Aar04](#), [ACH⁺18](#)]. However, a significant drawback of known methods for this task is that they require exponential classical runtime to extract the expected value of a given Pauli observable from the compressed classical representation.

A consequence of [Lemma 11](#) is that this exponential cost can be avoided, at least for any small constant precision $\epsilon = \Omega(1)$. That is, one can compress an n -qubit state ρ into $\text{poly}(n)$ classical bits. Given this classical data and an n -qubit Pauli P , there is an *efficient* classical algorithm to estimate $\text{Tr}(\rho P)$ to within ϵ -error. Moreover, such a representation can be learned from $\text{poly}(n)$ samples of ρ .

Corollary 12 (Rapid-retrieval Pauli compression). *Let ρ be an n -qubit quantum state. Let $\epsilon \in (0, 1)$ be a constant independent of n . Using two-copy Clifford measurements on $\text{poly}(n)$ copies of ρ , along with $2^{O(n)}$ runtime, we can (with high probability) learn a compressed classical representation of ρ , call it $D(\rho, \epsilon)$, that consists of $\text{poly}(n)$ bits. An ϵ -approximation to the expected value $\text{Tr}(\rho P)$ of any Pauli observable $P \in \mathcal{P}^{(n)}$ can be extracted from $D(\rho, \epsilon)$ using a classical algorithm with $\text{poly}(n)$ runtime.*

The classical description $D(\rho, \epsilon)$ consists of a list of all the Clifford measurement bases and measurement outcomes used in the two-copy learning algorithm discussed above, see [Section 2](#) for details. In particular, [Corollary 12](#) is obtained by combining [Lemma 11](#) and [Lemma 21](#).

1.3 Discussion and open questions

In this paper we have provided the first triply efficient shadow tomography protocols for the set of k -body fermionic observables and the set of all Pauli operators. We have also provided a route to strengthening and generalizing our results via a connection between two-copy tomography and graph theory techniques related to chi-boundedness. In [Table 2](#) we provide a comparison of our results with other known protocols for shadow tomography.

There are many questions left open by our work. Is it possible to improve the upper bounds from [Eq. \(13\)](#) and [Eq. \(15\)](#)—e.g., can we establish better chi-binding functions for the (families of) commutation graphs of interest? Is rapid-retrieval compression possible for smaller error parameters, e.g., $\epsilon = 1/\text{poly}(n)$? Can we devise triply efficient learning algorithms for any subset of Pauli observables?

One route towards resolving these questions would be via improved algorithms for coloring the commutation graph $G(S_\epsilon)$. The following conjecture asserts an efficient fractional coloring of the commutation graph of any subset of Paulis that has simultaneously large expected values in a quantum state.

Conjecture 13. *Let ρ be an n -qubit state, $\delta \in (0, 1)$, and let $B \subseteq \mathcal{P}^{(n)}$ be the set of all Paulis P such that $|\text{Tr}(\rho P)| \geq \delta$. There is a fractional coloring of the commutation graph $G(B)$ of size $O(1/\delta^2)$.*

If this conjecture holds, and in addition the fractional coloring is suitably efficient,⁴ then we would obtain a triply efficient Pauli shadow tomography algorithm for *any* subset S of Pauli observables. Moreover, the learning algorithm would also output a rapid-retrieval Pauli compression of all observables in S , of size $O(n^2(\log |S|)/\epsilon^4)$, see [Section 2.3](#).

To address [Conjecture 13](#), it is natural to ask if [Lemma 8](#) can be strengthened by showing that $O(1/\epsilon^2)$ is in fact an upper bound on the *fractional clique number*—a well-known linear programming relaxation of the clique number [SU11]. The existence of a suitable fractional coloring stated in [Conjecture 13](#) would then follow from linear programming duality, which asserts that the fractional clique number of any graph equals its fractional chromatic number.

The remainder of the paper is organized as follows. In [Section 2](#) we describe the commutation graph and its properties, as well as our frameworks for one- and two-copy Clifford learning. In [Section 3](#) we describe techniques that go into our algorithms for learning all Paulis: in [Section 3.1](#) we describe the algorithm to compute a mimicking state, which completes the proof of [Theorem 7](#), and in [Section 3.2](#) we give the proof of [Lemma 11](#). Finally in [Section 4](#) we establish our results concerning shadow tomography for k -body fermionic operators. In [Section 4.1](#) we establish the lower bound [Theorem 3](#) for single-copy learning. In [Section 4.2](#) we describe the simple triply efficient learning algorithm for the special case $k = 1$, and then in [Section 4.3](#) we consider the case $k \geq 2$ and prove [Lemma 9](#).

⁴In particular, we require that a fractional coloring of size $O(1/\delta^2)$ for any subset $R \subseteq B$ can be sampled in time $\text{poly}(|R|, n, 1/\delta)$

Observables	Copies per measurement	Algorithm or lower bound	Sample complexity	Time efficient?
k -local Pauli operators	1	Naive	$O(3^k \binom{n}{k} k(\log n)/\epsilon^2)$	✓
		Various methods [CW20, EHF19, BMBO20] [JKMN20, HKP20]	$O(3^k k(\log n)/\epsilon^2)$	✓
		Lower bound (Theorem 26) for $k \leq \log_3(2n+1)$.	$\Omega(3^k/\epsilon^2)$	—
	2	Theorem 6	$O(k(\log n)/\epsilon^4)$	✗
	unrestricted	Bell sampling and gentle measurements [HKP21]	$O(k(\log n)/\epsilon^4)$	✓
k -body fermionic operators	1	Naive	$O(\binom{2n}{2k} k(\log n)/\epsilon^2)$	✓
		Various methods [BMBO20, JKMN20, WHLB23]	$O_k(n^k(\log n)/\epsilon^2)^*$	✓
		Lower bound (Theorem 27)	$\Omega_k(n^k/\epsilon^2)^*$	—
	2	Theorem 6	$O(k(\log n)/\epsilon^4)$	✗
		Theorem 10 for $k=1$	$O((\log n)/\epsilon^4)$	✓
		Theorem 10	$\log n \cdot \text{poly}_k(1/\epsilon)^\dagger$	✓
	unrestricted	Bell sampling and gentle measurements [HKP21]	$O(k(\log n)/\epsilon^4)$	✓
All Pauli operators	1	Naive	$O(4^n n/\epsilon^2)$	✓
		Random Clifford [‡]	$O(2^n n/\epsilon^2)$	✓
		Lower bound [CCHL22]	$\Omega(2^n/\epsilon^2)$	—
	2	Theorem 7	$O(n \log(n/\epsilon)/\epsilon^4)$	✓
		Corollary 12	$n^{O(1/\epsilon^2)}$	✓
	unrestricted	Bell sampling and gentle measurements [HKP21]	$O(n/\epsilon^4)$	✓

Table 2: A summary of Pauli shadow tomography algorithms for the sets of observables we study. Cells in blue are new results in this paper.

*The constant depends on k .

[†]The degree of the polynomial depends on k .

[‡]Based on single-copy measurements in uniformly random Clifford bases.

Independent work. While writing this paper, we learned of related, independent work by Chen, Gong, and Ye [CGY24] that studies the sample complexity of Pauli shadow tomography with limited quantum memory. We have coordinated our arXiv submissions so that both papers appear on the same day.

Acknowledgments

We thank Bill Huggins for valuable comments on this manuscript. DG thanks Jim Geelen for a helpful discussion about chi-boundedness, and Sophie Spirkl for explaining how the result of Gyárfás [Gyá87] can be turned into an algorithm. RB thanks Nicholas Rubin and Joonho Lee for discussions about using fermionic RDMs in embedding and perturbation theory contexts. DG is a CIFAR fellow in the quantum information science program.

2 Commutation, learning, and coloring

In this section we describe properties of the commutation graph $G(S)$ of a set of Pauli observables $S \subseteq \mathcal{P}^{(n)}$, and the connection between these properties and shadow tomography algorithms. In Section 2.1 we describe the tension between learning and anticommutativity of a set of observables. This is quantified by a *commutation index* which was shown in Ref. [CCHL22] to lower bound the sample complexity of single-copy learning. Then in Sections 2.2 and 2.3 we describe our frameworks for one- and two-copy learning with probabilistic Clifford measurement strategies.

2.1 Commutation index

The following result represents a kind of uncertainty principle, generalizing the familiar Bloch sphere constraint $\langle X \rangle^2 + \langle Y \rangle^2 + \langle Z \rangle^2 \leq 1$.

Lemma 14. *If Pauli operators $P_1, \dots, P_m$ pairwise anticommute, then for any state ρ*

$$\sum_j \text{Tr}(P_j \rho)^2 \leq 1. \quad (16)$$

Versions of Lemma 14 appear in various papers, for example [AEHK16, Theorem 1]. Since the proof is simple, we reproduce it here.

Proof. Given ρ , let $a_j = \text{Tr}(P_j \rho)$. We aim to show $\sum_j a_j^2 \leq 1$. Consider the observable

$$Q = \sum_j a_j P_j. \quad (17)$$

We will use the inequality $\text{Tr}(Q^2 \rho) - \text{Tr}(Q \rho)^2 = \text{Var}_\rho(Q) \geq 0$. Formally, this holds since ρ is positive semi-definite so $\text{Tr}(O \rho) \geq 0$ for any positive semi-definite operator O and

$$\text{Tr}\left((Q - \text{Tr}(Q \rho) \mathbb{1})^2 \rho\right) \geq 0 \implies \text{Tr}(Q \rho)^2 \leq \text{Tr}(Q^2 \rho). \quad (18)$$

Due to anticommutativity of $P_1, \dots, P_m$, we have

$$Q^2 = \sum_j a_j^2 P_j^2 + \sum_{j \neq l} a_j a_l P_j P_l = \sum_j a_j^2 \cdot \mathbb{1} + \frac{1}{2} \sum_{j \neq l} a_j a_l \{P_j, P_l\} = \sum_j a_j^2 \cdot \mathbb{1}. \quad (19)$$

Note $P_j^2 = \mathbb{1}$ since they are Hermitian unitaries. Thus $\text{Tr}(Q^2 \rho) = \sum_j a_j^2$. On the other hand $\text{Tr}(Q \rho) = \sum_j a_j^2$, thus

$$\text{Tr}(Q \rho)^2 \leq \text{Tr}(Q^2 \rho) \implies \left(\sum_j a_j^2 \right)^2 \leq \sum_j a_j^2 \implies \sum_j a_j^2 \leq 1. \quad \square \quad (20)$$

In [Lemma 14](#), we saw that pairwise anticommuting Pauli operators cannot all have large expected values in a quantum state, as the sum of their squared expected values cannot exceed 1. More generally, the average of squares of expected values of Paulis in a set S is defined to be its *commutation index*.

Definition 15. For a set S of Pauli operators, define their commutation index by

$$\Delta(S) = \frac{1}{|S|} \max_{\rho} \sum_{P \in S} \text{Tr}(P \rho)^2. \quad (21)$$

For example, [Lemma 14](#) shows that if all the Paulis in S anticommute, then $\Delta(S) \leq 1/|S|$.

Ref. [\[CCHL22\]](#) shows that the inverse of the commutation index is a lower bound on the sample complexity of single-copy shadow tomography for S . It can be interpreted as describing a tension between anticommutation and learnability. (Maximization over states as stated in [\[CCHL22, Equation 79\]](#) can be replaced by maximization over density matrices via convexity.)

Theorem 16 (Theorem 5.5, [\[CCHL22\]](#)). *Shadow tomography to precision ϵ for a set S of Pauli observables with single-copy measurements requires at least*

$$\Omega\left(\frac{1}{\epsilon^2 \Delta(S)}\right) \quad (22)$$

*copies of ρ . This holds even for adaptive measurement strategies.*⁵

The commutation index $\Delta(S)$ can be upper bounded in terms of the *Lovasz ϑ -function* [\[dGHG23, XSW23, HO22\]](#).

Definition 17. Let G be a graph on m vertices. The Lovasz ϑ -function $\vartheta(G)$ is defined by the following semidefinite program of dimension m :

$$\max \{ \text{Tr}(JX), X \in \mathbb{R}^{m \times m} \text{ s.t. } X \succeq 0, \text{Tr } X = 1, X_{jl} = 0 \ \forall (j, l) \in E \}, \quad (23)$$

where E denotes the edge set of the graph G and J the all-ones matrix. It has dual

$$\min \{ \lambda \in \mathbb{R} \text{ s.t. } \exists A \in \mathbb{R}^{m \times m}, A_{jj} = 1 \ \forall j, A_{jl} = 0 \ \forall (j, l) \notin E, \lambda A \succeq J \}. \quad (24)$$

The following result is a generalization of [Lemma 14](#), with proof in [Appendix A](#).

Lemma 18. ([\[dGHG23, XSW23, HO22\]](#)) *Let S be a set of Pauli operators. Then*

$$\Delta(S) \leq \frac{\vartheta(G(S))}{|S|}. \quad (25)$$

⁵The lower bound holds even if the shadow tomography scheme is only able to output the absolute values of the expectation values to precision ϵ .

2.2 Fractional coloring and single-copy Clifford learning

Here we describe the connection between shadow tomography algorithms that learn Pauli observables $S \subseteq \mathcal{P}^{(n)}$ using probabilistic Clifford measurements, and fractional colorings of the commutation graph $G(S)$. A fractional coloring is a relaxation of the usual notion of graph coloring [SU11].

Definition 19. *Let $G = (V, E)$ be a graph. A fractional coloring of G of size χ is a probability distribution q over independent sets $I \subseteq V$ with the property that*

$$\forall v \in V : \Pr_{I \sim q}(v \in I) \geq 1/\chi. \quad (26)$$

The fractional chromatic number $\chi_f(G)$ of G is the size of the smallest fractional coloring of G .

Note that the size χ of a fractional coloring need not be an integer. Also note that a (standard, non-fractional) coloring of G with χ colors can be regarded as a fractional coloring of size χ , corresponding to a uniform distribution over color classes. So the fractional chromatic number of a graph is upper bounded by its chromatic number.

[Theorem 5](#), restated below, asserts that if we have a fractional coloring of the commutation graph $G(S)$ of small size, then we can learn the expectation values of all observables in S with few single-copy Clifford measurements. In particular, the sample complexity of the algorithm scales linearly with the size of the fractional coloring.

In the following, samples from the fractional coloring are represented as binary vectors of length $|S|$ whose support is an independent set in $G(S)$. In this setting, the runtime to produce a single sample from a fractional coloring always satisfies $T \geq |S|$.

Theorem 5. *Let $S \subseteq \mathcal{P}^{(n)}$. Suppose the commutation graph $G(S)$ admits a fractional coloring of size χ that can be sampled by a classical randomized algorithm with runtime T . Then there is an algorithm using only single-copy Clifford measurements of ρ which can estimate $\text{Tr}(P\rho)$ within error ϵ for all $P \in S$ with high probability using*

$$O(\chi(\log |S|)/\epsilon^2) \quad (7)$$

copies of ρ . The runtime of the algorithm is $O((T + n^3) \cdot \chi(\log |S|)/\epsilon^2)$.

Proof. An independent set I in the commutation graph $G(S)$ consists of a set of mutually commuting Pauli operators, which can be simultaneously measured by applying a Clifford circuit and measuring in the computational basis. Such a Clifford circuit can be computed using a classical algorithm with $O(n^3)$ runtime, via standard techniques in the stabilizer formalism [AG04].

If we draw an independent set I from a fractional coloring q with size χ and then measure ρ in the corresponding Clifford basis, the result gives us a measurement of Pauli $P \in S$ whenever $P \in I$. Note that it is also possible that we get more useful measurements than this—the Clifford unitary may diagonalize some Paulis P that are not in I .

Suppose we repeat this process independently N times, sampling independent sets $I_1, I_2, \dots, I_N$ and measuring N independent identical copies of ρ in the corresponding Clifford

bases $C_1, C_2, \dots, C_N$. For each $P \in S$, let $x_P^j \in \{-1, 0, 1\}$ be the random variable that is equal to the measured outcome of P if it is diagonalized by C_j , and zero otherwise.

We have

$$\Pr[x_P^j \in \{-1, 1\}] \geq \Pr[P \in I_j] \geq \frac{1}{\chi} \quad j \in [N] \quad P \in S, \quad (27)$$

where we used the fact that q is a fractional coloring of size χ .

For each Pauli $P \in S$, let

$$R_P = \{j : x_P^j \in \{-1, 1\}\} \quad \text{and} \quad N_P = |R_P|. \quad (28)$$

Consider the sample mean

$$\tilde{P} = \frac{1}{N_P} \sum_{j \in R_P} x_P^j. \quad (29)$$

Conditioned on a fixed value $N_P \geq 1$, this sample mean $\tilde{P}$ is an average of N_P independent ± 1 -valued random variables. It satisfies

$$\mathbb{E}(\tilde{P}) = \text{Tr}(\rho P) \quad \text{and} \quad \text{Var}(\tilde{P}) \leq \frac{1}{N_P}. \quad (30)$$

By Chebyshev's inequality we have

$$\Pr \left[|\tilde{P} - \text{Tr}(\rho P)| \geq \epsilon \mid N_P \geq 100/\epsilon^2 \right] \leq 0.01. \quad (31)$$

From Eq. (27) we see that by taking $N = O(\chi/\epsilon^2)$ we can ensure that, for a given Pauli P , we have $N_P \geq 100/\epsilon^2$ with probability at least 0.99 (say). Therefore,

$$\Pr \left[|\tilde{P} - \text{Tr}(\rho P)| \leq \epsilon \right] \geq \Pr \left[N_P \geq 100/\epsilon^2 \right] \cdot \Pr \left[|\tilde{P} - \text{Tr}(\rho P)| \leq \epsilon \mid N_P \geq 100/\epsilon^2 \right] \quad (32)$$

$$\geq 0.99^2 \quad (33)$$

for each Pauli $P \in S$.

Now repeat the above process L times, generating sample means $\tilde{P}_1, \dots, \tilde{P}_L$ for each $P \in S$, and consider the median-of-means estimator

$$\lambda_P = \text{median}(\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_L). \quad (34)$$

By choosing $L = O(\log |S|)$ we can ensure that, for each $P \in S$ we have $|\lambda_P - \text{Tr}(\rho P)| \leq \epsilon$ with probability at least $0.01/|S|$. By a union bound we get all the expected values in S to within ϵ with probability at least 0.99. The total number of samples of ρ and the total number of samples from the fractional coloring q used in the algorithm, are both at most $LN = O(\chi(\log |S|)/\epsilon^2)$.

Now consider the runtime of the protocol. The independent sets I in the fractional coloring are specified explicitly as subsets of S , so the median-of-means estimator λ_P can be computed for all $P \in S$ with a runtime $O(LN|S|)$ once we have already obtained all the measurement data $\{x_P^j\}$. The total runtime is therefore upper bounded as $O((T + n^3 + |S|)LN)$, where the first term is the cost of sampling the fractional colorings, the second term is the cost of computing a Clifford circuit for each sample (and applying this circuit to measure the state), and the third term is the cost of postprocessing. Since $T \geq |S|$ the runtime simplifies to $O((T + n^3)LN)$. $\square$

Finally, let us show that single-copy measurement strategies from [Theorem 5](#) learn a *compressed classical representation* of ρ that encodes the expected values of all Pauli observables from S (to within error ϵ). Indeed, each Clifford measurement basis has an efficient classical description consisting of $O(n^2)$ bits. We can imagine a version of the learning algorithm described in [Theorem 5](#) where, after the measurements are performed using [Eq. \(7\)](#) copies of ρ , the resulting measurement outcomes and measurement bases are packaged up into a classical description $D(\rho, S, \epsilon)$ of size

$$O(n^2 \chi(\log |S|)/\epsilon^2). \quad (35)$$

Here there is a factor of n^2 for each measurement basis (each measurement outcome only requires n bits to describe and so describing the outcomes requires asymptotically fewer bits than describing the bases). The efficient protocol for extracting expected values $\text{Tr}(\rho P)$ with $P \in S$ (up to ϵ error) can be performed using only the compressed classical description $D(\rho, S, \epsilon)$.

2.3 Cliques and two-copy Clifford learning

As discussed in the Introduction, we propose a framework for two-copy learning that uses an initial stage of Bell sampling (as in Ref. [\[HKP21\]](#)) to determine a set $S_\epsilon \subseteq S$ that with high probability satisfies [Eq. \(10\)](#), which we restate:

$$|\text{Tr}(\rho P)| \geq \epsilon/2 \quad \text{for all } P \in S_\epsilon. \quad (36)$$

This step uses $O((\log |S|)/\epsilon^4)$ Clifford measurements on copies of $\rho \otimes \rho$. Then we aim to learn all observables in S_ϵ using the fractional coloring approach described in the previous section. In particular, we aim to find an (efficiently sampleable) fractional coloring of the commutation graph $G(S_\epsilon)$.

Lemma 20 (Template for two-copy Clifford shadow tomography). *Suppose that $G(S_\epsilon)$ admits a fractional coloring of size χ that can be sampled by a randomized algorithm with runtime T . Then there is an algorithm which performs shadow tomography for S using*

$$O((\log |S|)/\epsilon^4 + \chi(\log |S|)/\epsilon^2) \quad (37)$$

two-copy Clifford measurements and runtime $O(|S|(\log |S|)/\epsilon^4 + (T + n^3)\chi(\log |S|)/\epsilon^2)$.

Proof. The first step uses $O((\log |S|)/\epsilon^4)$ two-copy Bell measurements and classical runtime $O(|S|(\log |S|)/\epsilon^4)$ to compute the set S_ϵ . We output zero as our estimate for the expected value of any Pauli in $S \setminus S_\epsilon$. Then we use the single-copy learning protocol from [Theorem 5](#) to compute estimates of $\text{Tr}(\rho P)$ for all $P \in S_\epsilon$. This second step uses runtime $O((T + n^3)\chi(\log |S|)/\epsilon^2)$. $\square$

The protocol from [Lemma 20](#) is based on Clifford measurements, which have an efficient classical description. Because of this, and the fact that the classical postprocessing is simple, we obtain the following classical compressed representation of ρ .

Lemma 21 (Rapid-retrieval compression). *The shadow tomography protocol described in Lemma 20 learns a compressed classical representation of ρ consisting of B bits, where*

$$B = O(n(\log |S|)/\epsilon^4 + n^2\chi(\log |S|)/\epsilon^2). \quad (38)$$

With high probability, this compressed representation has the following rapid-retrieval property. There is a classical algorithm which, given this classical data and any Pauli $P \in S$, outputs an estimate of $\text{Tr}(\rho P)$ to within ϵ error. The runtime of the algorithm is $O(B)$.

Proof. The compressed representation consists of the $O((\log |S|)/\epsilon^4)$ Bell samples (each one is $2n$ bits) as well as a classical description of each Clifford measurement basis and measurement outcome used in the second stage of the learning protocol. That is, $N = O(\chi(\log |S|)/\epsilon^2)$ Clifford measurement bases $C_1, C_2, \dots, C_N$ (each described by a circuit with $O(n^2)$ one- and two-qubit Clifford gates) and corresponding measurement outcomes $z^1, z^2, \dots, z^N \in \{0, 1\}^n$. Given a Pauli $P \in S$ and this classical data, we can compute an ϵ -error estimate of $\text{Tr}(\rho P)$ in the following way. First, using the Bell samples, we determine if $P \in S_\epsilon$. This step requires us to compute a sample mean over the Bell samples, using runtime $O(n(\log |S|)/\epsilon^4)$. If $P \notin S_\epsilon$, we output 0 as our estimate. If $P \in S_\epsilon$ then we compute the median-of-means estimator from Eq. (34). To do this we have to compute indicator functions $x_P^j \in \{1, 0, -1\}$ that describe the measured outcome of Pauli P for each Clifford measurement basis j , which is given by

$$x_P^j = \langle z^j | C_j^\dagger P C_j | z^j \rangle. \quad (39)$$

The RHS is computed using the stabilizer formalism: we update the Pauli P by conjugating each gate in the circuit C_j one-by-one, and this process takes a total runtime $O(n^2)$ since there are $O(n^2)$ one- and two-qubit Clifford gates in the circuit. The total runtime to extract the estimate of $\text{Tr}(\rho P)$ is therefore

$$O(n(\log |S|)/\epsilon^4 + n^2\chi(\log |S|)/\epsilon^2). \quad \square \quad (40)$$

Lemma 20 forms the basis of several of the two-copy protocols that we present in this work. To use this framework one needs to find an efficiently sampleable fractional coloring of $G(S_\epsilon)$.

A challenge here is that the set S_ϵ and its commutation graph depend in a potentially complicated way on the unknown state ρ . Ideally, we would like to understand any structural properties of this graph that can be leveraged to compute good fractional colorings. In this paper we will only exploit two simple properties: (A) $G(S_\epsilon)$ is an induced subgraph of $G(S)$ and (B) with high probability, $G(S_\epsilon)$ does not have large cliques, as described in Lemma 8, which we restate and prove below.

Lemma 8. *The largest clique in the commutation graph $G(S_\epsilon)$ has size at most $4/\epsilon^2$ with high probability.*

Proof. Recall that S_ϵ satisfies Eq. (36) with high probability. We show that in this case the largest clique in $G(S_\epsilon)$ has size at most $4/\epsilon^2$.

Suppose there is a clique in $G(S_\epsilon)$ of size ω . The vertices of the clique are a set of pairwise anticommuting Pauli operators $P_1, P_2, \dots, P_\omega$. Applying Lemma 14 with these operators and

the state ρ gives

$$\sum_{j=1}^{\omega} \text{Tr}(P_j \rho)^2 \leq 1. \quad (41)$$

On the other hand from [Eq. \(36\)](#) we have $\text{Tr}(P_j \rho)^2 \geq \epsilon^2/4$ for each $1 \leq j \leq \omega$. Plugging into the above gives $\omega \epsilon^2/4 \leq 1$, and therefore the size of the maximal clique is upper bounded as $\omega \leq 4/\epsilon^2$. $\square$

To use our framework to learn Pauli observables $S \subseteq \mathcal{P}^{(n)}$, it suffices to establish a so-called *chi-binding function* for the family of induced subgraphs of $G(S)$. That is, we seek a function $g(\omega)$ such that:

For any induced subgraph G' of $G(S)$ with largest clique of size ω , there is a fractional coloring of G' with size $\chi \leq g(\omega)$.

A statement of this form implies—via [Lemma 20](#) and [Lemma 8](#)—a shadow tomography algorithm that uses two-copy Clifford measurements and has sample complexity

$$O((\log |S|)/\epsilon^4 + g(4/\epsilon^2)(\log |S|)/\epsilon^2). \quad (42)$$

[Lemma 20](#) also gives an upper bound on the runtime of the protocol in terms of the time required to sample from the fractional coloring.

3 Learning all Pauli observables

In [Section 3.1](#) we provide further details of the shadow tomography algorithm for all Paulis stated in [Theorem 7](#). Then, in [Section 3.2](#), we give the proof of [Lemma 11](#) which is used to establish the rapid retrieval Pauli compression stated in [Corollary 12](#).

3.1 Computing a mimicking state

Here we complete the proof of [Theorem 7](#) by showing that a mimicking state can be computed using $O(n \log(n/\epsilon)/\epsilon^4)$ samples of ρ and runtime $\text{poly}(2^n, 1/\epsilon)$.

Algorithm 1 Compute a mimicking state

Input: A precision parameter $\epsilon \in (0, 1)$, $O(n \log(n/\epsilon)/\epsilon^4)$ copies of an unknown n -qubit state ρ , and estimates $\{u_P\}_{P \in \mathcal{P}^{(n)}}$, such that

$$u_P \geq 0 \quad \text{and} \quad |u_P - |\text{Tr}(P\rho)|| \leq \epsilon/4 \quad P \in \mathcal{P}^{(n)}. \quad (43)$$

Output: A classical description of a density matrix σ satisfying the following mimicking state condition with high probability:

$$|\text{Tr}(P\sigma)| \geq \epsilon/4 \quad \text{for all } P \in \mathcal{P}^{(n)} \text{ with } u_P \geq 3\epsilon/4. \quad (44)$$

Algorithm:

1. Set $T = \lceil 64n/\epsilon^2 \rceil + 1$ and $\beta = \sqrt{n/T}$.
 2. Initialize $\omega^{(0)} = \mathbb{1}/2^n$, the maximally mixed state.
 3. For $t = 0, \dots, T - 1$,
 - (a) Search for Pauli $P \in \mathcal{P}^{(n)}$ such that $u_P \geq 3\epsilon/4$ and $|\text{Tr}(P\omega^{(t)}) - u_P| > \epsilon/2$ and $|\text{Tr}(P\omega^{(t)}) + u_P| > \epsilon/2$.
 - (b) If there is no such Pauli P then we are done, since we are guaranteed $|\text{Tr}(P\omega^{(t)})| \geq \epsilon/4$ whenever $u_P \geq 3\epsilon/4$. Output $\sigma = \omega^{(t)}$.
 - (c) Else, do:
 - i. Use $O((\log T)/\epsilon^2)$ copies of ρ to compute an estimate $r_P \in \{+1, -1\}$ such that $r_P = \text{sign}(\text{Tr}(\rho P))$ with probability at least $1 - 0.01/|T|$.
 - ii. Set
$$M^{(t)} = \text{sign}\{\text{Tr}(P\omega^{(t)}) - r_P u_P\} \cdot P. \quad (45)$$
 - iii. Set
$$\omega^{(t+1)} = \frac{\exp\left(-\beta \sum_{\tau=1}^t M^{(\tau)}\right)}{\text{Tr}\left(\exp\left(-\beta \sum_{\tau=1}^t M^{(\tau)}\right)\right)}. \quad (46)$$
 4. Output $\sigma = \omega^{(T)}$.
-

Before showing correctness of the algorithm, let us discuss its runtime and sample complexity. During the course of the algorithm we store a classical representation of the state $\omega^{(t)}$ as a matrix of size $2^n \times 2^n$, and in each of the $T = O(n/\epsilon^2)$ steps we need to exhaustively search over the set of Paulis $\mathcal{P}^{(n)}$, compute Pauli expected values in the state $\omega^{(t)}$, and compute matrix exponentials Eq. (46). The runtime of the algorithm is polynomial in the Hilbert space dimension 2^n and $1/\epsilon$. The sample complexity is $O(T(\log T)/\epsilon^2)$.

Algorithm 1 is a variant of the *matrix multiplicative weights* algorithm, which has been used previously in the context of quantum learning [BKL⁺19, ACH⁺18]. To show correctness we follow the standard analysis, with some modifications.

Lemma 22. ([AK07] Theorem 3.1) Suppose [Algorithm 1](#) reaches step 4. Then the states [Eq. \(46\)](#) satisfy

$$\sum_{t=1}^T \text{Tr}(M^{(t)} \omega^{(t)}) - \lambda_{\min} \left(\sum_{t=1}^T M^{(t)} \right) \leq 2\sqrt{nT}. \quad (47)$$

Here $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue.

Lemma 23. With high probability the output of [Algorithm 1](#) satisfies [Eq. \(44\)](#).

Proof. Let $P^{(t)}$ be the Pauli found at step t , and define

$$y^{(t)} = \text{Tr}(P^{(t)} \omega^{(t)}), \quad r^{(t)} = r_{P^{(t)}}, \quad u^{(t)} = u_{P^{(t)}}. \quad (48)$$

By a union bound, with probability at least 0.99, all estimates $r^{(t)}$ computed during the course of the algorithm satisfy $r^{(t)} = \text{sign}(\text{Tr}(\rho P^{(t)}))$. In the following we show that in this case the output of the algorithm is guaranteed to satisfy [Eq. \(44\)](#).

Note that if the algorithm outputs a state $\sigma = \omega^{(t)}$ for some $t < T$ then the algorithm is correct, as guaranteed by the output condition. To complete the proof, below we show that this is the case—the algorithm never reaches step 4.

Toward a contradiction, suppose the algorithm does reach step 4. By [Lemma 22](#), we have

$$\sum_{t=1}^T \text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} \cdot y^{(t)} \leq \lambda_{\min} \left(\sum_{t=1}^T \text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} \cdot P^{(t)} \right) + 2\sqrt{nT} \quad (49)$$

$$\leq \sum_{t=1}^T \text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} \cdot \text{Tr}(P^{(t)} \rho) + 2\sqrt{nT}. \quad (50)$$

In the last equation, we substituted our state ρ . Equivalently, we have

$$\sum_{t=1}^T \text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} \cdot \left(y^{(t)} - \text{Tr}(P^{(t)} \rho) \right) \leq 2\sqrt{nT}. \quad (51)$$

For each t we have $|y^{(t)} - r^{(t)} u^{(t)}| > \epsilon/2$. But also $|r^{(t)} u^{(t)} - \text{Tr}(P^{(t)} \rho)| \leq \epsilon/4$. Together these conditions imply

$$\text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} = \text{sign}\{y^{(t)} - \text{Tr}(P^{(t)} \rho)\} \quad (52)$$

and

$$|y^{(t)} - \text{Tr}(P^{(t)} \rho)| > \frac{\epsilon}{4}. \quad (53)$$

From [Eq. \(52\)](#) we get

$$\text{sign}\{y^{(t)} - r^{(t)} u^{(t)}\} \cdot \left(y^{(t)} - \text{Tr}(P^{(t)} \rho) \right) = \left| y^{(t)} - \text{Tr}(P^{(t)} \rho) \right| \quad (54)$$

and combining this with [Eq. \(51\)](#) we get

$$\sum_{t=1}^T |y^{(t)} - \text{Tr}(P^{(t)} \rho)| \leq 2\sqrt{nT}. \quad (55)$$

Plugging in Eq. (53) gives

$$T \cdot \frac{\epsilon}{4} \leq 2\sqrt{nT}, \quad (56)$$

and therefore $T \leq 64n/\epsilon^2$. This is a contradiction, since $T = \lceil 64n/\epsilon^2 \rceil + 1$ is larger than this. So the algorithm never reaches step 4, as claimed. $\square$

3.2 Coloring commutation graphs with bounded clique number

In this section we prove Lemma 11, restated below.

Lemma 11. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{P}^{(n)})$, and let ω be the size of the largest clique in G' . The chromatic number of G' is upper bounded as*

$$\chi(G') \leq (2n + 1)^{\omega-1}. \quad (15)$$

Moreover, a coloring with this many colors can be computed by a classical algorithm with runtime $\text{poly}(|G'|, n^\omega)$.

We shall use the following algorithmic version of a result of Gyárfás, which we prove below.

Lemma 24 (Algorithmic version of Thm. 2.4 of [Gyá87]). *Suppose G is a graph on m vertices whose longest induced path has ℓ vertices, with $\ell \geq 1$, and clique number ω . Then there is a classical algorithm which colors G using $\ell^{\omega-1}$ colors and runtime $O(m^2\omega)$.*

Proof of Lemma 11. Let G' be an induced subgraph of $G(\mathcal{P}^{(n)})$. Below we show that $G(\mathcal{P}^{(n)})$, and therefore also G' , does not contain any induced paths with more than $2n + 1$ vertices. The claim then follows by applying Lemma 24.

Suppose $P_1, P_2, \dots, P_s$ is an induced path in $G(\mathcal{P}^{(n)})$, i.e.,

$$P_i P_{i+1} = -P_{i+1} P_i \text{ for } 1 \leq i \leq s-1 \text{ and } [P_i, P_j] = 0 \text{ for } |i-j| \geq 2. \quad (57)$$

Define Pauli operators

$$Q_r = P_1 P_2 \dots P_r \quad 1 \leq r \leq s. \quad (58)$$

We now use Eq. (57) to show that these operators are pairwise anticommuting. To see this note that

$$Q_{r+a} = (P_1 P_2 \dots P_r) (P_{r+1} P_{r+2} \dots P_{r+a}) = - (P_{r+1} P_{r+2} \dots P_{r+a}) (P_1 P_2 \dots P_r), \quad (59)$$

where we used $P_r P_{r+1} = -P_{r+1} P_r$ and the fact that $[P_i, P_j] = 0$ whenever $i \leq r-1$ and $j \geq r+1$. Therefore

$$Q_r Q_{r+a} = -Q_r (P_{r+1} P_{r+2} \dots P_{r+a}) (P_1 P_2 \dots P_r) = -Q_{r+a} Q_r, \quad (60)$$

which shows that $\{Q_j\}_{j \in [s]}$ are pairwise anticommuting Pauli operators.

It is a well known fact that the n qubit Hilbert space does not contain any set of pairwise anticommuting Pauli operators with size greater than $2n + 1$ (see for example Appendix G of Ref. [BMBO20]). Thus $s \leq 2n + 1$. $\square$




Figure 1: An example showing the spanning trees generated by breadth-first search (BFS), depth-first search (DFS), and our algorithm neighbour-first search (NFS) for the cycle on 5 vertices.

The algorithm of [Lemma 24](#) relies on a non-standard graph traversal algorithm, which can be interpreted as a combination of depth-first-search and breadth-first-search. The graph search algorithm begins with an arbitrary seed vertex v in G , and generates a spanning tree of G with root v . We call it *neighbour-first search*.

Algorithm 2 Neighbour-first search (NFS)

Input: Connected graph G , seed vertex v .

Output: Spanning tree T of G with root v .

NFS(G, v):

1. If T is empty, initialize $T = \{v\}$.
 2. For each neighbour w of v which is not yet in T , add w to T as a child of v .
 3. For each child w of v :
 - (a) Do NFS(G, w).
-

The spanning tree T output by [Algorithm 2](#) is associated with a partition of the vertex set of G into levels, which are the vertices at a fixed distance from the root of T . (The number of levels is the depth of the T plus one.)

Lemma 25. *Let graph $G = (V, E)$ have m vertices, clique number ω , and longest induced path with ℓ vertices. [Algorithm 2](#) has runtime $O(|V| + |E|) = O(m^2)$ and outputs a spanning tree T of G with the following properties:*

- The depth of T is no larger than $\ell - 1$.
- The vertices in any level of T induce a subgraph of G with clique number at most $\omega - 1$.

Proof. No vertex can share an edge with any ancestors in T other than its parent, since if it shared an edge with an ancestor higher than its parent then it would have appeared at a higher level as a neighbour of the ancestor. This means a path from the root down T forms an induced path, and the depth of T cannot be longer than the longest induced path in G .

Consider two vertices w_1 and w_2 in the same level t of the spanning tree T . Then the children of w_1 cannot share any edges in G with the children of w_2 . This is because the children of w_1 constitute a connected component of the subgraph of G induced by all vertices that are not in the first t levels of T . Armed with this observation, consider a clique of size ω within a level of T . When combined with the common parent, this would form a clique of size $\omega + 1$ in G , a contradiction. Thus the clique number of any level of T is at most $\omega - 1$.

Finally, similar to breadth-first search or depth-first search, since each edge is examined at most twice, the time complexity is $O(|V| + |E|) = O(m^2)$. $\square$

Proof of Lemma 24. We can prove the theorem by induction on ω . Let $\mathcal{A}_\omega$ denote the coloring algorithm which applies to graphs of clique number ω . When $\omega = 1$, there are no edges and there is an algorithm $\mathcal{A}_1$ which can color the graph using a single color in $O(m^2)$ time. For the inductive step, assume there is a coloring algorithm $\mathcal{A}_{\omega-1}$ using $\ell^{\omega-2}$ colors and runtime $O(m^2\omega)$ for any graph of clique number $\omega - 1$ and longest induced path ℓ .

The coloring algorithm $\mathcal{A}_\omega$ for graphs of clique number ω is as follows. First apply the neighbour-first search algorithm to find spanning tree T of G . Then for each level of T , apply $\mathcal{A}_{\omega-1}$. For each level, we use a disjoint set of colors. Since there are at most ℓ levels in the T , the number of colors used by $\mathcal{A}_\omega$ is at most $\ell^{\omega-1}$ by the induction hypothesis.

It remains to analyze the runtime of $\mathcal{A}_\omega$. Say the neighbour-first search step has runtime at most Cm^2 in the worst case for some constant C . We will show that the runtime of $\mathcal{A}_\omega$ is at most $Cm^2\omega$. From the induction hypothesis, the applications of $\mathcal{A}_{\omega-1}$ have total runtime $\sum_i Cm_i^2(\omega - 1) \leq Cm^2(\omega - 1)$, where m_i is the number of vertices in the i^{th} layer. Here we used $\sum_i m_i^2 \leq (\sum_i m_i)^2 = m^2$. Thus the runtime of $\mathcal{A}_\omega$ is $Cm^2 + Cm^2(\omega - 1) \leq Cm^2\omega$. $\square$

4 Learning fermionic observables

In this section we consider shadow tomography for local fermionic observables.

A system of n fermionic modes is associated with a set of $2n$ Majorana fermion operators, which are mutually anticommuting Hermitian observables $\{c_a\}_{a \in [2n]}$ that act on a Hilbert space of dimension 2^n . We can represent them by a set of Pauli operators satisfying

$$c_1, c_2, \dots, c_{2n} \in \mathcal{P}^{(n)} \quad c_a c_b + c_b c_a = 2\delta_{ab} \mathbb{1}. \quad (61)$$

There are a variety of specific choices (fermion-to-qubit mappings) that satisfy the above, including the Jordan-Wigner mapping [JW28], the Bravyi-Kitaev mapping [BK02], and the ternary tree mapping [Vla19, JKMN20]. The latter two mappings have the desirable property that each Majorana fermion operator is represented by a Pauli operator of low weight $O(\log n)$.

For concreteness, we now review the ternary tree mapping, which is simple enough that it can be understood by looking at Figure 2. We place a qubit at each non-leaf vertex of a complete rooted ternary tree. The total number of qubits is $n = (3^\ell - 1)/2$ where ℓ is

the depth of the tree. The three edges that connect the vertex for a given qubit q to its children are associated with the three single-qubit Pauli operators X, Y, Z acting on qubit q . Each path from the root to a leaf defines an operator which is the tensor product of the Pauli operators which appear on all the edges along the path. One can easily check that these operators are pairwise anticommuting. Moreover, each operator defined in this way has weight exactly $\ell = \log_3(2n + 1)$, and there are exactly $3^\ell = 2n + 1$ of them. We can take all but one of them to be the Majorana fermion operators $\{c_j\}_{j \in [2n]}$.

Below we are interested in k -body fermionic observables as defined in [Section 1.1](#). We write

$$\Gamma(x) = i^{|x| \cdot (|x|-1)/2} c_1^{x_1} c_2^{x_2} \dots c_{2n}^{x_{2n}} \quad x \in \{0, 1\}^{2n} \quad (62)$$

for the Majorana monomials, and

$$\mathcal{F}_k^{(n)} = \{\Gamma(x) : |x| = 2k\} \quad (63)$$

for the set of k -body Majorana operators on n fermionic modes.




Figure 2: The ternary tree fermion-to-qubit mapping of Refs. [\[Vla19, JKMN20\]](#) for $n = 4$ and $\ell = 2$. For example, $c_1 = X_1 X_2 = X \otimes X \otimes \mathbb{1} \otimes \mathbb{1}$.

4.1 Single-copy lower bounds

Theorem 26. *Suppose $1 \leq k \leq \log_3(2n + 1)$. Any (possibly adaptive) single-copy protocol which learns $\text{Tr}(P\rho)$ to precision ϵ for all k -local n -qubit Paulis $P \in \mathcal{P}_k^{(n)}$ with constant probability requires $\Omega(3^k/\epsilon^2)$ copies of ρ .*

Proof. The ternary tree mapping [\[Vla19, JKMN20\]](#) which we reviewed in the previous section shows that, for each $1 \leq k \leq \log_3(2n + 1)$ there is a subset $S \subseteq \mathcal{P}_k^{(n)}$ of 3^k pairwise anticommuting operators within the set of k -local Paulis $\mathcal{P}_k^{(n)}$. See [Figure 2](#). On the other hand, [Lemma 14](#) shows that the commutation index of any set of m pairwise anticommuting observables is at most $1/m$. Applying [Theorem 16](#), we conclude that $\Omega(3^k \epsilon^{-2})$ copies of ρ are needed to learn the expected values of all observables from the set S . Therefore at least this many samples are needed to learn all expected values of observables in $\mathcal{P}_k^{(n)}$. $\square$

The lower bound in [Theorem 26](#) matches the sample complexity of known single-copy protocols such as classical shadows, up to a factor of $k \log n$ [\[HKP20\]](#), see [Table 2](#).

For local fermionic observables, one can obtain a much stronger single-copy sample complexity lower bound which scales polynomially in the system size n . This is telling us that local fermionic observables are much harder to learn than local Pauli observables. The following Theorem implies that there is no sample efficient single-copy protocol for k -body fermionic observables, as stated in [Theorem 3](#).

Theorem 27. *Any (possibly adaptive) single-copy protocol which learns $\text{Tr}(\Gamma\rho)$ to precision ϵ for all k -body Majorana operators Γ on n fermionic modes with constant probability requires number of copies scaling as $\Omega(n^k/\epsilon^2)$, for any fixed $k \geq 1$.*

Establishing [Theorem 27](#) requires a short detour into *Johnson association schemes*. For a subset $L \subseteq \{0, \dots, q-1\}$, define the *generalized Johnson graph* $G(m, q, L)$ to have vertices corresponding to the subsets of $x \subseteq \{1, \dots, m\}$ of size $|x| = q$, and an edge between any x and y such that $|x \cap y| \notin L$. Notice that the commutation graph of the k -body Majorana observables on n fermionic modes is precisely the generalized Johnson graph with L consisting of the *even* integers: $G(\mathcal{F}_k^{(n)}) = G(2n, 2k, \{0, 2, \dots, 2k-2\})$.

In Ref. [\[Lin24\]](#), they show that the Lovasz ϑ -function of the generalized Johnson graph scales like $\vartheta(G(m, q, L)) = \Theta(m^{|L|})$ as m grows for fixed q . Setting L equal to the even integers gives the following conclusion for the Lovasz ϑ -function of the commutation graph of degree- $2k$ Majoranas.

Theorem 28. ([\[Lin24, Theorem 1.2\]](#)) *For any fixed k , $\vartheta(G(\mathcal{F}_k^{(n)})) = \Theta(n^k)$ as n becomes large.*

This resolves Conjecture 4.13 of [\[HO22\]](#), up to the k -dependence of the constant factor in $\vartheta(G(\mathcal{F}_k^{(n)}))$.

Proof of [Theorem 27](#). [Theorem 28](#) and [Lemma 18](#) give

$$\Delta(\mathcal{F}_k^{(n)}) = O(n^k) / \binom{2n}{2k} = O(n^{-k}), \quad (64)$$

since $|\mathcal{F}_k^{(n)}| = \binom{2n}{2k}$. [Theorem 16](#) then implies the sample complexity lower bound $\Omega(n^k/\epsilon^2)$. $\square$

The $\Omega(n^k/\epsilon^2)$ single-copy sample complexity lower bound in [Theorem 27](#) matches what is achieved by the single-copy protocols in [\[BMBO20, JKMN20, WHLB23, HWM⁺22\]](#), up to a factor of $k \log(n)$, see [Table 2](#). It should also be noted that a matching lower bound of $\Delta(\mathcal{F}_k^{(n)}) = \Omega(n^{-k})$ can be shown by finding a large set of mutually commuting k -body fermionic observables.

4.2 Learning 1-body fermionic observables

In the case of 1-body observables there is a simple and practical algorithm for coloring induced subgraphs of $G(\mathcal{F}_1^{(n)})$ with bounded clique number.

Lemma 29. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_1^{(n)})$ of 1-body fermionic observables, and let ω be the size of the largest clique in G' . There is a classical algorithm with runtime $O(n^2\omega)$ that computes a coloring of G' with at most $\omega + 1$ colors.*

Proof. Let $G' = G(S)$ be the subgraph of $G(\mathcal{F}_1^{(n)})$ induced by some subset $S \subseteq \mathcal{F}_1^{(n)}$ of 1-body fermionic observables. Let ω be the maximum size of a clique in G' .

Consider an auxiliary graph $H(S)$ defined as follows. This graph $H(S)$ has $2n$ vertices labeled by the Majorana fermion operators $\{c_1, c_2, \dots, c_{2n}\}$. For each observable $ic_a c_b \in S$, we include an edge $\{c_a, c_b\}$ in $H(S)$. Two elements of S commute if and only if they do not share any Majorana fermion operators. For example, $ic_1 c_2$ anticommutes with $ic_2 c_3$ but commutes with $ic_3 c_4$. Thus a commuting set of 1-body fermionic observables corresponds to a matching in $H(S)$, and partitioning S into commuting sets corresponds to an edge coloring of $H(S)$.

Now observe that our graph of interest G' is the *line graph* of $H(S)$. An edge coloring of $H(S)$ gives a vertex coloring of G' . The edge coloring algorithm of Misra and Gries [MG92] computes an edge-coloring of a graph H using no more than $\deg(H) + 1$ colors, where $\deg(H)$ is the maximum degree of any vertex in H . But the edges connecting to a single vertex in our graph $H(S)$ form a clique in G' , so the degree of $H(S)$ is at most ω . The runtime of the edge coloring algorithm is asymptotically upper bounded by the number of vertices times the number of edges, which in our case is $O(n \cdot n\omega)$. $\square$

It is also possible to directly vertex color the given graph G' using Brooks' theorem, which states that the chromatic number of a graph is at most its maximum degree $+1$, since a high degree vertex also yields a large clique. This argument yields a slightly looser bound of 2ω .

4.3 Learning k -body fermionic observables

We now prove Lemma 9, restated below.

Lemma 9. *Let $k \geq 1$, and let G' be any induced subgraph of the commutation graph $G(\mathcal{F}_k^{(n)})$ of k -body fermionic observables, and let ω be the size of the largest clique in G' . Then the fractional chromatic number of G' satisfies*

$$\chi_f(G') \leq p_k(\omega). \quad (13)$$

where p_k is a polynomial. Moreover, for any $k = O(1)$ we can sample from a fractional coloring of G' with size $p_k(\omega)$ using a classical algorithm with runtime $\text{poly}(n)$. The polynomials for $k = 1, 2$ are $p_1(\omega) = \omega + 1$ and $p_2(\omega) = O(\omega^8)$.

Recall the definition of Majorana monomials from Eq. (62). In the following we shall use the commutation relations of these operators which we now derive. Using Equation (3) we get

$$c_j \Gamma(y) = (-1)^{|y|+y_j} \Gamma(y) c_j \quad y \in \{0, 1\}^{2n} \quad j \in [2n]. \quad (65)$$

Applying the above for all indices j in the support of $x \in \{0, 1\}^{2n}$ gives

$$\Gamma(x) \Gamma(y) = (-1)^{|x||y|+y \cdot x} \Gamma(y) \Gamma(x) \quad x, y \in \{0, 1\}^{2n}. \quad (66)$$

Claim 30. *Suppose $x, y \in \{0, 1\}^{2n}$ are such that $|x|, |y|$ are either both even, or both odd. If $x_j = y_j = 0$ for some $j \in [2n]$ then*

$$[c_j \Gamma(x), c_j \Gamma(y)] = 0 \quad \text{if and only if} \quad [\Gamma(x), \Gamma(y)] = 0. \quad (67)$$

Proof. Follows directly from [Eq. \(66\)](#). $\square$

In our proof, it will be helpful to consider Majorana monomials of both odd and even degree. Write

$$\mathcal{M}_r^{(n)} = \{\Gamma(x) : |x| = r, x \in \{0, 1\}^{2n}\} \quad (68)$$

for the set of degree- r Majorana monomials, so that $\mathcal{M}_{2k}^{(n)} = \mathcal{F}_k^{(n)}$ are the k -body fermionic observables of interest.

Proof. The proof is by induction in r . Our inductive hypothesis is that, for any induced subgraph H of the commutation graph $G(\mathcal{M}_r^{(n)})$ of degree- r Majorana monomials with largest clique of size at most ω , we can sample from a fractional coloring of H with $f_r(\omega)$ colors using a classical algorithm with runtime $t_r(n)$ such that $t_r(n) = \text{poly}(n)$ for any constant $r = O(1)$. Here $f_r(\omega)$ is a polynomial that we determine below. Ultimately we are interested in the even values of r and we have $p_k(\omega) = f_{2k}(\omega)$ where p_k is the polynomial in the statement of [Lemma 9](#).

The base case is $r = 2$. We saw in [Lemma 29](#) that if $S \subseteq \mathcal{M}_2^{(n)}$ and its commutation graph $G(S)$ has no cliques larger than ω , then $G(S)$ can be colored with $\omega + 1$ colors using a classical algorithm with runtime $O(n^2\omega) = O(\text{poly}(n))$ since $\omega \leq |\mathcal{M}_2^{(n)}| = O(n^2)$. Thus $f_2(\omega) = \omega + 1$ and we can efficiently sample from the coloring by selecting a color uniformly at random.

In the following two claims we handle the induction step separately for the odd and even values of r .

Claim 31. *Suppose $r \geq 3$ is odd. Then*

$$f_r(\omega) = r\omega f_{r-1}(\omega). \quad (69)$$

Proof. Let $r \geq 3$ be odd, let G' be an induced subgraph of $G(\mathcal{M}_r^{(n)})$, and suppose the largest clique in G' has size at most ω . Let $V \subseteq \mathcal{M}_r^{(n)}$ be the vertex set of G' . Let $\Gamma(x^1), \Gamma(x^2), \dots, \Gamma(x^L) \in V$ be a maximal set of pairwise anticommuting operators in V . We can construct such a set by starting at any vertex of G' and greedily adding vertices until this is no longer possible. By definition, this set is a clique in G' and therefore $L \leq \omega$.

Let $I \subseteq [2n]$ be the set of all indices of Majoranas that appear in these operators. Since each has weight r , we have

$$|I| \leq r\omega. \quad (70)$$

For convenience let us relabel the Majorana fermion operators so that $I = \{1, 2, \dots, T\}$ where $T \leq r\omega$. Then define

$$S_i = \{\Gamma(z) \in V : z_i = 1, \text{ and } z_j = 0 \text{ for all } 1 \leq j \leq i-1\}. \quad (71)$$

We now show that V can be partitioned as

$$V = S_1 \sqcup S_2 \sqcup \dots \sqcup S_T. \quad (72)$$

By definition, the sets on the RHS are disjoint and each contained in V so all we need to show is that for any $\Gamma(y) \in V$ there is some $i \in T$ such that $\Gamma(y) \in S_i$. So let $\Gamma(y) \in V$ be

given. Since the set $\Gamma(x^1), \Gamma(x^2), \dots, \Gamma(x^L) \in V$ is a maximal set of pairwise anticommuting operators, we must have

$$[\Gamma(y), \Gamma(x^j)] = 0 \quad \text{for some } j \in [L]. \quad (73)$$

Since $|y| = |x| = r$ are both odd we see from Eq. (66) that this implies $y \cdot x^j \neq 0$. Therefore $y_i = 1$ for some index $i \in \{1, 2, \dots, T\}$. Let $\ell \in [T]$ be the smallest index such that $y_\ell = 1$. Then $\Gamma(y) \in S_\ell$ and we have shown V can be partitioned as in Eq. (72).

Now for each $1 \leq i \leq T$ consider the commutation graph $G(S_i)$. Each operator $\Gamma(z) \in S_i$ has $z_i = 1$. From Claim 30, the commutation graph of S_i is therefore unchanged if we flip $z_i \leftarrow 0$ for all $\Gamma(z) \in S_i$. Define

$$S'_i = \{\Gamma(z \oplus \hat{e}_i) : \Gamma(z) \in S_i\}. \quad (74)$$

We have shown that the commutation graph $G(S_i)$ coincides with the commutation graph $G(S'_i)$, where the set $S'_i \subseteq \mathcal{M}_{r-1}^{(n)}$ only contains degree- $(r-1)$ Majorana monomials. Moreover, $G(V)$ does not contain any clique larger than ω , so neither does its induced subgraph $G(S_i)$. Therefore $G(S'_i)$ does not contain any clique of size greater than ω .

By our inductive hypothesis, for each $1 \leq i \leq T$, we can sample efficiently from a fractional coloring of $G(S'_i) = G(S_i)$ with size at most $f_{r-1}(\omega)$. Now let us define a fractional coloring of V in which we choose an index $i \in [T]$ uniformly at random and then sample an independent set in S_i according to the fractional coloring of $G(S_i)$. Note that any independent set in $G(S_i)$ is also an independent set in $G(V)$, so this defines a valid fractional coloring. Moreover, the probability of any vertex $u \in G(V)$ being sampled is equal to the probability that we choose i such that $u \in S_i$ (this probability is $1/T$) times the probability that the sampled independent set of S_i contains u (this is at least $1/f_{r-1}(\omega)$ by our inductive hypothesis). This procedure samples a fractional coloring of size

$$T \cdot f_{r-1}(\omega) \leq r\omega \cdot f_{r-1}(\omega) \quad (75)$$

as claimed. To sample from the fractional coloring, we need to first construct a maximal set of pairwise anticommuting operators $\Gamma(x^1), \dots, \Gamma(x^L)$, from which we can define the set I and the partition Eq. (72). As noted above, this step can be performed by starting at an arbitrary vertex $\Gamma(x^1)$ of G' and then growing the set one operator at a time until this is not longer possible. This step has $\text{poly}(n)$ runtime because the graph has at most $|\mathcal{M}_r^{(n)}| = \binom{2n}{r} = \text{poly}(n)$ vertices. The next step is to choose an index $1 \leq i \leq T$ at random and sample an independent set of S_i uniformly at random using a fractional coloring of $G(S_i)$ which by our inductive hypothesis can be done in $\text{poly}(n)$ time. $\square$

Claim 32. *Suppose $r \geq 4$ is even. Then*

$$f_r(\omega) = (f_{r-1}(\omega))^r. \quad (76)$$

Proof. Let $r \geq 4$ be even, let G' be an induced subgraph of $G(\mathcal{M}_r^{(n)})$, and let ω be the size of the largest clique in G' . Let $V \subseteq \mathcal{M}_r^{(n)}$ be the vertex set of G' . For each $1 \leq i \leq 2n$, define

$$W_i = \{\Gamma(x) \in V : x_i = 1\}. \quad (77)$$

Note any clique in $G(W_i)$ has size at most ω . From [Claim 30](#), the commutation graph of W_i is unchanged if we flip $z_i \leftarrow 0$ for all $\Gamma(z) \in W_i$. Define

$$W'_i = \{\Gamma(z \oplus \hat{e}_i) : \Gamma(z) \in W_i\}. \quad (78)$$

Then $G(W_i) = G(W'_i)$ and any clique in $G(W'_i)$ has size at most ω . Moreover, W'_i is a set of degree- $(r-1)$ Majorana monomials, and by our inductive hypothesis we can efficiently sample a coloring of $G(W'_i)$ with size at most $f_{r-1}(\omega)$. Let q_i be the corresponding fractional coloring of W_i , for each $1 \leq i \leq 2n$.

Now let us randomly sample a set $\Omega \subseteq V$ as follows. First, select independent sets $I_1 \sim q_1, I_2 \sim q_2, \dots, I_{2n} \sim q_{2n}$ according to the fractional colorings described above. Then let

$$\Omega = \{\Gamma(x) \in V : \Gamma(x) \in I_j \text{ for all } j \in [2n] \text{ such that } x_j = 1\} \quad (79)$$

Since $|x| = r$ for all $\Gamma(x) \in V$, we have

$$\Pr(\Gamma(x) \in \Omega) \geq \left(\frac{1}{f_{r-1}(\omega)} \right)^r \quad \Gamma(x) \in V \quad (80)$$

Now let us show that Ω is an independent set in V ; this implies that the above procedure samples from a fractional coloring of V with $(f_{r-1}(\omega))^r$ colors. So suppose $\Gamma(x), \Gamma(y) \in \Omega$. We will show that $\Gamma(x), \Gamma(y)$ commute; equivalently, there is no edge between the corresponding vertices in G' . First suppose $x \cap y = \emptyset$. In this case, since $|x| = |y| = r$ are both even, it follows directly that $[\Gamma(x), \Gamma(y)] = 0$. If on the other hand $x_j = y_j = 1$ for some $j \in [2n]$, Then $\Gamma(x), \Gamma(y) \in W_j \cap \Omega$. But $W_j \cap \Omega \subseteq I_j$ is an independent set in the commutation graph of W_j , and therefore $[\Gamma(x), \Gamma(y)] = 0$.

The algorithm we have described above only involves identifying the subsets of vertices W_i for $1 \leq i \leq 2n$ (which can be done in linear time in the number of vertices of G' , which is upper bounded polynomially in n), and then using $O(n)$ calls to the subroutine for sampling fractional colorings of commutation graphs of degree- $(r-1)$ Majorana monomials with clique number at most ω . Since this subroutine has $\text{poly}(n)$ runtime by our inductive hypothesis, so does the algorithm described above. $\square$

Putting together [Claims 31, 32](#), and [Lemma 29](#) we see that the sizes $f_r(\omega)$ of the fractional colorings are polynomial functions of ω with degree that depends only on r . For even values of r the polynomials $p_k(\omega) = f_{2k}(\omega)$ satisfy the recurrence

$$p_1(\omega) = \omega + 1 \quad \text{and} \quad p_k(\omega) = ((2k-1)\omega p_{k-1}(\omega))^{2k} \quad k \geq 2. \quad (81)$$

For the 2-body and 3-body fermionic observables we get

$$p_2(\omega) = O(\omega^8) \quad , \quad p_3(\omega) = O(\omega^{54}) \quad (82)$$

In general, we have the upper bound

$$p_k(\omega) \leq (2k\omega)^{(2k)^{k+1}}. \quad \square \quad (83)$$

References

- [Aar04] Scott Aaronson. Limitations of quantum advice and one-way communication. In *Proceedings of 19th IEEE Annual Conference on Computational Complexity (CCC 2004)*, pages 320–332, 2004. [p. 11]
- [Aar18] Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, pages 325–338, 2018. [pp. 1, 2]
- [ACH⁺18] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. In *Advances in Neural Information Processing Systems*, volume 31, 2018. [pp. 9, 11, 21]
- [AEHK16] Ali Asadian, Paul Erker, Marcus Huber, and Claude Klöckl. Heisenberg-Weyl observables: Bloch vectors in phase space. *Physical Review A*, 94(1):010301, 2016. [p. 14]
- [AG04] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004. [p. 16]
- [AK07] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *Proceedings of the 39th annual ACM Symposium on Theory of Computing (STOC 2007)*, pages 227–236, 2007. [pp. 9, 22]
- [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 (STOC 2019)*, pages 322–333, 2019. [p. 2]
- [BGKT19] Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3), 2019. [p. 5]
- [BK02] Sergey B Bravyi and Alexei Yu Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210–226, 2002. [pp. 5, 25]
- [BKL⁺19] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP Solvers: Large Speed-Ups, Optimality, and Applications to Quantum Learning. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:14, 2019. [pp. 2, 21]
- [BMBO20] Xavier Bonet-Monroig, Ryan Babbush, and Thomas E O’Brien. Nearly optimal measurement scheduling for partial tomography of quantum states. *Physical Review X*, 10(3):031064, 2020. [pp. 2, 3, 4, 5, 6, 13, 23, 27]

- [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)*, pages 1398–1411, 2021. [p. 2]
- [BWM⁺16] Bela Bauer, Dave Wecker, Andrew J. Millis, Matthew B. Hastings, and Matthias Troyer. Hybrid quantum-classical approach to correlated materials. *Phys. Rev. X*, 6:031045, Sep 2016. [p. 39]
- [CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential separations between learning with and without quantum memory. In *62nd Annual Symposium on Foundations of Computer Science (FOCS 2022)*, pages 574–585. IEEE, 2022. [pp. 4, 6, 13, 14, 15]
- [CGY24] Sitan Chen, Weiyuan Gong, and Qi Ye. Optimal tradeoffs for estimating Pauli observables, 2024. To appear. [p. 14]
- [CW20] Jordan Cotler and Frank Wilczek. Quantum overlapping tomography. *Physical Review Letters*, 124(10):100401, 2020. [pp. 2, 3, 4, 6, 13]
- [dGHG23] Carlos de Gois, Kiara Hansenne, and Otfried Gühne. Uncertainty relations from graph theory. *Physical Review A*, 107(6):062211, 2023. [pp. 15, 36]
- [DKBC21] Charles Derby, Joel Klassen, Johannes Bausch, and Toby Cubitt. Compact fermion to qubit mappings. *Physical Review B*, 104(3):035118, 2021. [p. 5]
- [EHF19] Tim J. Evans, Robin Harper, and Steven T. Flammia. Scalable Bayesian Hamiltonian learning. *arXiv preprint arXiv:1912.07636*, 2019. [pp. 2, 3, 4, 6, 13]
- [FB18] Edoardo Fertitta and George H. Booth. Rigorous wave function embedding with dynamical fluctuations. *Physical Review B*, 98(23):235132, 12 2018. [p. 39]
- [GKK⁺07] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald De Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of the 39th annual ACM Symposium on Theory of Computing (STOC 2007)*, pages 516–525, 2007. [p. 11]
- [GS19] David Gosset and John Smolin. A Compressed Classical Description of Quantum States. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:9, 2019. [p. 11]
- [GWH⁺16] Sheng Guo, Mark A. Watson, Weifeng Hu, Qiming Sun, and Garnet Kin-Lic Chan. Electron Valence State Perturbation Theory Based on a Density Matrix Renormalization Group Reference Function, with Applications to the Chromium Dimer and a Trimer Model of Poly-Phenylenevinylene. *Journal of Chemical Theory and Computation*, 12(4):1583–1591, 4 2016. [p. 5]

- [Gyá87] András Gyárfás. Problems from the world surrounding perfect graphs. *Applicaciones Mathematicae*, 19(3-4):413–441, 1987. [pp. [9](#), [10](#), [14](#), [23](#)]
- [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. [pp. [2](#), [3](#), [4](#), [6](#), [13](#), [26](#)]
- [HKP21] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Information-theoretic bounds on quantum advantage in machine learning. *Physical Review Letters*, 126(19):190505, 2021. [pp. [2](#), [7](#), [8](#), [9](#), [13](#), [18](#)]
- [HO22] Matthew B. Hastings and Ryan O’Donnell. Optimizing strongly interacting fermionic Hamiltonians. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*, page 776–789, 2022. [pp. [15](#), [27](#), [36](#)]
- [HOR⁺22] William J. Huggins, Bryan A. O’Gorman, Nicholas C. Rubin, David R. Reichman, Ryan Babbush, and Joonho Lee. Unbiasing fermionic quantum Monte Carlo with a quantum computer. *Nature*, 603(7901):416–420, 3 2022. [p. [5](#)]
- [HWM⁺22] William J. Huggins, Kianna Wan, Jarrod McClean, Thomas E. O’Brien, Nathan Wiebe, and Ryan Babbush. Nearly optimal quantum algorithm for estimating multiple expectation values. *Physical Review Letters*, 129(24):240501, 2022. [pp. [6](#), [27](#)]
- [JGM19] Andrew Jena, Scott Genin, and Michele Mosca. Pauli partitioning with respect to gate sets. *arXiv preprint arXiv:1907.07859*, 2019. [p. [7](#)]
- [JKMN20] Zhang Jiang, Amir Kalev, Wojciech Mruczkiewicz, and Hartmut Neven. Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning. *Quantum*, 4:276, 2020. [pp. [2](#), [3](#), [4](#), [5](#), [6](#), [13](#), [25](#), [26](#), [27](#)]
- [JW28] P. Jordan and E. Wigner. Über das Paulische Äquivalenzverbot. *Zeitschrift für Physik*, 47(9):631–651, Sep 1928. [pp. [5](#), [25](#)]
- [KGZ15] Alexei A. Kananenka, Emanuel Gull, and Dominika Zgid. Systematically improvable multiscale solver for correlated electron systems. *Physical Review B*, 91(12):121111, 3 2015. [p. [39](#)]
- [Kim95] Jeong Han Kim. The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$. *Random Structures & Algorithms*, 7(3):173–207, 1995. [p. [9](#)]
- [Knu93] Donald E Knuth. The sandwich theorem. *arXiv preprint math/9312214*, 1993. [p. [37](#)]
- [KSH⁺06] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti. Electronic structure calculations with dynamical mean-field theory. *Reviews of Modern Physics*, 78(3):865–951, 8 2006. [p. [39](#)]

- [KWM24] Robbie King, Kianna Wan, and Jarrod McClean. Exponential learning advantages with conjugate states and minimal quantum memory. *arXiv preprint arXiv:2403.03469*, 2024. [p. 9]
- [Lin24] William Linz. L -systems and the Lovász number. *arXiv preprint arXiv:2402.05818*, 2024. [p. 27]
- [LMC⁺23] Yuan Liu, Oinam R. Meitei, Zachary E. Chin, Arkopal Dutt, Max Tao, Isaac L. Chuang, and Troy Van Voorhis. Bootstrap Embedding on a Quantum Computer. *Journal of Chemical Theory and Computation*, 19(8):2230–2247, 4 2023. [p. 39]
- [MG92] Jayadev Misra and David Gries. A constructive proof of Vizing’s theorem. *Information Processing Letters*, 41(3):131–133, 1992. [p. 28]
- [MKCdJ17] Jarrod R. McClean, Mollie E. Kimchi-Schwartz, Jonathan Carter, and Wibe A. de Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A*, 95:042308, Apr 2017. [p. 5]
- [Mon17] Ashley Montanaro. Learning stabilizer states by Bell sampling. *arXiv preprint arXiv:1707.04012*, 2017. [p. 2]
- [PZC23] Linqing Peng, Xing Zhang, and Garnet Kin-Lic Chan. Fermionic Reduced Density Low-Rank Matrix Completion, Noise Filtering, and Measurement Reduction in Quantum Simulations. *Journal of Chemical Theory and Computation*, 19(24):9151–9160, 12 2023. [p. 5]
- [Raz99] Ran Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of the 31st annual ACM Symposium on Theory of Computing (STOC 1999)*, pages 358–367, 1999. [p. 11]
- [SBM06] V.V. Shende, S.S. Bullock, and I.L. Markov. Synthesis of quantum-logic circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25(6):1000–1010, 2006. [p. 9]
- [SKGA17] Sandeep Sharma, Gerald Knizia, Sheng Guo, and Ali Alavi. Combining internally contracted states and matrix product states to perform multireference perturbation theory. *Journal of Chemical Theory and Computation*, 13(2):488–498, 2017. [p. 5]
- [SR19] Ingo Schiermeyer and Bert Randerath. Polynomial χ -binding functions and forbidden induced subgraphs: a survey. *Graphs and Combinatorics*, 35(1):1–31, 2019. [p. 9]
- [SRL12] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *The Journal of Chemical Physics*, 137(22):224109, 12 2012. [p. 5]
- [SS20] Alex Scott and Paul Seymour. A survey of χ -boundedness, 2020. [p. 9]

- [SU11] E.R. Scheinerman and D.H. Ullman. *Fractional Graph Theory: A Rational Approach to the Theory of Graphs*. Dover books on mathematics. Dover Publications, 2011. [pp. 12, 16]
- [TRJ⁺20] Tyler Takeshita, Nicholas C. Rubin, Zhang Jiang, Eunseok Lee, Ryan Babbush, and Jarrod R. McClean. Increasing the representation accuracy of quantum simulations of chemistry without extra quantum resources. *Phys. Rev. X*, 10:011004, Jan 2020. [p. 5]
- [Vla19] Alexander Yu Vlasov. Clifford algebras, spin groups and qubit trees. *arXiv preprint arXiv:1904.09912*, 2019. [pp. 25, 26]
- [VYI20] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F. Izmaylov. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *The Journal of Chemical Physics*, 152(12), 3 2020. [pp. 5, 7]
- [WHLB23] Kianna Wan, William J. Huggins, Joonho Lee, and Ryan Babbush. Matchgate shadows for fermionic quantum simulation. *Communications in Mathematical Physics*, 404(2):629–700, Dec 2023. [pp. 6, 13, 27]
- [WJHSC16] Sebastian Wouters, Carlos A Jiménez-Hoyos, Qiming Sun, and Garnet K.-L. Chan. A Practical Guide to Density Matrix Embedding Theory in Quantum Chemistry. *Journal of Chemical Theory and Computation*, 12(6):2706–2719, 2016. [p. 5]
- [XSW23] Zhen-Peng Xu, René Schwonnek, and Andreas Winter. Bounding the joint numerical range of Pauli strings by graph parameters. *arXiv preprint arXiv:2308.00753*, 2023. [pp. 15, 36]
- [YHM⁺22] Nobuyuki Yoshioka, Hideaki Hakoshima, Yuichiro Matsuzaki, Yuuki Tokunaga, Yasunari Suzuki, and Suguru Endo. Generalized Quantum Subspace Expansion. *Physical Review Letters*, 129(2):020502, 7 2022. [p. 5]
- [ZRM21] Andrew Zhao, Nicholas C. Rubin, and Akimasa Miyake. Fermionic Partial Tomography via Classical Shadows. *Physical Review Letters*, 127(11):110504, 9 2021. [p. 5]

A Proof of Lemma 18

Lemma 18. ([dGHG23, XSW23, HO22]) *Let S be a set of Pauli operators. Then*

$$\Delta(S) \leq \frac{\vartheta(G(S))}{|S|}. \quad (25)$$

Proof. Denote $S = \{P_1, \dots, P_m\}$. Given ρ , let

$$a_j = \text{Tr}(P_j \rho). \quad (84)$$

We aim to show $\sum_j a_j^2 \leq \vartheta(G(S))$.

Consider the observable

$$Q = \sum_j a_j P_j. \quad (85)$$

We have

$$Q^2 = \sum_{j,l} a_j a_l P_j P_l = \frac{1}{2} \sum_{j,l} a_j a_l \{P_j, P_l\}. \quad (86)$$

Note $P_j^2 = \mathbb{1}$ since they are Hermitian unitaries.

Now take the trace with ρ . We get

$$\text{Tr}(Q^2 \rho) = \sum_{j,l} a_j a_l B_{jl} \leq \lambda_{\max}(B) \sum_j a_j^2, \quad (87)$$

where we defined the matrix

$$B_{jl} = \frac{1}{2} \text{Tr}(\{P_j, P_l\} \rho). \quad (88)$$

By positivity of the state ρ , we have $\text{Tr}((Q - \text{Tr}(Q\rho) \mathbb{1})^2 \rho) \geq 0$ and therefore

$$\text{Tr}(Q\rho)^2 \leq \text{Tr}(Q^2 \rho) \quad (89)$$

$$\implies \left(\sum_j a_j^2 \right)^2 \leq \lambda_{\max}(B) \sum_j a_j^2 \quad (90)$$

$$\implies \sum_j a_j^2 \leq \lambda_{\max}(B). \quad (91)$$

B satisfies $B_{jj} = 1 \ \forall j$ and $B_{jl} = 0$ for all edges (j, l) . The latter holds since (j, l) is an edge precisely when $\{P_j, P_l\} = 0$. Positivity of the state ρ implies that B is positive semidefinite, since for any vector $v \in \mathbb{R}^m$

$$v^T B v = \text{Tr} \left(\left(\sum_j v_j P_j \right)^2 \rho \right) \geq 0. \quad (92)$$

Let's now take the supremum of the right-hand-side over all such B to get

$$\sum_j a_j^2 \leq \tilde{\vartheta}(G(S)), \quad (93)$$

where

$$\begin{aligned} \tilde{\vartheta}(G) = \max \{ & \lambda_{\max}(B), \ B \in \mathbb{R}^{m \times m} \\ & \text{s.t., } B_{jj} = 1 \ \forall j, \ B_{jl} = 0 \ \forall (j, l) \in E, \ B \succeq 0 \}. \end{aligned} \quad (94)$$

[Lemma 33](#) completes the proof. □

Lemma 33. ([\[Knu93\]](#)) *The function $\tilde{\vartheta}(G)$ from [Equation \(94\)](#) satisfies $\tilde{\vartheta}(G) \leq \vartheta(G)$.*

Proof. We will use the dual description [Equation \(24\)](#). Let (λ, A) achieve the optimal dual value $\lambda = \vartheta(G)$. Define the $m \times (m+1)$ matrix

$$U = (\vec{1}, \sqrt{\lambda A - J}), \quad (95)$$

where we padded with the all-ones column vector $\vec{1}$ on the left. (Recall J denotes the all-ones matrix.) Let B be any matrix feasible for $\tilde{\vartheta}(G)$. Decompose

$$B = Q^T D Q = V^T V, \quad V = \sqrt{D} Q, \quad (96)$$

where Q is orthogonal and D is diagonal with $D_{11} = \lambda_{\max}(B)$. (The entries of D are the eigenvalues of B .) Now consider the collection of m matrices $\{Y^{(j)}\}$ of size $m \times (m+1)$ given by

$$Y_{ab}^{(j)} = V_{aj} U_{jb}. \quad (97)$$

We have

$$\text{Tr}((Y^{(j)})^T Y^{(l)}) = \left(\sum_a V_{aj} V_{al} \right) \left(\sum_b U_{jb} U_{lb} \right) = \lambda B_{jl} A_{jl}. \quad (98)$$

If $j \neq l$, this is zero, since if (j, l) is an edge in G then $B_{jl} = 0$, and if not then $A_{jl} = 0$. If $j = l$, we get $\text{Tr}((Y^{(j)})^T Y^{(j)}) = \lambda$. Thus $\{Y^{(j)}/\sqrt{\lambda}\}$ are orthonormal when viewed as vectors of dimension $m(m+1)$, and

$$1 \geq \sum_j (Y_{11}^{(j)}/\sqrt{\lambda})^2 = \frac{D_{11}}{\lambda} \sum_j Q_{1j}^2 = \frac{D_{11}}{\lambda} \implies D_{11} \leq \lambda. \quad \square \quad (99)$$

It is in fact true that $\tilde{\vartheta}(G) = \vartheta(G)$, but we only need $\tilde{\vartheta}(G) \leq \vartheta(G)$ for our purposes.

B Learning Greens functions

A quantity of fundamental interest in the chemistry and physics of fermionic system is the Greens function. The k -body Greens function is similar to the k -RDM except that some of the operators have been evolved forward in time. Accordingly, the Greens function can be used to characterize the response of a fermionic system to external perturbations.

Definition 34. *The 1-body Greens function of a state ρ with respect to Hamiltonian H is defined as*

$$G_{ab}(t) = \text{Tr}(i c_a(t) c_b(0) \rho), \quad (100)$$

where

$$c_a(t) = e^{iHt} c_a e^{-iHt}. \quad (101)$$

Examples of dynamical properties one can compute from the 1-body Greens function but not the 1-RDM include electrical conductivity, magnetic and electric susceptibility, and dynamic structure factor. The time-dependent part of the Greens function is often essential for characterizing interesting phases of matter and the presence of certain quasiparticles. Many impurity model schemes for converging finite quantum simulations of condensed fermionic systems towards their thermodynamic limit also require the time-dependent part of the

Greens function. Such methods include dynamical mean-field theory (DMFT) [KSH⁺06] and self-energy embedding theory [KGZ15]. Using quantum computers as impurity model solvers in this context has been explored in papers such as [BWM⁺16, LMC⁺23].

In this work we do not give a particularly efficient method for computing Greens functions at non-zero times (in the limit of zero time, the one-body Greens function is the 1-RDM). However, we are able to show that one can compute time-derivatives of the Greens function at $t = 0$ for sparse Hamiltonians. One can then use these time derivatives to reconstruct the Greens function using a Taylor expansion. Prior work developing methods for DMFT has used this same approach to reconstructing and embedding Greens functions [FB18].

The value $G_{ab}(0)$ of the Greens function at time zero is simply the 1-RDM, which was tackled in Lemma 29. The q^{th} derivative at time zero is given by

$$G_{ab}^{(q)}(0) = \text{Tr}(i\mathcal{L}_H^q(c_a)c_b\rho), \quad (102)$$

where $\mathcal{L}_H(X) = i[H, X] = i(HX - XH)$ denotes the *Lie bracket*, or *commutator*. $\mathcal{L}_H^q$ denotes the t -fold commutator; for example $\mathcal{L}_H^2(X) = -[H, [H, X]]$.

We will require the Hamiltonian H to be sparse, according to the following definition.

Definition 35. *A Hamiltonian H is s -sparse if each Majorana mode appears in at most s terms.*

One should view $\{G_{ab}^{(q)}(0)\}_{a,b}$ as a $n \times n$ matrix for each q , and we would like to learn each entry to precision ϵ . For a given t , the naive strategy of measuring one-at-a-time requires $O(n^2/\epsilon^2)$ copies of ρ . Let Hamiltonian H be k -body and s -sparse. In this section, we give a quantum algorithm which exploits entangled measurements on $\rho \otimes \rho$ to achieve a sample complexity of $\tilde{O}(\log n/\epsilon^4)$, depending only logarithmically on the system size n .

Theorem 36. *Suppose ρ is an unknown state on n fermion modes, and H a k -body and s -sparse Hamiltonian. There is an algorithm using entangled measurements on two copies $\rho \otimes \rho$ at a time which can estimate all $\{G_{ab}^{(q)}(0)\}_{a,b}$ to precision ϵ with high probability using $\tilde{O}((2skq)^{5q} \log n/\epsilon^4)$ total copies of ρ . Moreover, the algorithm runs in time poly n .*

Proof of Theorem 36. Let's examine the operators $\{i\mathcal{L}_H^q(c_a)c_b\}_{a,b}$ more closely by expanding

$$\mathcal{L}_H^q(c_a) = \sum_{\Gamma \in S_{H,a}^{(q)}} h_{a,\Gamma}^{(q)} \Gamma \quad (103)$$

$$i\mathcal{L}_H^q(c_a)c_b = \sum_{\Gamma \in S_{H,a}^{(q)}} ih_{a,\Gamma}^{(q)} \Gamma c_b \quad (104)$$

$S_{H,a}^{(q)}$ denotes the Majorana monomials on which $\mathcal{L}_H^q(c_a)$ has support. Note that all $\Gamma \in S_{H,a}^{(q)}$ have odd degree. Let's assume the original Hamiltonian H was normalized so that the coefficients in the Majorana basis have absolute value at most 1; this implies all $|h_{a,\Gamma}^{(q)}| \leq 1$.

The following lemma makes a crucial observation that the number of terms from H which survive in the expansion of $\mathcal{L}_H^q(c_a)$ is bounded independent of the system size n . For example, the terms which survive in $[H, c_a]$ are those which act on the fermion associated to Majorana mode c_a , of which there are at most s .

Lemma 37. $|S_{H,a}^{(q)}| \leq s^q (2k)^{q-1} (q-1)!$.

Proof of Lemma 37. The proof is a short combinatorial calculation. The degree of the operators in $S_{H,a}^{(q)}$ are upper bounded by $(2k-2)q+1$. This is because we increase the degree by $(2k-2)$ each time we take the Lie bracket, and initially the degree is 1. Using s -sparsity of H , we can write a recursion upper bounding $|S_{H,a}^{(q)}|$:

$$|S_{H,a}^{(q)}| \leq |S_{H,a}^{(q-1)}| \cdot s \cdot ((2k-2)(q-1) + 1). \quad (105)$$

Using $|S_{H,a}^{(0)}| = 1$, we get

$$|S_{H,a}^{(q)}| \leq s^q ((2k-2)(q-1) + 1)((2k-2)(q-2) + 1) \dots (2k-1) \quad (106)$$

$$\leq s^q (2k)^{q-1} (q-1)! \quad (107)$$

□

The goal is to estimate $\text{Tr}(i\mathcal{L}_H^q(c_a)c_b\rho)$ to precision ϵ for all (a, b) . By a triangle inequality on Equation (104), it is sufficient to estimate $\text{Tr}(\Gamma c_b\rho)$ to precision $\epsilon/|S_{H,a}^{(q)}|$ for every $\Gamma \in S_{H,a}^{(q)}$.

Thus we focus on learning the set of Majorana operators

$$\mathcal{S} = \{\Gamma c_b : \Gamma \in S_{H,a}^{(q)}, j = 1, \dots, m\}. \quad (108)$$

The number of Majorana operators we are required to learn could be as large as $|S_{H,a}^{(q)}|m$. Thus we do *not* want to measure these one at a time; rather we would like to parallelize the learning of these operators by using entangled measurements. To this end, we establish the following chi-binding result.

Lemma 38. *Let G' be any induced subgraph of the commutation graph $G(\mathcal{S})$, and let ω be the size of the maximal clique in G' . Then we can efficiently find a coloring of G' with at most $O(s^q (2k)^{q+2} q^2 (q!) \omega)$ colors.*

Proof of Lemma 38. We will show that the degree of G' is bounded by $O(s^q (2k)^{q+2} q^2 (q!) \omega)$. Then a greedy coloring algorithm is sufficient to establish the result.

We begin with an initial lemma.

Lemma 39. *Fix $\Gamma \in S_{H,a}^{(q)}$. The number of indices b for which Γc_b is a vertex in G' is upper bounded by 2ω .*

Proof of Lemma 39. This argument follows the same idea as the degree bound in the proof of Lemma 29. Let $B = \{b : \Gamma c_b \in G'\}$. We seek to bound $|B| = O(\omega)$. Let's split B into two subsets:

$$B \setminus \Gamma = \{b : \Gamma c_b \in G', c_b \notin \Gamma\} \quad (109)$$

$$B \cap \Gamma = \{b : \Gamma c_b \in G', c_b \in \Gamma\} \quad (110)$$

(The notation $c_b \in \Gamma$ indicates that c_b appears as a factor in Γ .) The operators $\{\Gamma c_b\}_{b \in B \setminus \Gamma}$ form a mutually anticommuting set, and likewise for the operators $\{\Gamma c_b\}_{b \in B \cap \Gamma}$. Using the bound on clique size ω , we get $|B \setminus \Gamma| = \omega$ and $|B \cap \Gamma| = \omega$ completing the proof. □

Now fix a single operator $\Gamma_{0c_{b_0}}$ in G' . We will upper bound the degree of $\Gamma_{0c_{b_0}}$ in G' . Suppose vertex Γ_{c_b} forms an edge with $\Gamma_{0c_{b_0}}$ in G' . This means Γ_{c_b} anticommutes with $\Gamma_{0c_{b_0}}$. In order for Γ_{c_b} and $\Gamma_{0c_{b_0}}$ to anticommute, they must overlap on an odd number of Majorana modes.

Case 1. Γ overlaps with $\Gamma_{0c_{b_0}}$ on at least one Majorana mode. By a similar combinatorial argument as the proof of [Lemma 37](#), the number of $\Gamma \in \bigcup_a S_{H,a}^{(q)}$ overlapping with any given Majorana mode is upper bounded by

$$s^q((2k-2)q+2)((2k-2)(q-1)+2)\dots(2k) \leq s^q(2k)^q q!. \quad (111)$$

Multiplying by the number of single Majorana factors in $\Gamma_{0c_{b_0}}$, we get that there are at most

$$s^q(2k)^q q! \cdot ((2k-2)q+2) \leq s^q(2k)^{q+1} q(q!) \quad (112)$$

$\Gamma \in S_{H,a}^{(q)}$ which overlap with $\Gamma_{0c_{b_0}}$ on at least one mode. With Γ fixed, there are at most 2ω choices for b such that $\Gamma_{c_b} \in G'$ by [Lemma 39](#). Thus there are overall at most

$$2s^q(2k)^{q+1} q(q!) \omega \quad (113)$$

vertices Γ_{c_b} forming an edge with $\Gamma_{0c_{b_0}}$ such that Γ overlaps with $\Gamma_{0c_{b_0}}$ on at least one mode.

Case 2. Γ is disjoint from $\Gamma_{0c_{b_0}}$. Since Γ_{c_b} and $\Gamma_{0c_{b_0}}$ anticommute, necessarily c_b must appear as a factor in $\Gamma_{0c_{b_0}}$; we can write $c_b \in \Gamma_{0c_{b_0}}$.

Fix b such that $c_b \in \Gamma_{0c_{b_0}}$, and consider a new graph $G^{(b)}$ whose vertices are the Γ such that Γ_{c_b} forms an edge with $\Gamma_{0c_{b_0}}$ and Γ is disjoint from $\Gamma_{0c_{b_0}}$. Include an edge between Γ and Γ' in $G^{(b)}$ if the Γ and Γ' anticommute. Note that the operators Γ and Γ' have odd degree as products of Majoranas.

The graph $G^{(b)}$ is very dense, and is close to the complete graph. In fact, the *co-degree* of $G^{(b)}$ is bounded; that is, the degree of the *complement* of $G^{(b)}$. This is because disjoint odd-degree Majoranas anticommute, so in order for Γ and Γ' *not* to form an edge, the operators Γ and Γ' must overlap. To examine the co-degree of $G^{(b)}$, we can perform a similar combinatorial calculation as in Case 1 to get an upper bound

$$\text{codeg}(G^{(b)}) \leq s^q(2k)^{q+1} q(q!). \quad (114)$$

The co-degree bound means that we can partition $G^{(b)}$ into cliques using at most $\text{codeg}(G^{(b)}) + 1$ cliques; this corresponds to a greedy coloring of the complement of $G^{(b)}$. To each clique we can apply the assumption on the clique number ω to get a bound on the overall size of $G^{(b)}$:

$$|G^{(b)}| \leq (\text{codeg}(G^{(b)}) + 1) \cdot \omega \leq O(s^q(2k)^{q+1} q(q!) \omega). \quad (115)$$

With at most $(2k-2)q+2$ choices for b , we conclude that there are at most

$$O(s^q(2k)^{q+2} q^2(q!) \omega) \quad (116)$$

vertices Γ_{c_b} forming an edge with $\Gamma_{0c_{b_0}}$ such that Γ is disjoint from $\Gamma_{0c_{b_0}}$. $\square$

Combining [Lemma 38](#) with [Lemma 20](#) completes the proof of [Theorem 36](#). At last invoking [Lemma 37](#), we can bound the final sample complexity as follows:

$$N = \tilde{O}(|S_{H,a}^{(q)}|^4 \log m / \epsilon^4) + O(s^q (2k)^{q+2} q^2 (q!) \cdot |S_{H,a}^{(q)}|^2 / \epsilon^2) \cdot \tilde{O}(|S_{H,a}^{(q)}|^2 \log m / \epsilon^2) \quad (117)$$

$$= \tilde{O}(s^q (2k)^{q+2} q^2 (q!) \cdot |S_{H,a}^{(q)}|^4 \cdot \log m / \epsilon^4) \quad (118)$$

$$= \tilde{O}(s^{5q} (2k)^{5q-2} q^3 (q-1)!^5 \log m / \epsilon^4) \quad (119)$$

$$= \tilde{O}((2skq)^{5q} \log m / \epsilon^4). \quad (120)$$

□