Abelian State Hidden Subgroup Problem: Learning Stabilizer Groups and Beyond

Marcel Hinsche,* Jens Eisert, and Jose Carrasco†
Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin 14195, Germany

(Received 29 October 2025; revised 19 February 2026; accepted 25 February 2026; published 26 May 2026)

Identifying the symmetry properties of quantum states is a central theme in quantum information theory and quantum many-body physics. In this work, we investigate quantum learning problems in which the goal is to identify a hidden symmetry of an unknown quantum state. Building on the recent formulation of the state hidden subgroup problem (StateHSP), we focus on abelian groups and develop an efficient quantum algorithm that learns any hidden symmetry subgroup using a generalized form of Fourier sampling. We showcase the versatility of the approach in three concrete applications: These are learning (i) qubit and qudit stabilizer groups, (ii) cuts along which a state is unentangled, and (iii) hidden translation symmetries. Through these applications, we reveal that well-known quantum learning primitives, such as Bell sampling and Bell difference sampling, are, in fact, special cases of Fourier sampling. Our results highlight the broad potential of the StateHSP framework for symmetry-based quantum learning tasks and provide protocols that are easier to implement on near-term quantum devices.

DOI: 10.1103/6frk-891j

I. INTRODUCTION

Symmetry lies at the heart of modern physics, shaping everything from the conservation laws that govern classical and quantum mechanics to the rich classification of phases of matter in quantum many-body systems. At its core, symmetry captures what remains invariant under specific transformations—typically described by mathematical groups—providing a unifying framework to study the fundamental properties of complex quantum systems.

In the rapidly evolving field of quantum information science, symmetries continue to play a particularly prominent role. They inform the design and analysis of diverse quantum protocols offering both conceptual clarity and practical advantages. In quantum simulation, symmetry can dramatically reduce resource overhead and enhance accuracy. Symmetries are likewise crucial for quantum error correction [1] and mitigation strategies, particularly those that rely on entangled copies of quantum states or exploit invariant subspaces [2], as well as for the robust and scalable benchmarking and certification of quantum devices [3,4]. In quantum many-body and condensed matter physics, symmetries are essential for understanding conventional phases and exploring exotic topological and symmetry-protected order [5].

Furthermore, symmetries are a driving force behind many of the most powerful quantum algorithms. For instance, group-theoretic structures underpin the celebrated *quantum Fourier transform* (QFT), a key component in algorithms such as Shor's for factoring large integers [6,7] and other hidden subgroup problems. By systematically exploiting group symmetries, researchers are able to craft algorithms that outperform their classical counterparts.

At the same time, quantum property testing [8] and quantum learning theory [9,10] have emerged as central frameworks for understanding the capabilities and limitations of extracting information from quantum systems. The construction of learning and testing algorithms—as well as the proof of their optimality—often relies on the structural constraints introduced by symmetry, which enable the design of more efficient procedures that exploit invariance to reduce complexity.

A. State hidden subgroup problem

This ubiquity of symmetries in quantum states naturally raises the question: can we *learn* an unknown symmetry group directly from copies of a quantum state? In this vein, Ref. [11] has recently introduced the *state hidden subgroup problem* (StateHSP) as the task of identifying a hidden symmetry subgroup of some parent group G leaving a quantum state invariant.

*Contact author: m.hinsche@fu-berlin.de
†Contact author: jose.carrasco@fu-berlin.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Definition 1 (State hidden subgroup problem (StateHSP)). Let G be a finite group with a unitary representation $R : G \to \mathrm{U}(\mathcal{H})$ acting on the Hilbert space $\mathcal{H}$ and let $H \leq G$ be a subgroup of G. Assume that you have access to copies of an unknown quantum state vector $|\psi\rangle \in \mathcal{H}$ that is promised to have the following properties:

(1) $\forall\, h \in H, \quad R(h)\,|\psi\rangle = |\psi\rangle.$
(2) $\forall\, g \notin H, \quad |\langle\psi|R(g)|\psi\rangle| \leq 1 - \epsilon.$

The problem is to identify H.

The promise ensures that exact symmetries leave $|\psi\rangle$ invariant while any other group element perturbs it by at least ϵ. As a concrete application of the StateHSP framework, Ref. [11] has given an efficient algorithm for the *hidden cut problem*—a generalization of entanglement testing. This problem asks to identify the cuts across which a given state takes the form of a product state. Their algorithm extends to an efficient algorithm for the StateHSP over arbitrary finite abelian groups using $\mathrm{poly}(\log|G|, 1/\epsilon)$ copies of the state. The result demonstrates the framework's power for quantum learning tasks.

The StateHSP formulation builds a bridge between quantum learning theory and traditional quantum algorithms. It generalizes the standard *hidden subgroup problem* (HSP) [6] by hiding subgroup structure in the invariance properties of an unknown quantum state rather than in a classical oracle. In particular, in the extreme case of $\epsilon = 1$, any $g \notin H$ yields an orthogonal state, recovering the familiar coset-state formulation of the standard HSP. However, unlike the HSP, the StateHSP admits no classical analog, since its input is inherently quantum.

Despite these parallels, the new lens offered by the StateHSP formulation remains largely unexplored: Can other symmetry-related learning problems be formulated as instances of the StateHSP? What is the optimal way to solve the StateHSP and how does it depend on the parameter ϵ?

B. Main results

In this work, we focus on the *abelian* StateHSP and make three main technical and conceptual contributions (see Fig. 1):

(1) We show that *Fourier sampling*, a standard quantum algorithmic technique for the HSP, enables an efficient solution to any abelian StateHSP using only $O(\log|G|/\epsilon)$ copies. Our analysis is straightforward and significantly simpler than that in Ref. [11]. As such, it is expected to be more amenable to experimental implementation.
(2) We demonstrate applications of the abelian StateHSP framework to new and existing problems—including learning stabilizer groups for qubits and qudits, the hidden cut problem, and learning translational symmetries.
(3) We show that for certain groups and representations, Fourier sampling can be implemented without generalized phase estimation [12]. Besides enabling more efficient implementations—and hence contributing to the quest of finding simpler quantum algorithms that are more feasible to be implemented [13]—this observation reveals that established measurement routines in the quantum learning literature,

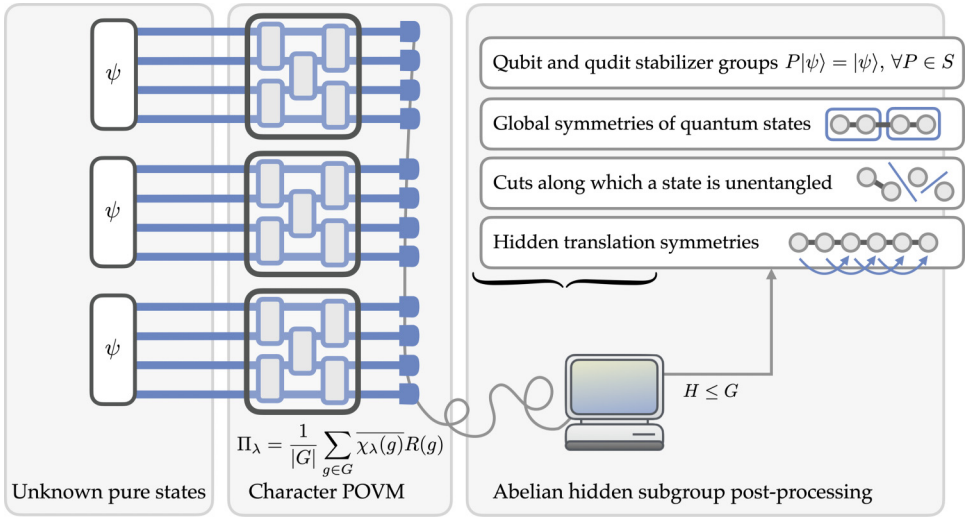


FIG. 1. Overview of our work on the abelian StateHSP. We perform Fourier sampling by measuring the character positive operator valued measure on $m = O(\log|G|/\epsilon)$ copies of the unknown input state vector $\psi = |\psi\rangle\langle\psi|$, followed by abelian hidden subgroup problem postprocessing to recover the hidden subgroup $H \leq G$. The character positive operator valued measure is efficiently implementable and corresponds to well-known quantum learning measurement routines.

such as *Bell sampling* [14–16] and *Bell difference sampling* [17–19], are in fact special cases of Fourier sampling.

We now elaborate on each of these contributions in more detail.

1. Solving the abelian StateHSP via Fourier sampling

The key technique for solving the standard abelian HSP is *Fourier sampling* [6,20,21], which uses the regular representation of G to extract information about the hidden subgroup. In the StateHSP setting, which involves an arbitrary representation R of G, a natural generalization of Fourier sampling can be applied, as observed in Ref. [11]. Specifically, Fourier sampling with respect to R corresponds to measuring the positive operator valued measure (POVM) that we call *character POVM*, whose outcomes λ correspond to characters χ_λ of the abelian group G. Measuring

$$\Pi_\lambda = \frac{1}{|G|}\sum_{g\in G}\overline{\chi_\lambda(g)}R(g) \tag{1}$$

on a copy of the unknown state vector $|\psi\rangle$ yields the output distribution $q_\psi(\lambda) = \mathrm{tr}(\Pi_\lambda|\psi\rangle\langle\psi|)$, which is supported only on $H^\perp$, the dual of the hidden subgroup H. Hence, repeated sampling from q_ψ gradually reveals H.

The key technical question—also studied in Ref. [11] for the hidden cut problem—is how the parameter ϵ, absent in the standard HSP, affects sample complexity. For the hidden cut problem, Ref. [11] relied on measuring multiple copies of $|\psi\rangle$ simultaneously to amplify the orthogonality condition $|\langle\psi|R(g)|\psi\rangle| \le 1-\epsilon$ to $\left|\langle\psi|^{\otimes t}R(g)^{\otimes t}|\psi\rangle^{\otimes t}\right| \le (1-\epsilon)^t$ for $g \notin H$, where t is the number of coherently accessed copies. They combined this idea with an adaptive variant of Fourier sampling and a technically involved analysis.

In contrast, we show that such amplification and adaptivity are unnecessary. For any abelian StateHSP, nonadaptively sampling $\lambda \sim q_\psi$ a total of $O(\log|G|/\epsilon)$ times and applying standard abelian HSP post-processing (see Ref. [6]), one can recover H efficiently in polynomial time. This is proved via a straightforward anticoncentration argument.

Theorem 1 (Efficient algorithm for the abelian StateHSP). For any finite abelian group G and $\epsilon > 0$, there is a polynomial-time quantum algorithm that, with $O(\log|G|/\epsilon)$ copies of the unknown state vector $|\psi\rangle$, identifies the hidden subgroup $H \le G$.

2. Implementing Fourier sampling

The standard technique for implementing Fourier sampling is *generalized phase estimation* [12, Chapter 8], which uses an auxiliary register and controlled application of $R(g)$ together with the quantum group Fourier transform (see Sec. II C). While broadly applicable to all finite abelian groups, this technique requires additional qubits and complex controlled operations, posing practical challenges for near-term devices.

An alternative strategy, conceptually simpler but potentially less general, is to directly measure in a joint eigenbasis of the representation $\{R(g)\}_{g\in G}$. Whether this approach is feasible depends on both the group G and the representation R, and in general, it is unclear how to implement such a measurement. Nevertheless, we show that in several applications this approach is natural and easily realized. Crucially, this viewpoint reveals that quantum learning primitives like *Bell sampling* [14–16], *Bell difference sampling* [17–19], and *computational difference sampling* [18,22] are, in fact, special cases of Fourier sampling, providing a unifying framework and enabling more efficient implementations.

3. Learning stabilizer groups

Learning stabilizer states is a well-studied problem in the qubit case: The first algorithm to identify an unknown stabilizer state has been given by Aaronson and Gottesman in Ref. [23] featuring a sample complexity of $O\left(n^2\right)$. Following up on this work, Montanaro [16] has proposed an algorithm based on Bell sampling that achieves the optimal $O(n)$ sample complexity. Further studies [14,18,36] have demonstrated how to learn the stabilizer group of an unknown input state that is not necessarily a stabilizer state but possesses a nontrivial stabilizer group. The key measurement routine used in this body of work is that of *Bell difference sampling*, first introduced in Ref. [19]. More recent work [17,24] has applied Bell difference sampling to *agnostic learning* stabilizer states, where the goal is to output the closest stabilizer state without any promise on the unknown state.

In contrast, stabilizer learning for qudits (with local dimension $d > 2$) remains underexplored. The sole work addressing qudits explicitly [25] shows that the key subroutine for qubits—Bell difference sampling—fails to provide useful information in the qudit case.

In this work, we demonstrate that learning stabilizer groups can be formulated as an instance of the abelian StateHSP framework. The key trick is to focus on Pauli group up to phases which is isomorphic to $G = \mathbb{Z}_d^{2n}$ where d is the local dimension of the qudits. Formally, we define the following problem: Let $\mathcal{P}_n$ denote the generalized n-qudit Pauli group. Then a stabilizer group S is an abelian subgroup of $\mathcal{P}_n$ that does not contain nontrivial multiples of the identity operator. Denote by $\mathrm{P}_n := \mathcal{P}_n/\langle\omega I\rangle$ and $\mathrm{S} := S/\langle\omega I\rangle$ the corresponding *phaseless* or *projective* Pauli group and its induced stabilizer subgroup, where $\omega = e^{2\pi i/d}$.

Definition 2 (Hidden stabilizer group problem). Let $|\psi\rangle$ be a state vector on n qudits and suppose that there is a nontrivial stabilizer group $S \subseteq \mathcal{P}_n$ such that

$$P|\psi\rangle = |\psi\rangle, \quad \forall P \in S, \tag{2}$$

while for all $P \in \mathrm{P}_n \setminus \mathrm{S}$, $|\langle\psi|P|\psi\rangle| \leq 1 - \epsilon$. The hidden stabilizer group problem asks to identify S given access to copies of $|\psi\rangle$.

By applying the Fourier sampling approach, we provide a unified algorithm for this problem for qubits and qudits. Moreover, we show that in the qubit case, Fourier sampling corresponds exactly to the familiar Bell difference sampling routine [16,18] on four copies of $|\psi\rangle$, whereas in the qudit case, it gives rise to a novel measurement routine.

Theorem 2 (Unified algorithm for learning stabilizer groups—Informal version of Theorem 6). There is an efficient nonadaptive quantum algorithm for the hidden stabilizer group problem on n qudits which uses $O(n/\epsilon)$ copies of the state and runs in polynomial time acting coherently on at most $O(d)$ copies.

We note that this algorithm can be applied to learn qudit stabilizer and t-doped stabilizer states [26] for $t = O(\log n)$, since, as shown in Refs. [14,18], learning such states reduces to learning their stabilizer group, which allows to compress the nonstabilizerness in a small part of the state. In this sense, our algorithm is more versatile than the one presented in Ref. [25] for qudits which is limited to learning nondoped stabilizer states. Moreover, by detecting global stabilizer-like symmetries, our algorithm naturally extends to identifying features associated with *symmetry-protected topological* (SPT) order [5,27], offering a connection to the study of quantum phases of matter in condensed matter physics.

In addition, we note that Refs. [17,24] have recently constructed algorithms for agnostic learning of stabilizer states in the qubit case. These algorithms crucially rely on Bell difference and, in particular, exploit certain concentration properties of the Bell difference sampling distribution. We observe that the distribution q_ψ arising in our qudit setting exhibits analogous concentration properties. We thus expect that the agnostic stabilizer learning algorithms from Refs. [17,24] can be generalized to the qudit case by suitably replacing Bell difference sampling with our qudit measurement routine.

Lastly, we point out that the same algorithm can also be applied in a scenario where one is given many different inputs states sharing a common hidden stabilizer group instead of many copies of a single state. Hence, one can learn a stabilizer code from access to different states in the code space.

4. Improvement for the hidden cut problem via Bell sampling

We revisit the hidden cut problem introduced in Ref. [11] as another instance of the StateHSP framework. We directly focus on the general version of the problem involving potentially many cuts, called the hidden many-cut problem, defined as follows:

Definition 3 (Hidden many-cut problem). Let $|\psi\rangle$ be a state vector on n qubits. Suppose that $|\psi\rangle$ is a product of $m \geq 1$ factor state vectors

$$|\psi\rangle = |\phi_1\rangle_{C_1} \otimes \cdots \otimes |\phi_m\rangle_{C_m}, \tag{3}$$

for some partition $C_1 \sqcup \cdots \sqcup C_m = [n]$ such that each factor state vector $|\phi_k\rangle_{C_k}$ is at least $\tilde{\epsilon}$-far from any separable state on $|C_k|$ qubits. The hidden many-cut problem asks to identify the set partition $C_1 \sqcup \cdots \sqcup C_m$, given copies of $|\psi\rangle$.

We emphasize that the parameter $\tilde{\epsilon}$ used in the hidden many-cut problem does not coincide with the ϵ in Definition 5. Rather, the two are quantitatively related, with $\tilde{\epsilon}^2$ scaling on the order fo ϵ (see also Proposition 1). We adopt the notation $\tilde{\epsilon}$ here to remain consistent with the formulation and complexity bounds in Ref. [11], where the resource requirements are expressed in terms of this parameter. In Ref. [11], the authors have given an algorithm based on generalized phase estimation requiring controlled-SWAP operations. In particular, their algorithm solves this problem using $O\left(n/\tilde{\epsilon}^2\right)$ many samples obtained by measuring the output of circuits of depth $O\left(n^2\right) + O\left(\log \tilde{\epsilon}^{-1}\right)$ acting coherently on $O\left(1/\tilde{\epsilon}^2\right)$ copies and $O(n)$ auxiliary qubits at a time. Here, we improve upon this significantly by showing that the problem can be solved by a direct implementation of the character POVM as hinted at earlier. Interestingly, this direct implementation corresponds to another well-known quantum learning subroutine, namely *Bell sampling* [16].

Theorem 3 (Improved hidden-cut algorithm—Informal version of Theorem 8). There is an efficient nonadaptive quantum algorithm for the hidden many-cut problem on n qubits which uses $O\left(n/\tilde{\epsilon}^2\right)$ copies of the state and runs in polynomial time, requiring circuits of constant depths acting on two copies at a time using no additional auxiliary qubits.

This improves upon Ref. [11] in all requirements, namely the circuit depth, the number of coherently accessed copies as well as removing the need for auxiliary registers while keeping the sample and time complexity the same.

5. Identifying translational symmetries

As a novel application of the StateHSP framework, we formulate the problem of identifying a hidden translational symmetry of a state. Here, for simplicity, we focus on qubits on a ring in one spatial dimension, although the framework immediately generalizes to translations in higher dimensions, e.g., qubits on a two-dimensional cubic lattice.

Definition 4 (Hidden translation problem). Let $|\psi\rangle$ be a state vector on n qubits on a ring. Suppose that $|\psi\rangle$ is translation-invariant under some subgroup of translations of the qubits. The hidden translation symmetry problem asks to identify this subgroup.

Since translations form a representation of the abelian cyclic group $\mathbb{Z}_n$, this problem is an instance of the abelian StateHSP and by virtue of our Theorem 1 can be solved efficiently by measuring with the corresponding character POVM.

Theorem 4 (Efficient algorithm for the hidden translation problem—Informal version of Theorem 9). There is an efficient nonadaptive quantum algorithm for the hidden translation problem on n qubits which uses $O(\log n)$ copies of the state and runs in polynomial time acting on a single copy of $|\psi\rangle$ at a time.

C. Technical overview

1. Dependence of the sample complexity on ϵ

To establish our main result (Theorem 1), we combine standard facts from the representation theory of finite abelian groups with a straightforward anticoncentration argument. We begin with stating the output distribution of Fourier sampling

$$q_\psi(\lambda) = \frac{1}{|G|}\sum_{g\in G}\overline{\chi_\lambda(g)}\langle\psi|R(g)|\psi\rangle. \tag{4}$$

By character orthogonality, q_ψ is supported only on $H^\perp$, the dual of the hidden subgroup H. Crucially, we show that q_ψ anticoncentrates over all proper subgroups of $H^\perp$. Specifically, for any proper subgroup $K^\perp < H^\perp$, we have

$$q_\psi\left(K^\perp\right) \le 1 - \Omega(\epsilon).$$

This anticoncentration ensures that after $O(\log|G|/\epsilon)$ samples from q_ψ, one obtains a complete generating set for $H^\perp$ with high probability, and thus efficiently recovers H.

2. Learning abelianization

In the context of learning stabilizer groups, we adopt the known strategy of first learning the stabilizer group up to phases [14,16,18]. Within our framework, this can be interpreted as learning the *abelianization* of the group. More generally, even for non-abelian groups G, one can consider their abelianization $G/[G,G]$, which is always abelian. For example, the Pauli group $\mathcal{P}_n$ abelianizes to $\mathbb{Z}_2^{2n}$, corresponding to Pauli operators modulo phases. This perspective suggests a broader principle: viewing learning tasks through the lens of abelianizations may help uncover new applications of the abelian StateHSP and simplify existing ones.

3. Known measurement routines are instances of Fourier sampling

A separate conceptual contribution of this work is to show that important quantum learning routines such as Bell sampling and Bell difference sampling are special cases of Fourier sampling for appropriate choices of group G and representation R. In these cases, we show that the character POVM in Eq. (17) associated with Fourier sampling exactly matches the measurement implemented by these routines. Notably, these choices of G and R arise naturally from the structure of the corresponding learning problems.

D. Related work

Most closely related works have already been discussed in the context of our applications. Beyond these, we note recent works on quantum algorithms for *testing* symmetries rather than learning them [28,29]. These works employ similar representation-theoretic tools. They also clarify distinct notions of symmetry for quantum states in the mixed-state setting that we pick up on in Sec. II D.

Additionally, Ref. [30] develops a general circuit framework for implementing character projection, targeting mostly more involved group actions than those considered in this work. Lastly, Ref. [31] explores the hidden subgroup problem and Fourier sampling approach in the context of designing heuristic quantum machine learning algorithms.

II. ABELIAN STATE HIDDEN SUBGROUP PROBLEM

In this section, we present the Fourier sampling-based approach to the abelian StateHSP. This approach is applicable to identifying hidden symmetries associated to subgroups of any finite abelian group G. In Sec. II A, we give the necessary background on the representation theory of such groups. In Sec. II B, we outline the algorithm, discuss its key properties and prove our first main result, the upper bound of $O(\log|G|/\epsilon)$ on its sample complexity. In Sec. II C, we discuss how to implement the algorithm as efficiently as possible. Lastly, in Sec. II D, we comment on extending the StateHSP framework to mixed-state inputs.

A. Basics of representation theory of abelian groups

Let G be a finite abelian group, then a *character* of G is a group homomorphism $\chi : G \to \mathbb{S}^1$ into the circle group $\mathbb{S}^1$. The set of all characters of G, denoted $\hat{G}$, forms the dual group of G and is itself abelian. We have that $G \cong \hat{G}$ and, therefore, the same set of labels can be used for the group and its dual. Any subgroup $H \leq G$ defines a dual subgroup $H^\perp \leq \hat{G}$, referred to as the *annihilator* of H, given by

$$H^\perp = \{\chi_\lambda \in \hat{G} : \chi_\lambda(h) = 1, \forall\, h \in H\}, \tag{5}$$

with $|H| \cdot |H^\perp| = |G|$. Equivalently, H is uniquely determined by $H^\perp$ via

$$H = \{h \in G : \chi_\lambda(h) = 1, \forall\, \chi_\lambda \in H^\perp\} \tag{6}$$

and, therefore, if $\{\chi_{\lambda_1}, \ldots, \chi_{\lambda_k}\}$ is a generating set for $H^\perp$, then

$$H = \bigcap_{i=1}^{k} \ker(\chi_{\lambda_i}), \tag{7}$$

$$\text{where} \quad \ker(\chi_\lambda) = \{g \in G : \chi_\lambda(g) = 1\}. \tag{8}$$

Further, the characters satisfy orthogonality relations. Here, we only state a special case of these relations that is directly relevant to this work.

Fact 1 (Character sum). Let $H \leq G$ be a subgroup of G and $H^\perp \leq \hat{G}$ the corresponding annihilator subgroup in the dual $\hat{G}$, then

$$\sum_{\lambda \in H^\perp} \chi_\lambda(g) = \begin{cases} |H^\perp|, & g \in H, \\ 0, & g \notin H. \end{cases} \tag{9}$$

The characters are also central to the representation theory of finite abelian groups. In particular, since every *irreducible representation* of a finite abelian group is one-dimensional, each character corresponds to such a representation. Formally, let $R : G \to \mathrm{U}(\mathcal{H})$ be a unitary representation of the finite group G on the Hilbert space $\mathcal{H}$. Then $\mathcal{H}$ decomposes into irreducible representations of G as

$$\mathcal{H} \cong \bigoplus_\lambda V_\lambda^{\oplus n_\lambda} \cong \bigoplus_\lambda V_\lambda \otimes \mathbb{C}^{n_\lambda}, \tag{10}$$

where n_λ is the multiplicity of the irreducible representation V_λ and $V_\lambda \otimes \mathbb{C}^{n_\lambda}$ is called the λ-*isotypic component* of $\mathcal{H}$. In the abelian case, each V_λ is one-dimensional and corresponds to the character χ_λ. An orthonormal basis for the Hilbert space can be labeled as

$$|\lambda, \nu_\lambda\rangle := |\lambda\rangle \otimes |\nu_\lambda\rangle, \tag{11}$$

where λ runs over irreducible representations and $\nu_\lambda \in \{1, 2, \ldots, n_\lambda\}$. This basis is also a common eigenbasis for the mutually commuting operators $\{R(g)\}_{g \in G}$ with eigenvalues given by the characters,

$$R(g)|\lambda, \nu_\lambda\rangle = \chi_\lambda(g)|\lambda, \nu_\lambda\rangle. \tag{12}$$

We focus on projectors onto λ-isotypic components, which in the abelian case are given by

$$\Pi_\lambda = \sum_{\nu_\lambda=1}^{n_\lambda} |\lambda, \nu_\lambda\rangle\langle\lambda, \nu_\lambda|, \tag{13}$$

where $|\lambda, \nu_\lambda\rangle$ is the orthonormal basis described in Eq. (11). In general, these projectors can also be expressed as follows.

Fact 2 (Character projection formula). Let G be a finite group, R a representation on $\mathcal{H}$ and let χ_λ denote the irreducible characters. Then $\mathcal{H}$ decomposes as in Eq. (10) and

$$\Pi_\lambda = \frac{\dim V_\lambda}{|G|} \sum_{g \in G} \overline{\chi_\lambda(g)} R(g). \tag{14}$$

We finish this section with a general statement about probability distributions p defined on finite abelian groups G. Given a set of group elements $\{g_1, \ldots, g_m\}$, we denote the subgroup generated by them by $\langle g_1, \ldots, g_m\rangle$. Further, we denote the probability mass on a subgroup $K \leq G$ by $p(K) = \sum_{g \in K} p(g)$.

Lemma 1 (Probability mass on generated subgroup). Let $\epsilon, \delta \in (0, 1)$ and let p be a distribution over a finite abelian group G. Suppose $g_1, \ldots, g_m$ are i.i.d. samples from p with

$$m \geq \frac{2(\log|G| + \log(1/\delta))}{\epsilon}, \tag{15}$$

then with probability at least $1 - \delta$,

$$p(\langle g_1, \ldots, g_m\rangle) \geq 1 - \epsilon. \tag{16}$$

Similar statements have been used in Refs. [17,18] for distributions over $\mathbb{F}_2^n$ in the context of stabilizer learning. Intuitively, with every new independent sample, the chance of making progress in growing the generated subgroup is at least ϵ but the number of generators is bounded by $\log|G|$. The result then follows from a Chernoff bound.

B. Solving the abelian StateHSP via Fourier sampling

The state hidden subgroup problem as introduced in Ref. [11] is the following problem:

Definition 5 (State hidden subgroup problem (StateHSP) [11]). Let G be a finite group with a unitary representation $R : G \to \mathrm{U}(\mathcal{H})$ acting on the Hilbert space $\mathcal{H}$ and let $H \leq G$ be a subgroup. Assume that you have access to copies of an unknown quantum state vector $|\psi\rangle \in \mathcal{H}$ with the promise that

(1) $R(h)|\psi\rangle = |\psi\rangle, \forall h \in H$, and
(2) $|\langle\psi|R(g)|\psi\rangle| < 1 - \epsilon$, whenever $g \notin H$, where $\epsilon > 0$.

The problem is to identify H.

Below we will describe our Fourier sampling-based approach to solving the StateHSP, in the case where G is abelian. Fourier sampling corresponds to measuring $|\psi\rangle$ repeatedly using the *character POVM*

$$\Pi_\lambda = \frac{1}{|G|} \sum_{g\in G} \overline{\chi_\lambda(g)} R(g), \tag{17}$$

which is precisely the projector onto the λ-isotypic component of the representation R by Fact 2.

Throughout, we will denote the output distribution from measuring the POVM Π_λ as

$$q_\psi(\lambda) = \mathrm{tr}(\Pi_\lambda |\psi\rangle\langle\psi|) = \frac{1}{|G|} \sum_{g\in G} \overline{\chi_\lambda(g)} \langle\psi|R(g)|\psi\rangle. \tag{18}$$

As a consequence of Fact 1, we find that for any subgroup $K \leq G$ the probability mass of q_ψ on the annihilator $K^\perp$ satisfies the following relation.

Lemma 2 (q_ψ-mass on subgroups). Let $K \leq G$ be a subgroup of G and let $K^\perp$ be its corresponding annihilator. Then,

$$q_\psi(K^\perp) = \mathrm{tr}(\Pi_K |\psi\rangle\langle\psi|) = \frac{1}{|K|} \sum_{k\in K} \langle\psi|R(k)|\psi\rangle, \tag{19}$$

where $\Pi_K = \frac{1}{|K|} \sum_{k\in K} R(k)$.

Note that $\Pi_K = (1/|K|) \sum_{k\in K} R(k)$ is the projector onto the subspace that is invariant under the K-action. So, the q_ψ-mass on some $K^\perp$ measures precisely how close $|\psi\rangle$ is to being invariant under K.

Proof. Using Fact 1, we find that

$$\begin{aligned}
\sum_{\lambda\in K^\perp} \Pi_\lambda &= \frac{1}{|G|} \sum_{g\in G} \left(\sum_{\lambda\in K^\perp} \overline{\chi_\lambda(g)} \right) R(g), \\
&= \frac{|K^\perp|}{|G|} \sum_{k\in K} R(k), \\
&= \frac{1}{|K|} \sum_{k\in K} R(k), \\
&= \Pi_K,
\end{aligned} \tag{20}$$

where we have used $|K^\perp||K| = |G|$. By linearity of the trace, we find

$$\sum_{\lambda\in K^\perp} q_\psi(\lambda) = \sum_{\lambda\in K^\perp} \mathrm{tr}(\Pi_\lambda |\psi\rangle\langle\psi|) = \mathrm{tr}(\Pi_K |\psi\rangle\langle\psi|). \tag{21}$$

■

Hence, if $|\psi\rangle$ is invariant under the R-action of the hidden subgroup $H \leq G$ as assumed in Definition 5, then q_ψ is supported only on $H^\perp$, meaning

$$q_\psi(H^\perp) = 1. \tag{22}$$

From an algorithmic viewpoint, this suggests that repeatedly sampling from q_ψ gradually reveals $H^\perp$. In particular, a sufficiently large number of i.i.d. samples is likely to form a generating set for $H^\perp$. But exactly how many samples are sufficient to ensure this?

The key factor here is the parameter ϵ from Definition 5. Generally, the larger ϵ, the more uniformly q_ψ is distributed over $H^\perp$, increasing the likelihood that new samples remain independent of those already obtained. In what follows, we present a straightforward analysis that quantifies this relationship, showing that $O(\log(|H^\perp|/\epsilon))$ samples suffice to form a generating set for $H^\perp$. This analysis is based on demonstrating that q_ψ is anticoncentrated on the proper subgroups of $H^\perp$.

Lemma 3 (Anticoncentration of q_ψ). Let G, H, R and $|\psi\rangle$ be defined as in Definition 5. In particular, assume that $|\psi\rangle$ is such that

(1) $R(h)|\psi\rangle = |\psi\rangle, \forall h \in H$, and
(2) $|\langle\psi|R(g)|\psi\rangle| < 1 - \epsilon$, whenever $g \notin H$.

Let q_ψ be defined as in Eq. (18). Then, for all proper subgroups $K^\perp < H^\perp$ of $H^\perp$, it holds that

$$q_\psi(K^\perp) \leq 1 - \frac{\epsilon}{2}. \tag{23}$$

Proof. First, from Lemma 2, we have

$$q_\psi(K^\perp) = \frac{1}{|K|}\sum_{k\in K}\langle\psi|R(k)|\psi\rangle. \tag{24}$$

Note that by assumption $H < K$, in fact, K is strictly larger than H. Next, we split the sum over K and use the assumptions on $|\psi\rangle$,

$$\begin{aligned} q_\psi(K^\perp) &= \frac{1}{|K|}\left(\sum_{k\in H} 1 + \sum_{k\in K\setminus H}\langle\psi|R(k)|\psi\rangle\right) \\ &\leq \frac{1}{|K|}\left(|H| + (1-\epsilon)\left(|K| - |H|\right)\right), \\ &\leq 1 - \epsilon/2, \end{aligned} \tag{25}$$

where in the last step, we have used that $|H| \leq |K|/2$ by Lagrange's theorem. ■

Combining this anticoncentration result with Lemma 1, we arrive at the following sample complexity upper bound for the abelian StateHSP.

Theorem 5 (Sample complexity upper bound). There exists an algorithm for solving the abelian StateHSP with sample complexity $O(\log|G|/\epsilon)$.

Proof. Let $\epsilon' > \epsilon/2$ and assume that we have sampled

$$m \geq \frac{2\log|G| + 2\log(1/\delta)}{\epsilon'} \tag{26}$$

times from q_ψ to obtain i.i.d. samples $\chi_{\lambda_1}, \ldots, \chi_{\lambda_m}$. Then, it follows from Lemma 1 that with probability $1-\delta$, the mass on the group generated by the samples satisfies $q_\psi\left(\langle\chi_{\lambda_1}, \ldots, \chi_{\lambda_m}\rangle\right) > 1 - \epsilon/2$. Now, using Lemma 3, we conclude that, in this case $\langle\chi_{\lambda_1}, \ldots, \chi_{\lambda_m}\rangle = H^\perp$ because all proper subgroups of $H^\perp$ only account for at most $1-\epsilon/2$ of the total mass of q_ψ. ■

The StateHSP as formulated in Definition 5 asks us to find a hidden symmetry in a quantum state vector $|\psi\rangle$ promised to have one. However, sometimes it might be more natural to adopt a slightly different perspective: rather than assuming the existence of a hidden symmetry, we simply ask whether the state exhibits any nontrivial symmetry at all and if so to find it. In this vein, we point out two basic properties about the output of the algorithm:

Lemma 4 (Properties of H output by the algorithm). Let G be a finite group with a unitary representation $R : G \to \mathrm{U}(\mathcal{H})$ on $\mathcal{H}$. Assume that you have access to copies of an unknown quantum state vector $|\psi\rangle$. Let $\{\chi_{\lambda_1}, \ldots, \chi_{\lambda_m}\}$ be m independent samples from q_ψ and let $H = \cap_i^m \ker(\chi_{\lambda_i})$ be the output of the algorithm. Then,

(1) For all $g \in G$ such that $R(g)|\psi\rangle = |\psi\rangle$, $g \in H$.
(2) If $m \geq 2(\log|G| + \log(1/\delta))/\epsilon$, then with probability at least $1-\delta$ all elements $h \in H$ satisfy, $|\langle\psi|R(h)|\psi\rangle| \geq 1-\epsilon$.

The first property ensures that the algorithm's output always includes the exact symmetry group, while the second implies that any elements g not corresponding to exact symmetries are rapidly suppressed as the number of samples increases.

Proof. Let g be an arbitrary element of G and let $K := \langle g\rangle$ be the group generated by g. Then, the dual of K is given by

$$K^\perp = \{\lambda : \chi_\lambda(g') = 1, \forall\, g' \in \langle g\rangle\} = \{\lambda : \chi_\lambda(g) = 1\}. \tag{27}$$

The probability that g belongs to the subgroup H output by our algorithm is given by

$$\begin{aligned} \Pr(g \in H) &= \Pr(g \in \cap_i^m \ker(\chi_{\lambda_i})) = \prod_{i=1}^m \Pr(g \in \ker(\chi_{\lambda_i})) \\ &= \left(q_\psi(K^\perp)\right)^m, \end{aligned} \tag{28}$$

where we have used that the λ_i have been sampled independently. We will now find a bound on this probability in terms of $|\langle\psi|R(g)|\psi\rangle|$. By the orthogonality relation of the characters, we can invert Eq. (18) to obtain

$$\begin{aligned} \langle\psi|R(g)|\psi\rangle &= \sum_\lambda \chi_\lambda(g) q_\psi(\lambda) = q_\psi(K^\perp) \\ &\quad + \sum_{\lambda\notin K^\perp}\chi_\lambda(g) q_\psi(\lambda), \end{aligned} \tag{29}$$

where we have used the definition of $K^\perp$ and written $q_\psi(K^\perp) = \sum_{\lambda\in K^\perp} q_\psi(\lambda)$. From Eq. (29), the first claimed property follows: By taking the real part of this equation, we observe that $\langle\psi|R(g)|\psi\rangle = 1$ implies that $q(K^\perp) = 1$ and so $\Pr(g \in H) = 1$ by Eq. (28).

To show the second property, we can lower bound the absolute value of the LHS of Eq. (29) as

$$\begin{aligned} |\langle\psi|R(g)|\psi\rangle| &\geq q_\psi(K^\perp) \\ &\quad - \left|\sum_{\lambda\notin K^\perp}\chi_\lambda(g) q_\psi(\lambda)\right| \geq 2q_\psi(K^\perp) - 1, \end{aligned} \tag{30}$$

where we have used that q_ψ is normalized such that

$$\left|\sum_{\lambda\notin K^\perp}\chi_\lambda(g)q_\psi(\lambda)\right| \le \sum_{\lambda\notin K^\perp} q_\psi(\lambda) = 1 - q_\psi(K^\perp). \quad (31)$$

Thus, we have shown

$$q_\psi(K^\perp) \le \frac{1+|\langle\psi|R(g)|\psi\rangle|}{2}, \quad (32)$$

and hence

$$\Pr(g\in H) \le \left(\frac{1+|\langle\psi|R(g)|\psi\rangle|}{2}\right)^m. \quad (33)$$

Finally, consider the event that H contains at least a single g such that $|\langle\psi|R(g)|\psi\rangle| < 1-\epsilon$,

$$A = \bigcup_{\{g\in G:|\langle\psi|R(g)|\psi\rangle|<1-\epsilon\}} \{g\in H\}. \quad (34)$$

Now, using the union bound, the probability of this event is bounded as

$$\Pr(A) \le |G|(1-\epsilon/2)^m \le \exp(\ln|G| - \epsilon m/2) \le \delta, \quad (35)$$

whenever the sample complexity satisfies $m \ge (2\log|G| + 2\log(1/\delta))/\epsilon$. ■

C. Implementing Fourier sampling

While we have proved that Fourier sampling provides an efficient algorithm for solving the abelian StateHSP, we have not yet detailed how to implement it in practice. The canonical technique for this is known as *generalized phase estimation*. Here, for the sake of concreteness, we show that, in the abelian case, generalized phase estimation implements the character POVM in Eq. (17).

Let G be a finite abelian group, then the *quantum Fourier transform* (QFT) is defined as the following unitary [6]:

$$\mathrm{QFT}|g\rangle = \frac{1}{\sqrt{|G|}}\sum_\lambda \chi_\lambda(g)|\lambda\rangle. \quad (36)$$

To implement the POVM from Eq. (17) via a quantum Fourier sampling approach, we proceed as follows:

(1) Start with the state vector $|0\rangle\otimes|\psi\rangle$.
(2) Put the auxiliary register in superposition to obtain $(1/\sqrt{|G|})\sum_{g\in G}|g\rangle\otimes|\psi\rangle$.
(3) Apply the controlled group action $\sum_g |g\rangle\langle g|\otimes R(g)$ to obtain $(1/\sqrt{|G|})\sum_{g\in G}|g\rangle\otimes R(g)|\psi\rangle$.
(4) Apply the inverse QFT to the auxiliary register to obtain $(1/|G|)\sum_\lambda\sum_{g\in G}\overline{\chi_\lambda(g)}|\lambda\rangle\otimes R(g)|\psi\rangle$.
(5) Measure the auxiliary register in the $|\lambda\rangle$-basis.

This results in the following output distribution:

$$q_\psi(\lambda) = \frac{1}{|G|}\sum_{g\in G}\overline{\chi_\lambda(g)}\langle\psi|R(g)|\psi\rangle. \quad (37)$$

This is precisely the distribution $q_\psi(\lambda) = \mathrm{tr}(\Pi_\lambda|\psi\rangle\langle\psi|)$ from Eq. (18).

While the generalized phase estimation approach is broadly applicable to all finite abelian groups, it poses practical challenges. In particular, it requires auxiliary qubits and controlled operations involving many qubits, which can be difficult to implement on near-term quantum hardware, especially when locality constraints limit the connectivity.

An alternative strategy, conceptually simpler but potentially less general, is to directly measure in a joint eigenbasis of the representation $\{R(g)\}_{g\in G}$. Since the $R(g)$ commute for abelian G, there exists an orthonormal basis $|\lambda,\nu_\lambda\rangle$ of joint eigenvectors [see Eq. (11)] satisfying

$$R(g)|\lambda,\nu_\lambda\rangle = \chi_\lambda(g)|\lambda,\nu_\lambda\rangle. \quad (38)$$

Measuring in this basis and discarding the multiplicity label ν_λ yields a sample from the same distribution $q_\psi(\lambda)$ as in generalized phase estimation.

Whether this approach is feasible depends on both the structure of the group G and the explicit form of the representation R. In general, it is unclear how to construct and implement a measurement in the joint eigenbasis. Nevertheless, in the specific applications of the StateHSP framework explored in this paper, we encounter cases where this direct measurement strategy is natural and easily realizable, offering a more efficient route to Fourier sampling in practice. Interestingly, this perspective also allows us to connect Fourier sampling to well-established measurement routines such as Bell sampling and Bell difference sampling in the quantum learning literature. In particular, in Sec. IV, we show that the hidden cut problem defined in Ref. [11] can be solved via Bell sampling which precisely corresponds to such a direct implementation of Fourier sampling.

D. Mixed input states

An extension of the StateHSP is to allow *mixed* input states. Here, we provide a natural definition for this setting:

Definition 6 (Mixed-state hidden subgroup problem (MixedStateHSP)). Let G be a finite group with a unitary representation $R: G\to \mathrm{U}(\mathcal{H})$ acting on the Hilbert space $\mathcal{H}$ and let $H\le G$ be a subgroup. Assume that you have access to copies of an unknown quantum state $\rho\in\mathcal{L}(\mathcal{H})$ with the promise that

(1) $R(h)\rho R(h)^\dagger = \rho,\ \forall h \in H$, and
(2) $\|R(g)\rho R(g)^\dagger - \rho\| \geq \epsilon$, whenever $g \notin H$, for $\epsilon > 0$ and some suitable norm $\|\cdot\|$ such as the trace norm.

The problem is to identify H.

This formulation reduces to the StateHSP as defined in Definition 5 in case ρ is pure. To tackle an abelian MixedStateHSP with a mixed input ρ, a first thought is to apply the same character-measurement and postprocessing strategy as in the pure state case, setting

$$q_\rho(\lambda) = \mathrm{tr}(\Pi_\lambda \rho). \tag{39}$$

However, we note that, in general this strategy fails, as the invariance condition $R(h)\rho R(h)^\dagger = \rho$ in Definition 6 does not guarantee that q_ρ fully concentrates all its mass on the true annihilator $H^\perp$ such that, in general, $q_\rho(H^\perp) < 1$. This breaks the algorithm, as we have no efficient way of distinguishing samples belonging to $H^\perp$ from those that do not.

On the other hand, our algorithm *does* succeed when we impose a stronger requirement than that stated in Definition 6. Namely, if ρ satisfies

(1) $\mathrm{tr}(R(h)\rho) = 1, \forall\, h \in H$, and
(2) $|\mathrm{tr}(R(g)\rho)| \leq 1 - \epsilon$, whenever $g \notin H$, for some $\epsilon > 0$,

then the character POVM-based approach will succeed in learning H, even with mixed input states. This condition is strictly stronger than the invariance requirement in Definition 6, as it implies—but is not implied by—that condition. Specifically,

$$\mathrm{tr}(R(h)\rho) = 1 \Rightarrow R(h)\rho R(h)^\dagger = \rho. \tag{40}$$

We note that these inequivalent notions of symmetries of mixed states have been discussed in detail in Ref. [28] in the context of testing symmetries rather than learning them. Designing an efficient algorithm that works under the weaker, more natural invariance condition of Definition 6 remains an open problem. Such a result would broaden the applicability of the StateHSP framework to abelian symmetry learning for general quantum channels. In particular, via the *Choi Jamiołkowski isomorphism*, these problems reduce naturally to instances of the StateHSP. However, the limitations with mixed input states discussed here restrict our current method to unitary channels.

III. LEARNING STABILIZER GROUPS

In this section, we explain how the problem of learning stabilizer groups can be viewed as an instance of the abelian StateHSP. We start by giving some essential background on Weyl operators in Sec. III A. Then, in Sec. III B, we show how to solve the problem via Fourier sampling in Sec. II B. In Sec. III C, we show that for this problem, Fourier sampling can always be implemented by a suitable Clifford circuit. Lastly, in Sec. III D, we argue that the algorithm can also be applied to learning a stabilizer code when given access to different code states rather than many copies of the same one.

A. Weyl operators and stabilizer groups

Let $d \geq 2$ be a prime. Consider a single qudit with computational basis elements $|q\rangle$ where $q \in \mathbb{Z}_d$. Define the unitary shift and clock operators X and Z, respectively, as

$$\begin{aligned} X|q\rangle &= |q+1\rangle, \\ Z|q\rangle &= \omega^q |q\rangle, \end{aligned} \tag{41}$$

for all $q \in \mathbb{Z}_d$, where $\omega = \mathrm{e}^{2\pi i/d}$ is the d-th root of unity. There are some slight differences in the algebra of the shift and clock operators between the qubit case (d even) and the qudit case (d odd). To treat both cases simultaneously, we also introduce $\tau = \mathrm{e}^{i\pi(d^2+1)/d}$. We have that $\tau^2 = \omega$ and

$$\tau^d = \begin{cases} 1 & d \text{ odd}, \\ -1 & d \text{ even}. \end{cases} \tag{42}$$

Next, we introduce the n-qudit Weyl operators as

$$W_x = W_{(a,b)} = \tau^{-a\cdot b}\left(Z^{a_1}X^{b_1}\right) \otimes \cdots \otimes \left(Z^{a_n}X^{b_n}\right), \tag{43}$$

for all $x = (a,b) \in \mathbb{Z}_d^{2n}$. Each Weyl operator is also an element of the generalized Pauli group.

Consider the $\mathbb{Z}_d$-valued symplectic form defined in $\mathbb{Z}_d^{2n}$ for $x = (a,b)$ and $y = (a',b')$ as

$$[x,y] = \left[(a,b),(a',b')\right] = a\cdot b' - a'\cdot b \mod d. \tag{44}$$

Then, the Weyl operators compose as

$$W_x W_y = \tau^{[x,y]} W_{x+y}, \tag{45}$$

and their commutation relations are captured by the following equation:

$$W_x W_y = \omega^{[x,y]} W_y W_x. \tag{46}$$

For more background on Weyl operators and the symplectic formalism, see Ref. [19].

Tensor products of the shift and clock operators generate the generalized n-qudit Pauli group,

$$\mathcal{P}_n := \langle \tau I, X, Z\rangle^{\otimes n}. \tag{47}$$

A stabilizer group S is an abelian subgroup of the generalized Pauli group $\mathcal{P}_n$ that does not contain a nontrivial

multiple of the identity operator $I^{\otimes n}$. We say that a state vector $|\psi\rangle$ has a nontrivial stabilizer group if there exists $S \neq \{I^{\otimes n}\}$, such that

$$P|\psi\rangle = |\psi\rangle, \quad \forall\, P \in S. \tag{48}$$

For our purposes, it will be useful to consider the Pauli operators only up to phase. To this end, we denote by $\mathrm{P}_n := \mathcal{P}_n/\langle \omega I\rangle$ the phaseless or *projective* Pauli group which is isomorphic to $\mathbb{F}_d^{2n}$. For any state vector $|\psi\rangle$, we also define its corresponding phaseless stabilizer group as follows, adopting the notation introduced in Ref. [18]:

Definition 7 (Phaseless stabilizer group). Let $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ be a state vector on n qudits. We define $\mathrm{Weyl}\,(|\psi\rangle)$ as the following subspace of $\mathbb{F}_d^{2n}$,

$$\mathrm{Weyl}\,(|\psi\rangle) = \left\{x \in \mathbb{F}_d^{2n} : W_x|\psi\rangle = \omega^s|\psi\rangle \text{ for some } s \in \mathbb{F}_d\right\}. \tag{49}$$

We note that $\mathrm{Weyl}\,(|\psi\rangle) \subset \mathbb{F}_d^{2n}$ is isomorphic to $\mathbb{F}_d^k$ where $k = \dim \mathrm{Weyl}\,(|\psi\rangle) \leq n$.

B. Learning Weyl ($|\psi\rangle$) as an abelian StateHSP

We start by formulating the problem of learning stabilizer groups up to phases.

Definition 8 (Hidden phaseless stabilizer group problem). Let $|\psi\rangle \in \mathbb{C}^{d^n}$ be a state vector on n qudits and suppose that it has a nontrivial phaseless stabilizer group $\mathrm{Weyl}(|\psi\rangle)$ while for all $x \notin \mathrm{Weyl}(|\psi\rangle)$, $|\langle\psi|W_x|\psi\rangle| \leq 1 - \epsilon$. The hidden phaseless stabilizer group problem asks to identify $\mathrm{Weyl}(|\psi\rangle)$ given access to copies of $|\psi\rangle$.

We can identify the phase-full hidden stabilizer group S mentioned in Definition 2 from its corresponding phaseless version $\mathrm{Weyl}\,(|\psi\rangle)$ by determining the missing phases of the stabilizers from measuring a set of generators of $\mathrm{Weyl}\,(|\psi\rangle)$ on $|\psi\rangle$.

The phaseless version of the hidden stabilizer group problem in Definition 8 fits into the StateHSP framework by choosing $G = \mathbb{Z}_d^{2n}$. To treat the qubit and the qudit case simultaneously, let D denote the order of τ, i.e.,

$$D = \begin{cases} d & d \text{ odd}, \\ 2d & d \text{ even}, \end{cases} \tag{50}$$

and choose the representation

$$R : \mathbb{Z}_d^{2n} \to \mathrm{U}\left((\mathbb{C}^{d^n})^{\otimes D}\right), \tag{51}$$

$$x \mapsto R(x) = W_x^{\otimes D}. \tag{52}$$

It follows from Eq. (45) that this is a valid representation, however, we emphasize that R is a representation on $(\mathbb{C}^{d^n})^{\otimes D}$, i.e., on D copies of the n-qudit Hilbert space $\mathbb{C}^{d^n}$. The hidden subgroup is given by

$$H = \mathrm{Weyl}\,(|\psi\rangle)\,, \tag{53}$$

and we have for $|\psi\rangle^{\otimes D}$ that

(1) $W_x^{\otimes D}|\psi\rangle^{\otimes D} = |\psi\rangle^{\otimes D}$, for all $x \in \mathrm{Weyl}(|\psi\rangle)$,
(2) $|\langle\psi|W_x|\psi\rangle|^D \leq (1-\epsilon)^D$, for all $x \notin \mathrm{Weyl}(|\psi\rangle)$.

Since $\mathbb{Z}_d^{2n}$ is abelian, we can tackle this problem via the Fourier sampling approach to the abelian StateHSP outlined in Sec. II B. In particular, we will implement the character POVM corresponding to Eq. (17). The characters are labeled by $y = (a, b) \in \mathbb{Z}_d^{2n}$ and given by

$$\chi_y(x) = \omega^{[x,y]}, \tag{54}$$

for all $x \in \mathbb{Z}_d^{2n}$, and the POVM is given by

$$\Pi_y = \frac{1}{d^{2n}} \sum_{x \in \mathbb{Z}_d^{2n}} \omega^{-[y,x]} W_x^{\otimes D}. \tag{55}$$

Hence, the output distribution $q_\psi = \mathrm{tr}\left(\Pi_y (|\psi\rangle\langle\psi|)^{\otimes D}\right)$ is given by

$$q_\psi(y) = \frac{1}{d^{2n}} \sum_{x \in \mathbb{Z}_d^{2n}} \omega^{-[y,x]} \langle\psi|W_x|\psi\rangle^D. \tag{56}$$

In the qubit case, where $d = 2$ and $D = 4$, Eq. (55) is the Bell difference sampling POVM [18,19]. Hence, in the qubit case, Fourier sampling corresponds exactly to known algorithms for learning stabilizer groups. However, the approach also naturally generalizes to qudits, where such algorithms were not previously known.

The projective measurement Π_y from Eq. (55) can be implemented by measuring in the common eigenbasis of the $R(x) = W_x^{\otimes D}$ for all $x \in \mathbb{Z}_d^{2n}$. We comment on this eigenbasis in the subsection below. Here, by virtue of our general result, Theorem 5, we arrive at the following theorem.

Theorem 6 (Unified hidden stabilizer group algorithm). There is an efficient nonadaptive quantum algorithm for the hidden stabilizer group problem on n qudits which uses $O\,(n \log d \max\{d, 1/\epsilon\})$ copies of the unknown state and runs in polynomial time, requiring circuits of depths $O(d)$ acting coherently on at most d copies at a time using no additional auxiliary systems.

Lastly, we observe that computational difference sampling, another primitive from the stabilizer learning literature [18,22], also can be viewed as Fourier sampling. It

involves measuring two copies of an n-qubit state vector $|\psi\rangle$ in the computational basis and combining the outcomes via bitwise XOR. This approach enables learning Z-like stabilizers of $|\psi\rangle$ (up to phase) using only single-copy measurements. This measurement routine corresponds to Fourier sampling with respect to $G = \mathbb{Z}_2^n$ and the representation $R : a \mapsto W_{(a,0^n)}^{\otimes 2} = Z^a \otimes Z^a$ on two copies of the Hilbert space. Then, with the characters labeled by $b \in \mathbb{Z}_2^n$ and given by $\chi_b(a) = (-1)^{a\cdot b}$, the character POVM from Eq. (17) becomes

$$\Pi_b = \frac{1}{2^n}\sum_a (-1)^{a\cdot b} Z^a \otimes Z^a. \tag{57}$$

This is precisely the computational difference sampling POVM given in Ref. [18]. This correspondence readily also generalizes to qudits.

C. Common eigenbasis of the Weyl operators

To implement the POVM in Eq. (55), we can measure in the joint eigenbasis of the mutually commuting $W_x^{\otimes D}$ for all $x \in \mathbb{Z}_d^{2n}$. But what does this joint eigenbasis look like and how exactly do we implement a measurement in it?

In the qubit case, the eigenbasis of all the $W_x^{\otimes 2}$ is the well-known multiqubit Bell basis $\{|W_x\rangle\}_{x\in\mathbb{Z}_2^{2n}}$, with

$$|W_x\rangle = (W_x \otimes I)|\Omega\rangle, \quad |\Omega\rangle = \frac{1}{\sqrt{2^n}}\sum_{x\in\mathbb{Z}_2^n} |x\rangle|x\rangle. \tag{58}$$

Consequently, since $D = 2d = 4$, the eigenbasis of $W_x^{\otimes 4}$ consists of tensor products of Bell states $|W_x\rangle \otimes |W_y\rangle$ and the POVM in Eq. (55) corresponds to Bell difference sampling.

In the qudit case, the common eigenbasis of the $W_x^{\otimes D}$ is not given by the generalized Bell states $|W_x\rangle = (W_x \otimes I)|\Omega\rangle$ where $|\Omega\rangle$ is the qudit maximally entangled state. In fact, these states live on two copies, however we are looking for a basis of the d-copy Hilbert space $(\mathbb{C}^{d^n})^{\otimes d}$ since for odd d, we have $D = d$. This basis can be constructed as follows.

Since the $W_x^{\otimes d}$ mutually commute for all $x \in \mathbb{Z}_d^{2n}$, they form a stabilizer group S_d, namely

$$S_d = \langle X^{\otimes d}, \quad Z^{\otimes d}\rangle^{\otimes n}. \tag{59}$$

In general, to any stabilizer group, one can associate a joint eigenbasis of stabilizer states, a so-called *stabilizer basis* and a measurement in such a stabilizer basis can be implemented by application of a suitable Clifford circuit followed by measurement in the computational basis. Here, by the tensor product structure of S_d, the problem reduces to constructing a Clifford circuit on d qudits that maps $X^{\otimes d} \mapsto Z_1, Z^{\otimes d} \mapsto Z_2$. In tableau form, this Clifford can be found via standard Gaussian elimination over F_d, and consequently be compiled into a Clifford circuit [32].

More generally, the number of independent generators of S_d is $2n$, hence for $d > 2$, S_d is not a maximal stabilizer group, as the maximum possible number of independent generators is $dn > 2n$ on dn qudits. Consequently, the stabilizer basis associated to S_d is not uniquely determined. While the exact choice is not important here—any choice does the job—we emphasize again that the basis can be chosen to tensorize along the n qudits. In other words, since $((\mathbb{C}^d)^{\otimes n})^{\otimes d} \cong ((\mathbb{C}^d)^{\otimes d})^{\otimes n}$, the stabilizer basis can be chosen to only act simultaneously on d copies of the i-th qudit. Hence, assuming all-to-all connectivity, the measurement can be implemented by a Clifford circuit of depth at most $O(d)$.

D. Learning stabilizer codes from access to code states

In quantum learning theory, the standard access model typically involves learning properties of a single quantum state. However, for the symmetry identification tasks explored in this work, it is both natural and beneficial to consider access to a collection of unknown quantum states that exhibit the same underlying symmetry. A canonical example of this arises with states residing in the code space of a common stabilizer code. It is entirely plausible to encounter scenarios where the goal is to identify the stabilizer code based on access to multiple different states in the code space, rather than repeated access to a single one. This situation might, for instance, occur when intercepting a stream of encoded quantum information being transmitted over a quantum communication channel.

In this vein, here we define the following variant of the hidden stabilizer group problem from Definition 2:

Definition 9 (Learning hidden stabilizer codes). Let $\{|\psi_i\rangle\}_i \in \mathbb{C}^{d^n}$ be a collection of states on n qudits all belonging to the code space of some stabilizer code. In other words, assume that there is a stabilizer group $S \subseteq \mathcal{P}_n$, such that for all i,

$$P|\psi_i\rangle = |\psi_i\rangle, \quad \forall\, P \in S, \tag{60}$$

while for all $|\psi_i\rangle$ and $P \in \mathrm{P}_n \setminus \mathrm{S}$, $|\langle\psi_i|P|\psi_i\rangle| \leq 1 - \epsilon$. The hidden stabilizer code problem asks to identify S given access to copies of $\{|\psi_i\rangle\}_i$.

We can solve this problem by essentially following the same steps as outlined in Sec. III B. In particular, we again take the hidden subgroup H to be the phaseless version of the stabilizer group S which defines the code space. We implement the same POVM from Eq. (55) but on

$\bigotimes_{i=1}^{D} |\psi_i\rangle\langle\psi_i|$ in order to sample from the distribution

$$q_{\psi_1,\ldots,\psi_D}(y) = \frac{1}{d^{2n}} \sum_{x\in\mathbb{Z}_d^{2n}} \omega^{-[y,x]} \operatorname{tr}\left(W_x^{\otimes D} \bigotimes_{i=1}^{D} |\psi_i\rangle\langle\psi_i| \right). \tag{61}$$

Crucially, using the same analysis as outlined already in Sec. II B, we find that $q_{\psi_1,\ldots,\psi_D}$ is only supported on $H^\perp$ and not concentrated on any subgroup.

E. Detecting global symmetries of symmetry-protected topological ordered states

Beyond quantum information theory, our results also connect to problems in condensed-matter physics—specifically, to detecting *symmetry-protected topological* (SPT) phases, which we capture as instances of the StateHSP. At the heart of SPT order is the principle that, under a given symmetry group H, distinct SPT states cannot be connected by any constant depth, symmetry-preserving quantum circuit without undergoing a phase transition [5]. In contrast, if symmetry is allowed to be broken, any SPT state can be smoothly deformed into a trivial product state without encountering a phase transition.

There are simple instances of stabilizer states known to exhibit SPT order. A canonical example is the *cluster* state vector $|\psi\rangle$ on n qubits on a ring, with $d = 2$ and n even. It is a qubit stabilizer state with the stabilizer group S being generated by the n stabilizers $K_j = Z_{j-1} \otimes X_j \otimes Z_{j+1}$ (with indices taken mod n) each acting on three consecutive qubits labeled by $j \in \{0,\ldots,n-1\}$ [27]. Then $|\psi\rangle$ is the unique stabilizer state vector satisfying $P|\psi\rangle = |\psi\rangle, \forall\, P \in S$. It has the global symmetries

$$P_e := \bigotimes_{j=0}^{n/2-1} (I \otimes X)^{\otimes n/2}, \tag{62}$$

$$P_o := \bigotimes_{j=0}^{n/2-1} (X \otimes I)^{\otimes n/2}, \tag{63}$$

with Pauli X operators supported on even and odd sites, in that

$$P_e|\psi\rangle = |\psi\rangle, \quad P_o|\psi\rangle = |\psi\rangle \tag{64}$$

holds true. This symmetry is referred to as a $\mathbb{Z}_2 \times \mathbb{Z}_2$-symmetry. Although $|\psi\rangle$ admits an explicit matrix product state representation—obtained by applying a layer of controlled-Z gates to a product state—any constant depth quantum circuit mapping it to the trivial product state vector $|+\rangle^{\otimes n}$ must break the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry. Consequently, the cluster state and the trivial product state, while sharing the global symmetry group generated by P_e, P_o, belong to distinct SPT phases [5,33]. Motivated by this, we define the hidden global symmetry problem, closely related to the hidden stabilizer group problem from Definition 2.

Definition 10 (Hidden global symmetry problem). Let $p \in \mathbb{N}$ be a period and let $|\psi\rangle \in \mathbb{C}^{d^n}$ be a state vector on n qudits on a ring, where $n/p \in \mathbb{N}$. Let $\mathcal{M}_p < \mathcal{P}_p$ be an abelian subgroup of the Pauli group on p sites and let $\mathrm{P}_p, \mathrm{M}_p$ denote the respective phaseless versions. Then, suppose that $|\psi\rangle$ is invariant under the stabilizer group of the form

$$\{P^{\otimes(n/p)}, P \in \mathcal{M}_p\}, \tag{65}$$

while for all $P \in \mathrm{P}_p \setminus \mathrm{M}_p$, $|\langle\psi|P^{\otimes(n/p)}|\psi\rangle| \le 1 - \epsilon$. The hidden global symmetry problem asks to identify this stabilizer group.

This is an on-site symmetry upon "blocking". In the above example of the cluster state, the hidden global symmetry subgroup is $\langle P_e, P_o\rangle$. To view this as an abelian StateHSP, we observe that to learn the global symmetry, it suffices to learn the on-site symmetry group $\mathcal{M}_p$. Then, as in Sec. III B, we focus again on first learning it up to phases, i.e., we set H to be the phaseless version of $\mathcal{M}_p$. Further, we choose as a parent group $G = \mathbb{Z}_d^{2p}$ with the action $x \in \mathbb{Z}_d^{2p} \mapsto R(x) = (W_x^{\otimes n/p})^{\otimes D}$. Finally, it follows from Theorem 5 that this can be efficiently solved with a complexity independent of n, to capture the symmetry in a quantum phase of matter.

Theorem 7 (Sample complexity lower bound for the hidden global symmetry problem). There exists an algorithm for solving the hidden global symmetry problem with sample complexity $O(p \log d/(d\epsilon))$.

IV. REVISITING THE HIDDEN CUT PROBLEM

In this section, we are concerned with the following problem first studied in Ref. [11]:

Definition 11 (Hidden many-cut problem). Let $|\psi\rangle$ be a state vector on n qubits. Suppose that $|\psi\rangle$ is a product of $m \ge 1$ factor state vectors

$$|\psi\rangle = |\phi_1\rangle_{C_1} \otimes \cdots \otimes |\phi_m\rangle_{C_m} \tag{66}$$

for some partition $C_1 \sqcup \cdots \sqcup C_m = [n]$ such that each factor state vector $|\phi_k\rangle_{C_k}$ is at least ϵ-far in trace distance from any multipartite product state on $|C_k|$ qubits. The hidden many-cut problem asks to identify the set partition $C_1 \sqcup \cdots \sqcup C_m$, given copies of the state vector $|\psi\rangle$.

As noted previously in Ref. [11], this problem fits into the StateHSP framework by choosing $G = \mathbb{Z}_2^n$ and the

representation

$$R : \mathbb{Z}_2^n \to \mathrm{U}\left(((\mathbb{C}^2)^{\otimes n})^{\otimes 2}\right), \tag{67}$$

$$x \mapsto R(x) = \bigotimes_{i=1}^{n} \mathrm{SWAP}^{x_i}. \tag{68}$$

We emphasize that R is a representation on $\left((\mathbb{C}^2)^{\otimes n}\right)^{\otimes 2}$, i.e., on two copies of the n-qubit Hilbert space. The hidden subgroup H is given by

$$H = \mathrm{span}\{1^n, 0^{C_1}1^{\overline{C_1}}, \ldots, 0^{C_{m-1}}1^{\overline{C_{m-1}}}\}, \tag{69}$$

where $\overline{C_i}$ denotes the complement of C_i in $[n]$, and $H \cong \mathbb{Z}_2^m$. It can be verified that if $|\psi\rangle$ is of the form given in Eq. (66), then $|\psi\rangle^{\otimes 2}$ is invariant under $R(x)$ for all $x \in H$ as defined in Eq. (69).

Proposition 1 (cf. Proposition 5 in Ref. [11]). Let $|\psi\rangle$ be a state vector on n qubits. Suppose that $|\psi\rangle$ is a product of $m \geq 1$ factor state vectors

$$|\psi\rangle = |\phi_1\rangle_{C_1} \otimes \cdots \otimes |\phi_m\rangle_{C_m} \tag{70}$$

such that each factor state vector $|\phi_k\rangle_{C_k}$ is ϵ-far in trace distance from any multipartite product state on $|C_k|$ qubits. Let $S \subseteq [n]$ be a subset of qubits. Then,

$$\mathrm{tr}\left(\psi_S^2\right) = \begin{cases} 1 & S = C_i \quad \text{for } i \in [m], \\ \leq 1 - \epsilon^2 & \text{else.} \end{cases} \tag{71}$$

Here, $\mathrm{tr}\left(\psi_S^2\right)$ denotes the purity of the reduced state on S.

Since $\mathbb{Z}_2^n$ is abelian, we can tackle this problem via Fourier sampling as outlined in Sec. II B. In particular, we will implement the character POVM corresponding to Eq. (17). The characters are labeled by $y \in \mathbb{Z}_2^n$ and given by

$$\chi_y(x) = (-1)^{y \cdot x}, \tag{72}$$

for all $x \in \mathbb{Z}_2^n$ and the POVM is given by

$$\Pi_y = \frac{1}{2^n} \sum_{x \in \mathbb{Z}_2^n} (-1)^{y \cdot x} \bigotimes_{i=1}^{n} \mathrm{SWAP}^{x_i}. \tag{73}$$

Hence, the resulting output distribution $q_\psi(y) = \mathrm{tr}\left(\Pi_y (|\psi\rangle\langle\psi|)^{\otimes 2}\right)$ is given by

$$q_\psi(y) = \frac{1}{2^n} \sum_{x \in \mathbb{Z}_2^n} (-1)^{y \cdot x} \mathrm{tr}(\psi_x^2). \tag{74}$$

Now, applying Theorem 5 and Proposition 1, we find that the hidden many-cut problem as stated in Definition 11 on n qubits can be solved using $O(n/\epsilon^2)$ copies of the unknown state vector $|\psi\rangle$.

The projective measurement in Eq. (73) can be implemented by measuring in the common eigenbasis of the $R(x) = \bigotimes_{i=1}^{n} \mathrm{SWAP}^{x_i}$ which is the multiqubit Bell basis. Concretely, we divide the Bell state vectors into triplet and singlet as

$$|0,0\rangle = \frac{1}{\sqrt{2}}\left(|0,0\rangle + |1,1\rangle\right), \tag{75}$$

$$|0,1\rangle = \frac{1}{\sqrt{2}}\left(|0,0\rangle - |1,1\rangle\right), \tag{76}$$

$$|0,2\rangle = \frac{1}{\sqrt{2}}\left(|0,1\rangle + |1,0\rangle\right), \tag{77}$$

$$|1,1\rangle = \frac{1}{\sqrt{2}}\left(|0,1\rangle - |1,0\rangle\right). \tag{78}$$

To summarize, we arrive at the following theorem.

Theorem 8 (Improved hidden-cut algorithm). There is an efficient nonadaptive quantum algorithm for the hidden many-cut problem on n qubits which uses $O\left(n/\epsilon^2\right)$ copies of the unknown state and runs in polynomial time, requiring circuits of constant depths acting on two copies at a time using no additional auxiliary qubits.

As mentioned before, via the *Choi Jamiołkowski isomorphism*, the many-cut problem for quantum states naturally translates into a *many-cut problem for unitaries*. This translation allows us to apply our method by preparing the Choi state of the unitary channel. In this way, given oracle access to a unitary, one can efficiently learn in what hidden way the unitary is a product.

V. TRANSLATION INVARIANCE AS STATEHSP

Here we restrict ourselves to qubits for simplicity and consider the cyclic group $C_n = \langle g \rangle$ where g is some canonical generator such that $g^n = 1$. We can define a representation of this group acting on n qubits on a ring as

$$T|x_1, x_2, \ldots, x_n\rangle = |x_n, x_1, \ldots, x_{n-1}\rangle, \tag{79}$$

where $T = T(g)$ is the canonical translation operator corresponding to g and therefore $T(g^k) = T^k$. Equivalently, we can identify the cyclic group with $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$ with addition modulo n.

Definition 12 (Hidden translation problem). Let $|\psi\rangle \in \mathbb{C}^n$ be a state vector of n qubits. We assume there is a

subgroup H of $\mathbb{Z}_n$ such that

$$T^k|\psi\rangle = |\psi\rangle, \quad \forall\, k \in H \tag{80}$$

and also that $|\langle\psi|T^k|\psi\rangle| < 1-\epsilon$ whenever $k \notin H$. The hidden translation problem asks to identify the subgroup H.

This problem fits into the StateHSP framework by choosing $G = \mathbb{Z}_n$ and the representation

$$R : \mathbb{Z}_n \to \mathrm{U}\left(\mathbb{C}^n\right), \tag{81}$$

$$k \mapsto R(k) = T^k, \tag{82}$$

where T is the operator defined in Eq. (79). Since $\mathbb{Z}_n$ is abelian, we can tackle this problem via our approach to the abelian StateHSP outlined in Sec. II B. In particular, we will implement the character POVM corresponding to Eq. (17). The irreducible characters of $\mathbb{Z}_n$ can be labeled by $j \in \mathbb{Z}_n$ and are given by

$$\chi_j(k) = \omega^{jk}, \tag{83}$$

for all $k \in \mathbb{Z}_n$, where $\omega = \mathrm{e}^{2\pi i/n}$ is a primitive n-th root of unity. Therefore, in this case the POVM corresponding to Eq. (17) reads

$$\Pi_j = \frac{1}{n}\sum_{k\in\mathbb{Z}_n} \omega^{-jk} T^k. \tag{84}$$

Hence, the resulting output distribution $q_\psi(y) = \mathrm{tr}(\Pi_j |\psi\rangle\langle\psi|)$ is given by

$$q_\psi(j) = \frac{1}{n}\sum_{k\in\mathbb{Z}_n} \omega^{-jk}\langle\psi|T^k|\psi\rangle. \tag{85}$$

Note that all subgroups of $\mathbb{Z}_n$ can be labeled by their period since all those subgroups are of the form $\{k \in \mathbb{Z}_n : k = 0 \mod r\}$ with $r|n$ (denoting that r is a divisor of n). So, the hidden subgroup H is of the form

$$H_r = \{0, r, 2r, \ldots, (n/r-1)r\}, \tag{86}$$

with $|H_r| = n/r$ for every divisor r.

In principle, the projective measurement Π_j from Eq. (84) can be implemented without any auxiliary qubits by measuring in the common eigenbasis of the T^k. However, from a practical point of view, it is not clear how to implement this eigenbasis measurement via a quantum circuit. Instead, for a concrete implementation of Π_j, one can employ the generalized phase estimation approach laid out in Sec. II C. Concretely, we will need an auxiliary register with $m = \lceil \log_2(n) \rceil$ many qubits which will serve as the control register for implementing the controlled group action $\sum_{k=0}^{m-1} |k\rangle\langle k| \otimes T^k$. Assuming all-to-all connectivity, for all $k \in \{0, \ldots, n-1\}$, T^k can be implemented via c-SWAPS in depth $O(n)$. Hence, to implement the whole controlled group action, via concatenating c-T^{2^j} layers for $j = 0, \ldots, m-1$, we require a circuit of depth $O(n \log n)$. This circuit depth dominates the cost of subsequently implementing the QFT on the auxiliary register of size $O(\log n)$ qubits. Hence, by virtue of our general result Theorem 5, we arrive at the following theorem.

Theorem 9 (Efficient algorithm for the hidden translation problem). There is an efficient nonadaptive quantum algorithm for the hidden translation problem on n qubits which uses $O((\log n)/\epsilon)$ copies of the unknown state and runs in polynomial time, requiring circuits of depths $O(n \log n)$ acting coherently on only a single copy at a time using $O(\log n)$ additional auxiliary systems.

VI. CONCLUSIONS AND OUTLOOK

In this work, we have identified a unified, efficient approach for solving the abelian state hidden subgroup problem that not only subsumes known quantum-learning subroutines (such as Bell difference sampling for qubit stabilizer learning) but also extends naturally to new settings—including qudit stabilizer groups, the hidden many-cut problem, and translational symmetries on lattice systems. This approach is conceptually streamlined and features significantly reduced circuit complexity compared to prior work, making it even more feasible for practical implementation. It also constitutes a new bridge between sampling problems with the mere use of showing a quantum advantage [34] with proper quantum algorithms with practical applications. As such, it may bring quantum algorithms featuring some utility and showing a quantum advantage closer to feasibility [13,35].

Looking ahead, one natural direction is to generalize the StateHSP framework to continuous symmetry groups, thereby unlocking further applications across the board. It would also be interesting to identify further applications in the quantum many-body context. More broadly, we hope this work invites further studies that bring together ideas of quantum algorithms with those of quantum learning theory. In fact, this intersection seems to be a particularly fruitful avenue for future research.

ACKNOWLEDGMENTS

We would like to thank Marios Ioannou for helpful discussions. This project has been supported by the German Federal Ministry for Research, Technology, and Space (MUNIQC-Atoms, QSolid, QuSol, Hybrid++), Berlin Quantum, the Munich Quantum Valley (K-8), the QuantERA (HQCC), the Clusters of Excellence (MATH+,

ML4Q), the Quantum Flagship (Millenion, PasQuans2), the DFG (CRC 183, SPP 2541), and the European Research Council (DebuQC).

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] Barbara M. Terhal, Quantum error correction for quantum memories, Rev. Mod. Phys. **87**, 307 (2015).

[2] Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O'Brien, Quantum error mitigation, Rev. Mod. Phys. **95**, 045005 (2023).

[3] Jens Eisert, Dominik Hangleiter, Nathan Walk, Ingo Roth, Damian Markham, Rhea Parekh, Ulysse Chabaud, and Elham Kashefi, Quantum certification and benchmarking, Nat. Rev. Phys. **2**, 382 (2020).

[4] Martin Kliesch and Ingo Roth, Theory of quantum system certification, PRX Quantum **2**, 010201 (2021).

[5] Bei Zeng, Xie Chen, Duan-Lu Zhou, and Xiao-Gang Wen, *Quantum Information Meets Quantum Matter: From Quantum Entanglement to Topological Phases of Many-Body Systems* (Springer, New York, 2019).

[6] Andrew M. Childs, Lecture notes on quantum algorithms, 2025. URL: www.cs.umd.edu/amchilds/qa/qa.pdf.

[7] Peter W. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, Santa Fe, New Mexico, USA, 1994), pp. 124–134.

[8] Ashley Montanaro and Ronald de Wolf, *A Survey of Quantum Property Testing*. Graduate Surveys: Theory of Computing Library (Theory of Computing, 2016).

[9] Anurag Anshu and Srinivasan Arunachalam, A survey on the complexity of learning quantum states, Nat. Rev. Phys. **6**, 59 (2024).

[10] Srinivasan Arunachalam and Ronald de Wolf, Guest column: A survey of quantum learning theory, ACM SIGACT News **48**, 41 (2017).

[11] Adam Bouland, Tudor Giurgică-Tiron, and John Wright, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025), pp. 463–470. doi:10.1145/3717823.3718118.

[12] Aram W. Harrow, Applications of coherent classical communication and the Schur transform to quantum information theory, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 2005.

[13] Jens Eisert and John Preskill, Mind the gaps: The fraught road to quantum advantage, arXiv:2510.19928.

[14] Dominik Hangleiter and Michael J. Gullans, Bell sampling from quantum circuits, Phys. Rev. Lett. **133**, 020601 (2024).

[15] Hsin-Yuan Huang, Richard Kueng, and John Preskill, Information-theoretic bounds on quantum advantage in machine learning, Phys. Rev. Lett. **126**, 190505 (2021).

[16] Ashley Montanaro, Learning stabilizer states by Bell sampling, arXiv:1707.04012.

[17] Sitan Chen, Weiyuan Gong, Qi Ye, and Zhihan Zhang, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025), pp. 429–438. doi:10.1145/3717823.3718191

[18] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang, Efficient learning of quantum states prepared with few non-Clifford gates, Quantum **9**, 1907 (2025).

[19] David Gross, Sepehr Nezami, and Michael Walter, Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations, Commun. Math. Phys. **385**, 1325 (2021).

[20] Andrew M. Childs, Aram W. Harrow, and Paweł Wocjan, in *Proceedings of the 24th Annual Conference on Theoretical Aspects of Computer Science*, STACS'07 (Springer-Verlag, Berlin, Heidelberg, 2007), pp. 598–609.

[21] Andrew M. Childs and Wim van Dam, Quantum algorithms for algebraic problems, Rev. Mod. Phys. **82**, 1 (2010).

[22] Marcel Hinsche and Jonas Helsen, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25 (Association for Computing Machinery, New York, NY, USA, 2025), pp. 439–450. doi:10.1145/3717823.3718169.

[23] Scott Aaronson and Daniel Gottesman, *Identifying Stabilizer States* (Perimeter Institute, 2008).

[24] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024 (Association for Computing Machinery, New York, NY, USA, 2024), pp. 1352–1363.

[25] Jonathan Allcock, Joao F. Doriguello, Gábor Ivanyos, and Miklos Santha, Beyond Bell sampling: Stabilizer state learning and quantum pseudorandomness lower bounds on qudits, arXiv:2405.06357.

[26] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma, Learning t-doped stabilizer states, Quantum **8**, 1361 (2024).

[27] Dominic V. Else, Stephen D. Bartlett, and Andrew C. Doherty, Symmetry protection of measurement-based quantum computation in ground states, New J. Phys. **14**, 113016 (2012).

[28] Margarite L. LaBorde, Soorya Rethinasamy, and Mark M. Wilde, Testing symmetry on quantum computers, Quantum **7**, 1120 (2023).

[29] Soorya Rethinasamy, Margarite L. LaBorde, and Mark M. Wilde, Quantum computational complexity and symmetry, Can. J. Phys. **103**, 215 (2025).

[30] Victor M. Bastidas, Nathan Fitzpatrick, K. J. Joven, Zane M. Rossi, Shariful Islam, Troy Van Voorhis, Isaac L. Chuang, and Yuan Liu, Unification of finite symmetries in the simulation of many-body systems on quantum computers, Phys. Rev. A **111**, 052433 (2025).

[31] David Wakeham and Maria Schuld, Inference, interference and invariance: How the quantum Fourier transform can help to learn from data, arXiv:2409.00172.

[32] Erik Hostens, Jeroen Dehaene, and Bart De Moor, Stabilizer states and Clifford operations for systems of arbitrary dimensions, and modular arithmetic, Phys. Rev. A **71**, 042315 (2005).

[33] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, Rev. Mod. Phys. **93**, 045003 (2021).

[34] Dominik Hangleiter and Jens Eisert, Computational advantage of quantum random sampling, Rev. Mod. Phys. **95**, 035001 (2023).

[35] Hsin-Yuan Huang, Soonwon Choi, Jarrod R. McClean, and John Preskill, The vast world of quantum advantage, arXiv:2508.05720.

[36] Nai-Hui Chia, Ching-Yi Lai, and Han-Hsuan Lin, Efficient learning of *t*-doped stabilizer states with single-copy measurements, Quantum **8**, 1250 (2024).