

The state hidden subgroup problem and an efficient algorithm for locating unentanglement

Adam Bouland*
Stanford University

Tudor Giurgică-Tiron†
Stanford University, QuICS

John Wright‡
UC Berkeley

We study a generalization of entanglement testing which we call the “hidden cut problem.” Taking as input copies of an n -qubit pure state which is product across an unknown bipartition, the goal is to learn precisely *where* the state is unentangled, i.e. to determine which of the exponentially many possible cuts separates the state. We give a polynomial-time quantum algorithm which can find the cut using $O(n/\epsilon^2)$ many copies of the state, which is optimal up to logarithmic factors. Our algorithm also generalizes to learn the entanglement structure of arbitrary product states. In the special case of Haar-random states, we further show that our algorithm requires circuits of only constant depth. To develop our algorithm, we introduce a state generalization of the hidden subgroup problem (StateHSP) which might be of independent interest, in which one is given a quantum state invariant under an unknown subgroup action, with the goal of learning the hidden symmetry subgroup. We show how the hidden cut problem can be formulated as a StateHSP with a carefully chosen Abelian group action. We then prove that Fourier sampling on the hidden cut state produces similar outcomes as a variant of the well-known Simon’s problem, allowing us to find the hidden cut efficiently. Therefore, our algorithm can be interpreted as an extension of Simon’s algorithm to entanglement testing. We discuss possible applications of StateHSP and hidden cut problems to cryptography and pseudorandomness.

1 Introduction

Detecting the entanglement properties of states is a central theme in quantum information. In the standard formulation of product testing [MdW13], the goal is to determine if a state is product vs. far from product across a fixed bipartition, given as input copies of the state. The well-known “SWAP test” [GC01] provides a fundamental algorithmic primitive for entanglement testing in this setting. Variations of state product and separability testing have found many applications in quantum complexity theory [Gha08, GHMW13], and many problems which admit states as inputs can be reduced to questions of detecting entanglement — such as the proof that $\text{QMA}(k) = \text{QMA}(2)$ [HM13]. Recently a number of works have shown upper and lower bounds [BO20, SW22, FO24] for estimating various measures of entanglement/separability for quantum states.

In this work we study a generalization of product testing which we call the “hidden cut problem,” which removes the assumption of a pre-defined bipartition. Given as input copies of an n -qubit pure state, the task is not to determine *if* the state is unentangled, but rather to determine *where* it is unentangled. In other words, given a state which is promised to be a product of two $(n/2)$ -qubit states across some unknown “cut” (i.e. a bipartition of the n qubits), the goal is to learn the precise location of the cut. The combinatorics of set partitions makes the problem non-trivial: even though one can efficiently check any candidate cut via SWAP test on two state copies, there are exponentially many possible bipartitions of n qubits into two sets of size $n/2$, since $\frac{1}{2}\binom{n}{n/2} = 2^{\Theta(n)}$; therefore, a brute-force search is inefficient. See Figure 1 for an illustration. We define the hidden

*abouland@stanford.edu

†tgt@stanford.edu

‡jswright@berkeley.edu




Figure 1: A cartoon illustrating four out of the $\frac{1}{2}\binom{8}{4} = 35$ possible distinct ways (i.e. ‘cuts’) in which a pure state on $n = 8$ qubits can be formed as a product of two states $|\phi_1\rangle, |\phi_2\rangle$ on $\frac{n}{2} = 4$ qubits. The factor states $|\phi_1\rangle$ and $|\phi_2\rangle$ are depicted as contiguous clouds. The goal of the hidden cut problem is to determine which cut separates the product state.

cut problem more formally as follows:

Definition 1 (Hidden cut problem). *Let $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$ be a state on n qubits (where n is even) which is promised to be a product across an unknown cut $C \in \binom{[n]}{n/2}$, denoted $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$, such that the factor states $|\phi_1\rangle, |\phi_2\rangle \in (\mathbb{C}^2)^{\otimes n/2}$ are ϵ -far from being product states. The hidden cut problem asks to identify the cut C , given copies of the state $|\psi\rangle$.*

Note that in order for the cut to be well-defined, we require that the two unentangled factor states are far from product states themselves. Additionally, while we have defined the hidden cut problem to refer to equal-sized bipartitions of the qubits, one can of course generalize the problem to unequally sized cuts or to a product of more than two states, which we will address later in the paper.

A natural first question asks whether the hidden cut problem is even information-theoretically solvable given polynomially many copies of the input state. The answer is yes — a closely related problem has been analyzed in the property testing literature under the name of “multipartite productness”. If we interpret the hidden cut problem as a search task (which cut separates the state?), then multipartite productness is defined as the associated decision task (is there a cut which separates the state?). In [HLM17], Harrow, Lin, and Montanaro showed that only $O(n/\epsilon^2)$ state copies are information-theoretically required to decide if a state is multipartite product or ϵ -far from any such state, and this bound was recently shown to be optimal up to log factors by Jones and Montanaro [JM24]. A simple modification of Harrow, Lin and Montanaro’s argument shows the cut can also be located using only $O(n/\epsilon^2)$ copies of the input state (see Section 2.5). However, these property testing algorithms are computationally inefficient, and work by combining an exponentially long sequence of “gentle measurements” on the input state to try out the different possible cuts. This approach takes exponential time, and there is no obvious way to do better.

Our main result is a positive answer to this question: we construct an efficient algorithm for the hidden cut problem, which can learn the cut using $O(n/\epsilon^2)$ copies of the state and polynomial time. This is exponentially faster than prior property testing algorithms:

Theorem 1 (Hidden cut algorithm). *There is an efficient quantum algorithm for the hidden cut problem on n qubits which uses $O(n/\epsilon^2)$ copies of the state and runs in polynomial time, requiring circuits of depth $O(n^2) + O(\log \epsilon^{-1})$ which act coherently on $O(\epsilon^{-2})$ state copies at a time.*

The number of copies used by our algorithm is optimal up to log factors, in light of Jones and Montanaro’s decision lower bound [JM24]. Our algorithm works in an entirely different way than the prior information-theoretic approaches. In particular, we show the hidden cut problem can be formulated as a quantum state version of the well-known Abelian hidden subgroup problem (HSP), and then give an efficient algorithm to solve this quantum state HSP via a generalization of Simon’s

algorithm [Sim97]. Our work can thus be interpreted as an extension of Simon’s algorithm to entanglement testing.

In addition to being a natural question in entanglement testing, the hidden cut problem is also motivated by questions in quantum pseudorandomness and pseudoentanglement, since hiding the location of a separating cut could be a natural mechanism of hiding entanglement. A number of works have recently explored creating quantum pseudorandom states [JLS18] with low entanglement [ABF⁺24]. One natural recursive candidate construction¹ would consist of building larger states from products of two smaller pseudorandom states across a random partition. This could potentially result in a pseudorandom state construction with no entanglement across some partition, and which naturally lifts pseudorandom state construction on n qubits to pseudorandom state constructions on $n' > n$ qubits. Our result shows this approach does not work, as there is an efficient algorithm to locate unentanglement. Interestingly, our algorithm does not rule out the possibility of more general pseudorandom state constructions with low, but nonzero entanglement across hidden cuts (see Discussion section).

Our algorithm also admits a number of generalizations and improvements. First, our algorithm generalizes to the case in which the state is a product of two or more unentangled states which are not necessarily of the same size. Specifically, we naturally define the more generic “hidden many-cut problem”, in which the input state is allowed to be a product of several unentangled substates. We will show that the same algorithm can solve this version of the problem with minimal modifications, as long as the individual subsystems are entangled enough:

Corollary 1 (Algorithm for the many-cut problem — informal). *The same algorithm from Theorem 1 solves the “hidden many-cut problem” in which an n -qubit input state is product across an arbitrary set partition $C_1 \sqcup \dots \sqcup C_m = [n]$:*

$$|\psi\rangle = |\phi_1\rangle_{C_1} \otimes \dots \otimes |\phi_m\rangle_{C_m}, \tag{1}$$

such that each factor state $|\phi_k\rangle$ (where $k \in [m]$ indexes the m parts) is at least ϵ -far from any separable state on $|C_k|$ qubits. The algorithm identifies the set partition with the same resource and runtime requirements as in Theorem 1.

Second, a stronger promise about the internal entanglement structure of the input states can significantly reduce the runtime requirements of the algorithm. For example, if one assumes the input state is a product of two Haar random states, the algorithm works with constant-depth circuits, acting on only two copies at a time:

Theorem 2 (Hidden cut algorithm with Haar-random states). *Under the stronger promise of Haar-random factor states, the hidden cut can be found by the same algorithm with only $O(n)$ copies of the state, involving circuits of constant depth which coherently access only two state copies at a time.*

This highly efficient version of our algorithm still works when the factor states are not genuinely Haar, but rather computationally indistinguishable from Haar (i.e. pseudorandom states, for which known efficient constructions exist [JLS18]). We conjecture that the scope of this algorithm can be further extended to other families of states which obey a strong entanglement volume law. For these reasons, this highly efficient version of our algorithm could be of particular relevance to near-term experiments.

¹We thank Henry Yuen for raising this question.




Figure 2: An illustration of the permutational symmetries of the tn -qubit global state $|\psi\rangle^{\otimes t}$ when $|\psi\rangle$ is product across a cut $C \subset [n]$, depicted for $n = 8$ and $t = 6$. The same cross-copy permutation applied to all qubits inside each side of the cut leaves the global state invariant; to each hidden cut corresponds a hidden subgroup isomorphic to $S_t^{\times 2}$ inside the group $S_t^{\times n}$. In [Section 4.1](#), we show that Fourier sampling with this group action fails to identify the cut.

1.1 A hidden subgroup problem for states

We will now describe a conceptual framework which will motivate our design of the algorithm for the hidden cut problem. As with many other quantum algorithms, the key is to make critical use of the symmetries of the input states. The conceptual contribution is to recognize that the hidden cut problem, as well as potentially many other problems with state inputs, can be formulated within a quantum state generalization of the well-known hidden subgroup problem.


We start by recalling that a core algorithmic design principle for state-input problems is accounting for the symmetries of the global state $|\psi\rangle^{\otimes t}$ made up of copies of the input state $|\psi\rangle$ (see e.g. [\[MdW13, OW16\]](#)). In the hidden cut problem, we observe that the global state will have various internal symmetries which are determined by the location of the hidden cut. In other words, we can define a group action on the global state such that each hidden cut will correspond to a unique subgroup of hidden symmetries. This is reminiscent of the *hidden subgroup problem* (HSP), a central framework in quantum algorithms and complexity [\[NC10, Section 5.4.3\]](#), in which the task is also to identify a hidden subgroup H given a function on a parent group G which is invariant under H . However, there is a fundamental difference as the HSP takes as input a *function* with specific subgroup symmetries. In contrast, the hidden cut problem takes as input *quantum states* with particular sets of symmetries.

Motivated by this observation, we define a quantum state version of the HSP, which we call the *state hidden subgroup problem* (StateHSP). This problem takes as input (copies of) a state which admits an efficient action of a finite group, such that the state is preserved by an unknown subgroup; the task is once again to identify the symmetry subgroup:


Definition 2 (The state hidden subgroup problem (StateHSP) — informal). *Let G be a finite group with a unitary representation $R : G \rightarrow \text{U}(d)$. The goal is to identify the unknown hidden subgroup $H < G$, given access to the representation and (copies of) a quantum state $|\psi\rangle \in \mathbb{C}^d$ with the following properties:*

- $|\psi\rangle$ is invariant under the action of the subgroup H , i.e. for all $h \in H$, $R(h)|\psi\rangle = |\psi\rangle$.
- $|\psi\rangle$ is acted on nontrivially by elements outside the subgroup: for any $g \notin H$, $|\langle\psi|R(g)|\psi\rangle| \leq 1 - \epsilon$.

StateHSP can be interpreted as a generalization of the standard HSP in the following concrete sense: the canonical approach to the standard hidden subgroup problem already involves the construction



(a) The action of a group element $\mathbf{x} \in \mathbb{Z}_2^n$ on the global input state.



(b) The subgroup operations preserving the global input state.

Figure 3: The Abelian group action of the group $G = \mathbb{Z}_2^n$ on the tn -qubit global state $|\psi\rangle^{\otimes t}$ which underlies our hidden cut algorithm, depicted for $n = 8$ and $t = 4$. To each hidden cut $C \subset [n]$ corresponds an order-four hidden subgroup $H_C \simeq \mathbb{Z}_2^2$ which leaves the state invariant. The subgroup is generated by the two equivalent n -bit strings which encode the cut and its complement: $H_C = \langle 1^C 0^{\bar{C}}, 0^C 1^{\bar{C}} \rangle$.

of states with subgroup symmetries in the form of *coset states* [CVD10]. Whereas the coset states live in the regular representation of the group, StateHSP generalizes the state problem to arbitrary group representations. We will elaborate in more technical detail in Section 3. We hope the StateHSP problem might be of independent interest, as it is a natural generalization of the HSP.

In order to apply this framework to the hidden cut problem, one must (a) find an appropriate group action such that the hidden cut problem is formulated as a StateHSP, and (b) find an algorithm to solve the corresponding StateHSP. For the latter problem, we will later show that techniques for solving the standard HSP – such as an appropriate generalization of Fourier sampling – can be ported over to the StateHSP setting. But first we need to explain why the hidden cut problem is a StateHSP in the first place.

1.2 The Hidden cut problem as an Abelian StateHSP

In order to describe the hidden cut problem as a StateHSP, we must first find an appropriate group action on the global state $|\psi\rangle^{\otimes t}$ such that each hidden cut corresponds to a specific subgroup. In the hidden cut problem we make no assumption about the internal structure of the factor states $|\phi_{1,2}\rangle$ beyond high entanglement; therefore, the symmetries of the problem lie in acting across, not along, the copies of the input state. Any t -fold global state $|\psi\rangle^{\otimes t}$ has a trivial permutational symmetry group \mathbb{S}_t which permutes the t copies. However, since each copy is internally separable along a hidden cut $C \subset [n]$, there is a larger permutational symmetry group which leaves the global state invariant. In particular, permutations which act simultaneously on all qubits within each side of the cut also preserve the global state (see Figure 2). This would formulate the hidden cut problem as a StateHSP over the parent group $G = (\mathbb{S}_t)^{\times n}$, such that the hidden subgroups are promised to be isomorphic to $(\mathbb{S}_t)^{\times 2}$. This would seem to be the most general set of permutational symmetries of the global state. Furthermore, in the standard HSP, Fourier sampling is efficiently implementable given known circuits for the non-Abelian quantum Fourier transform on the symmetric group [Bea97], so there is hope this group action could result in an algorithm for hidden cuts. However, as we will show in Section 4.1, this Fourier sampling method fails to find the hidden cut, for similar reasons that Fourier sampling fails to solve the standard HSP over the symmetric group [MRS08]. Thus, the most obvious StateHSP approach to the hidden cut problem does not work.

Our key observation is that a much simpler Abelian group action can be used to define a StateHSP for the hidden cut problem. As it turns out, it is possible to restrict the permutational symmetries to a subset isomorphic to the group $G = \mathbb{Z}_2^n$, by considering simple SWAPs of pairs of qubits. Concretely, we consider dividing the t copies of the input state (assuming t is even) into pairs; the action of the i -th entry of the n -bit string $x \in \{0, 1\}^n$ is to SWAP the i -th qubits inside each pair (see Figure 3). The key point is that swapping all the qubits within each side of the cut leaves the global state invariant. Therefore, to each possible hidden cut $C \subset [n]$ corresponds a hidden subgroup of order four isomorphic to \mathbb{Z}_2^2 , which contains all operations acting simultaneously on all qubits on either side of the cut. For example, if the cut is between the first and second $n/2$ qubits, the hidden subgroup is the group with elements $\{0^n, 0^{n/2}1^{n/2}, 1^{n/2}0^{n/2}, 1^n\}$ under bitwise addition mod 2, because this subgroup of SWAP operations preserves the paired copies of the input state by exchanging the left/right halves of the paired states. Therefore, this choice of group action successfully formulates the hidden cut problem as an Abelian StateHSP instance.

1.3 Solving the Abelian StateHSP via Fourier Sampling

Having identified the hidden cut problem as an Abelian StateHSP over $G = \mathbb{Z}_2^n$, it remains to show how to efficiently solve it. Recall that standard Abelian HSP instances can be efficiently solved by Fourier sampling. We will show that a generalization of Fourier sampling can be transplanted to the StateHSP problem, resulting in an efficient algorithm to find the hidden cut, or more generally to solve any Abelian StateHSP (see Fact 3.3). The algorithm follows a familiar Fourier sampling workflow: we first prepare an equal superposition of group elements, then apply the controlled group action to the input state(s), and finally take a Fourier transform followed by a measurement on the group register. The circuit implementation of this approach is particularly simple, with the added benefit of parallelization over the n ancillary qubits which make up the \mathbb{Z}_2^n group register (see Figure 4).

The main question we need to answer next is how the output of this state Fourier sampling circuit relates to that of the equivalent standard HSP algorithm, i.e. the standard hidden subgroup problem defined with the same parent group, and the same set of valid hidden subgroups. The technical sections of this paper focus on precisely understanding the output distribution of Fourier samples arising from the hidden cut problem. A key observation is the way in which this measurement outcome distribution depends on the number of copies of the input states t one uses to produce each Fourier sample. Specifically, increasing the number of copies t behaves as a form of orthogonality amplification, resulting in the output distribution approaching the “ideal” distribution induced by the associated standard Abelian HSP. In other words, the ability to act coherently on multiple copies at a time makes the hidden cut problem behave more like the corresponding standard Abelian HSP. To see why, consider all the states obtained by group action from the initial input state, i.e. the group orbit of the initial state. A group element can either be inside the hidden subgroup (in which case it preserves the input state), or outside the hidden subgroup (in which case it does not), meaning that each distinct state in the orbit corresponds to a coset of the hidden subgroup. Acting coherently on several copies of the state at the same time exponentially suppresses the inner product between the states along the orbit of the group action. Intuitively, this effectively orthogonalizes the orbit states; the case of orthogonal coset states is precisely the regime of the standard HSP. This effect is crucial to our algorithm, because it essentially means the hidden cut problem can be reduced to an Abelian HSP, from the point of view of Fourier sampling.

Concretely, an input involving t copies of a specific state $|\psi\rangle$ will induce a specific distribution

$P_{\text{StateHSP}_{\psi,t}}[\mathbf{y}]$ over the measurement outcomes $\mathbf{y} \in \mathbb{Z}_2^n$ of the Fourier sampling circuit. In [Section 4](#) we describe the technical error analysis which allows us to appropriately choose the number of state copies t . Specifically, this number of copies is chosen such that the output distribution of the hidden cut Fourier sampling circuit $P_{\text{StateHSP}_{\psi,t}}$ becomes negligibly close to the equivalent standard HSP outcome distribution P_{HSP} in a multiplicative sense:

$$P_{\text{StateHSP}_{\psi,t}}[\mathbf{y}] = P_{\text{HSP}}[\mathbf{y}] (1 + \text{negl}(n)), \quad \text{for all } \mathbf{y} \in \mathbb{Z}_2^n. \quad (2)$$

Here, $P_{\text{HSP}}[\mathbf{y}]$ denotes the probability to obtain outcome $\mathbf{y} \in \mathbb{Z}_2^n$ via Fourier sampling in the associated standard HSP with the same group and subgroup specifications as our StateHSP. This associated HSP Fourier sampling distribution is particularly simple:

$$P_{\text{HSP}}[\mathbf{y}] = \begin{cases} 2^{-n+2} & \text{if } \mathbf{y} \cdot 1^C 0^{\bar{C}} = \mathbf{y} \cdot 0^C 1^{\bar{C}} = 0 \pmod{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Specifically, this means that the associated HSP distribution P_{HSP} is uniformly supported on the Boolean subspace of dimension $n - 2$ which is orthogonal to the two equivalent bit strings² $1^C 0^{\bar{C}}$, $0^C 1^{\bar{C}}$ which encode the hidden cut C . We remark that this is a variation of the classic Simon's problem [[Sim97](#)], in which the Fourier samples are also uniform over the subspace orthogonal to a secret string. The fact that we obtain a multiplicative error in the output distribution of Simon's algorithm means that we will never observe a string outside the orthogonal subspace when obtaining the Fourier samples.

Once the number of copies t is chosen such that the hidden cut problem returns similar outcomes as the associated Simon-like HSP, the original logic of Simon's algorithm allows us to efficiently find the hidden cut: after obtaining $O(n)$ independent Fourier samples, one has collected a complete basis for the supporting subspace with high probability, from which the secret string $1^C 0^{\bar{C}}$ which encodes the cut C can be learned by solving a simple Boolean linear algebra problem of size n . Therefore, our algorithm can be viewed as an extension of Simon's algorithm to entanglement testing, since finding the hidden cut reduces to solving a Simon-like Abelian HSP over the group $G = \mathbb{Z}_2^n$.

As will be detailed in [Section 4](#), a straightforward application of Abelian Fourier sampling to the hidden cut problem succeeds in finding the cut, however a number $O(n^2/\epsilon^2)$ of copies are required. A key observation is that it is possible to further reduce the requirement by a factor of n , down to the optimal $O(n/\epsilon^2)$ asymptotic of [Theorem 1](#), by an adaptive modification of the Fourier sampling procedure. Specifically, we show how this can be achieved by changing the initialization of the ancillary group register. Whereas the standard Fourier sampling approach involves starting with a uniform superposition over all group elements, in our second adaptive algorithm we will start with a superposition over the \mathbb{Z}_2^n elements which are orthogonal to previous samples. We will show how this choice boosts the probability of measuring new linearly independent samples, such that a number of copies $t = O(1/\epsilon^2)$ at each sampling round suffices to produce valid measurements from the cut subspace with constant success probability.

Different promises on the internal entanglement of the factor states will ultimately impact the number of state copies t required for orthogonality amplification. Following the formulation common

²The notation $1^C 0^{\bar{C}}$ denotes the n -bit string with 1's in the positions in $C \subseteq [n]$ and 0's elsewhere.

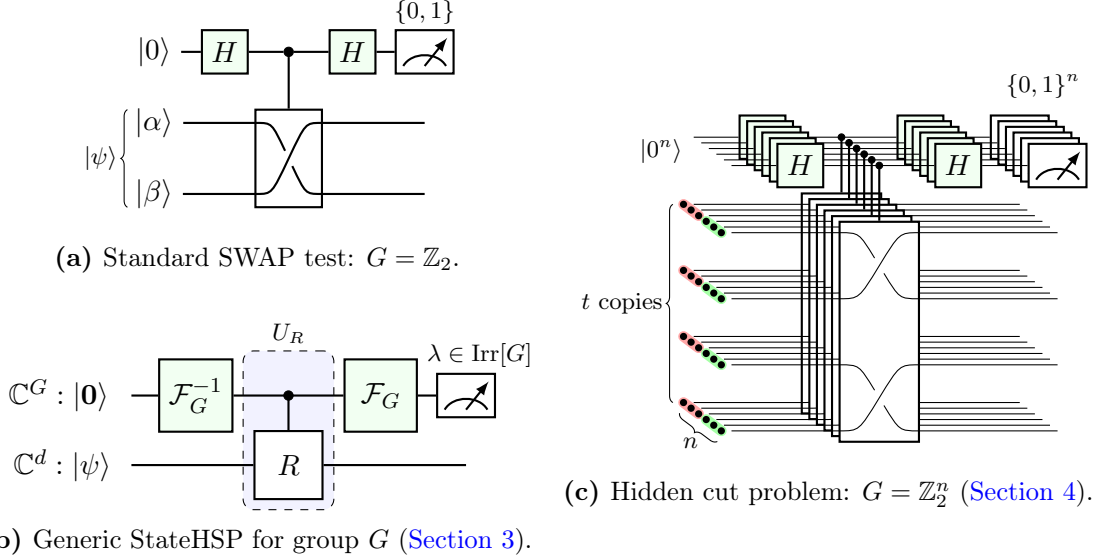


Figure 4: Group Fourier sampling circuits for specific cases of the state hidden subgroup problem. In all the examples, the top ancillary register supports a regular representation of the group, and the central operation is a controlled group action on the input state. (a) The standard SWAP test. (b) The generic StateHSP for arbitrary group G admitting circuit implementations of the quantum group Fourier transform \mathcal{F}_G , and which acts by a unitary representation R on the input state. The top register is initialized in the basis state associated with the trivial representation of G . The outcome is a label $\lambda \in \text{Irr}[G]$ of a group irreducible representation. (c) Schematic of the circuit which solves the hidden cut problem as part of the non-adaptive Algorithm 1.

to property testing scenarios, the generic version of the hidden cut problem (Definition 1) promises that the factor states are at least a constant trace distance ϵ away from separable. This will require a number of copies $t = O(1/\epsilon^2)$ per Fourier sample in order to exhaustively suppress the contributions coming from all of the possible false internal cuts. This count can be reduced to a constant of only $t = 2$ if the promise is strengthened to Haar-random factor states; furthermore, the corresponding circuits require only a constant depth. This improvement requires several changes to the analysis specific to the special case of Haar-random factor states (Theorem 2), which we will detail in Section 5. First, we show that in this case the Fourier sampling distribution self-averages in a particularly strong sense. Second, we relax the strong multiplicative error condition mentioned above, and generalize Simon’s algorithm to a setting which no longer involves uniform samples from the orthogonal subspace, but is skewed towards lower-weight strings. We will show that the Simon-like “basis coupon collection” process via Fourier sampling nonetheless succeeds to find the hidden cut under this modification.

We also note this algorithm directly generalizes to the multicut case — as this simply corresponds to larger Abelian subgroups of this same group action, where the subgroup is generated by $1^{C_i}0^{[n]\setminus C_i}$ for any sub-partition of the qubits C_i . The main challenge again is to carefully keep track of the errors in the Fourier sampling distribution in this more general setting. One can also observe our algorithm does not require knowing the number of cuts in advance, as this can be efficiently inferred from the linear algebra of the obtained Fourier samples.

We remark that one can interpret our algorithm as a combinatorial generalization of the standard SWAP test, the canonical primitive for state comparison in property testing. The SWAP test is

indeed a simple instance of StateHSP for the group $G = \mathbb{Z}_2$ which acts as an exchange operation, with Fourier sampling implementing the projective measurement against the symmetric and anti-symmetric subspaces. Our own Fourier sampling circuit for the hidden cut problem consists of n parallel amplified SWAP tests which are only entangled through the internal structure of the input state (see [Figure 4](#)).

1.4 Outline of the paper

[Section 2](#) contains basic preliminaries, as well as the exponential-time algorithm which solves the hidden cut problem using $O(n/\epsilon^2)$ copies. In [Section 3](#), we define the state version of the hidden subgroup problem. We adapt the Fourier sampling algorithm to the state problem, and describe a setting in which the state version and the standard version of the hidden subgroup problem produce similar outcomes. In [Section 4](#), we describe the efficient algorithm for the hidden cut problem and prove [Theorem 1](#). By taking advantage of the permutational symmetries of the global state, we design an Abelian group action which fits into the StateHSP framework of [Section 3](#), and show how the Fourier sampling outcomes concentrate towards a version of Simon’s algorithm. We will first describe a non-adaptive, Simon-like Fourier sampling algorithm ([Algorithm 1](#)) which finds the cut given $O(n^2/\epsilon^2)$ state copies. Subsequently, we will introduce an adaptive modification of the algorithm ([Algorithm 2](#)) and show how this decreases the state copy requirement to the optimal value of $O(n/\epsilon^2)$. [Section 5](#) is dedicated to the special case of Haar-random states, which will require a more in-depth technical analysis. Specifically, to find the cut with constant-depth circuits in the Haar-random case ([Theorem 2](#)), we will show a self-averaging property of the Fourier sampling distribution, by approximately diagonalizing the covariance matrix of internal purities of Haar-random states; additionally, this special case requires a modification of Simon’s algorithm to allow non-uniform samples. In [Section 6](#) we generalize our results to the “hidden many-cut problem”, showing how the same algorithm can solve not just for a single bipartition, but can similarly identify arbitrary product state structures. Finally, in [Section 7](#) we discuss open questions and possible applications of the hidden cut problem and StateHSP to cryptography and pseudorandomness.

2 Preliminaries

We start by collecting a few basic notions about the geometry of quantum states. For more background, we refer readers to a standard reference such as [\[NC10\]](#).

2.1 Distances

Definition 3 (Overlap). *The overlap of two pure states $|\psi\rangle, |\phi\rangle \in \mathbb{C}^d$ is given by $|\langle\psi|\phi\rangle|^2$.*

Definition 4 (Trace distance). *The trace distance between two mixed states $\rho, \sigma \in \mathbb{C}^{d \times d}$ is given by*

$$D_{\text{tr}}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1,$$

where $\|\cdot\|_1$ denotes the trace norm, also known as the Schatten 1-norm. If $\rho = |\psi\rangle\langle\psi|$ and $\sigma = |\phi\rangle\langle\phi|$ are both pure states, then

$$D_{\text{tr}}(\rho, \sigma) = \sqrt{1 - |\langle\psi|\phi\rangle|^2}. \tag{4}$$

Definition 5 (Distance from a subset). Let \mathcal{H} be a Hilbert space, and let \mathcal{P} be a subset of the pure states in \mathcal{H} . Then $|\psi\rangle$ is ϵ -far from \mathcal{P} if

$$D_{\text{tr}}(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \geq \epsilon$$

for all states $|\phi\rangle$ in \mathcal{P} . Via (4), this is equivalent to

$$|\langle\psi|\phi\rangle|^2 \leq 1 - \epsilon^2$$

for all states $|\phi\rangle$ in \mathcal{P} .

2.2 Product states

Definition 6 (Product states). Let \mathcal{H}_A and \mathcal{H}_B be Hilbert spaces. Then a product state on $\mathcal{H}_A \otimes \mathcal{H}_B$ is a state of the form $|a\rangle_A \otimes |b\rangle_B$. If the bipartition of the overall Hilbert space $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$ is clear from context, we will usually simply refer to $|\psi\rangle$ as a product state.

Although not every state is a product state, every state can be written as a superposition of product states which are orthogonal on both their A and B registers. This is given by the Schmidt decomposition.

Definition 7 (Schmidt decomposition). Let $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ be a bipartite quantum state. Suppose \mathcal{H}_A and \mathcal{H}_B have dimensions d_A and d_B , respectively, and write $r := \min\{d_A, d_B\}$. The Schmidt decomposition of $|\psi\rangle$ is given by

$$|\psi\rangle_{AB} = \sum_{i=1}^r \sqrt{\lambda_i} \cdot |u_i\rangle_A \otimes |v_i\rangle_B,$$

where (i) $\lambda_1, \dots, \lambda_r$ are nonnegative real numbers such that $\lambda_1 + \dots + \lambda_r = 1$, (ii) $|u_1\rangle, \dots, |u_r\rangle$ are orthonormal vectors in \mathcal{H}_A , and (iii) $|v_1\rangle, \dots, |v_r\rangle$ are orthonormal vectors in \mathcal{H}_B . The numbers $\lambda_1, \dots, \lambda_r$ are known as $|\psi\rangle$'s Schmidt coefficients.

Thus, $|\psi\rangle$ is a product state if and only if its largest Schmidt coefficient is equal to 1 and all other Schmidt coefficients are equal to 0. The next lemma shows a robust version of this statement, namely that $|\psi\rangle$'s maximum Schmidt coefficient is exactly its largest overlap with any product state.

Proposition 1. Suppose $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ has Schmidt coefficients $\lambda_1 \geq \dots \geq \lambda_r$. Then $|\psi\rangle$'s maximum squared overlap with any product state is equal to λ_1 .

Proof. First, we show that $|\psi\rangle$'s maximum overlap with any product state is at least λ_1 . Consider the product state $|u_1\rangle_A \otimes |v_1\rangle_B$. Then

$$|\langle u_1|_A \otimes \langle v_1|_B \cdot |\psi\rangle|^2 = \left| \langle u_1|_A \otimes \langle v_1|_B \cdot \left(\sum_{i=1}^r \sqrt{\lambda_i} \cdot |u_i\rangle_A \otimes |v_i\rangle_B \right) \right|^2 = \left| \sqrt{\lambda_1} \right|^2 = \lambda_1.$$

Next, we show that $|\psi\rangle$'s maximum overlap with any product state is at most λ_1 . Let $|a\rangle_A \otimes |b\rangle_B$ be a product state. Then

$$|\langle a|_A \otimes \langle b|_B \cdot |\psi\rangle|^2 = \left| \langle a|_A \otimes \langle b|_B \cdot \left(\sum_{i=1}^r \sqrt{\lambda_i} \cdot |u_i\rangle_A \otimes |v_i\rangle_B \right) \right|^2$$

$$\begin{aligned}
&= \left| \sum_{i=1}^r \sqrt{\lambda_i} \cdot \langle a|u_i\rangle \cdot \langle b|v_i\rangle \right|^2 \\
&\leq \left(\sum_{i=1}^r \lambda_i \cdot |\langle a|u_i\rangle|^2 \right) \cdot \left(\sum_{i=1}^r |\langle b|v_i\rangle|^2 \right), \tag{5}
\end{aligned}$$

where the last step used the Cauchy-Schwarz inequality. Because $|u_1\rangle, \dots, |u_r\rangle$ are orthonormal and $|v_1\rangle, \dots, |v_r\rangle$ are orthonormal, we have that

$$|\langle a|u_1\rangle|^2 + \dots + |\langle a|u_r\rangle|^2 \leq 1 \quad \text{and} \quad |\langle b|v_1\rangle|^2 + \dots + |\langle b|v_r\rangle|^2 \leq 1.$$

Plugging this into (5), we get that

$$(5) \leq \sum_{i=1}^r \lambda_i \cdot |\langle a|u_i\rangle|^2 \leq \sum_{i=1}^r \lambda_1 \cdot |\langle a|u_i\rangle|^2 \leq \lambda_1.$$

This completes the proof. □

Combining [Proposition 1](#) with (4) gives us the following immediate corollary.

Corollary 2. *Suppose $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ is ϵ -far from any product state. Then its maximum Schmidt coefficient is at most $1 - \epsilon^2$.*

An equivalent characterization of product states is that $|\psi\rangle_{AB}$ is product if and only if the reduced density matrix ψ_A is a pure state. The purity is an analytic measure for how pure a density matrix is.

Definition 8 (Purity). *Given a mixed state $\rho \in \mathbb{C}^{d \times d}$, its purity is the quantity $\text{Tr}(\rho^2)$.*

The following proposition shows that the purity of ψ_A can be used as a measure for how close $|\psi\rangle_{AB}$ is to being a product state.

Proposition 2. *Suppose $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ is ϵ -far from any product state. Then the purity of ψ_A is at most*

$$\text{Tr}(\psi_A^2) \leq 1 - \epsilon^2.$$

Proof. Write $\lambda_1 \geq \dots \geq \lambda_r$ for the Schmidt coefficients of $|\psi\rangle_{AB}$. Then [Corollary 2](#) implies that $\lambda_1 \leq 1 - \epsilon^2$. Thus, we can bound the purity of ψ_A by

$$\text{Tr}(\psi_A^2) = \sum_{i=1}^r \lambda_i^2 \leq \sum_{i=1}^r \lambda_1 \cdot \lambda_i = \lambda_1 \cdot \left(\sum_{i=1}^r \lambda_i \right) = \lambda_1 \leq 1 - \epsilon^2.$$

This completes the proof. □

2.3 Purity testing

Given a mixed state

$$\rho = \sum_{i=1}^d \alpha_i \cdot |v_i\rangle\langle v_i|,$$

testing if it is actually a pure state is impossible with only one copy of ρ because no matter how far from pure ρ is, a single copy of it can be always viewed as a mixture over pure states. It turns out, however, that *two* copies of ρ , i.e.

$$\rho^{\otimes 2} = \sum_{i=1}^d \sum_{j=1}^d \alpha_i \alpha_j \cdot |v_i\rangle\langle v_i| \otimes |v_j\rangle\langle v_j|, \quad (6)$$

suffice for this task, because if ρ is not a pure state, this mixture will contain nonzero weight on terms $|v_i\rangle\langle v_i| \otimes |v_j\rangle\langle v_j|$ for which $i \neq j$, consisting of two orthogonal pure states. We need only be able to detect when two pure states are orthogonal rather than equal, and this can be done via the well-known *SWAP test* procedure [GC01].

The most basic component of the SWAP test is the SWAP gate:

Definition 9 (The SWAP gate). *Let d be an integer. The SWAP gate is the unitary operator SWAP acting on $\mathbb{C}^d \otimes \mathbb{C}^d$ such that*

$$\text{SWAP} \cdot |i\rangle \otimes |j\rangle = |j\rangle \otimes |i\rangle,$$

for all $1 \leq i, j \leq d$.

Then the SWAP test acts as follows.

Definition 10 (The SWAP test). *Let ρ_{AB} be a mixed state in $\mathbb{C}^d \otimes \mathbb{C}^d$ (which will typically be a tensor product state $\rho_A \otimes \sigma_B$). The SWAP test is the quantum algorithm which acts as follows. Beginning with the input state ρ_{AB} , (i) append an ancilla qubit in the $|+\rangle_{\text{Anc}}$ state. Next, (ii) apply SWAP_{AB} conditioned on the ancilla qubit and then (iii) Hadamard the ancilla qubit. Finally, (iv) measure the ancilla qubit and accept if the outcome is “0”. An illustration of the SWAP test applied to a product state $|\alpha\rangle \otimes |\beta\rangle$ can be found in Figure 4a.*

The SWAP test can equivalently be stated in terms of a two-outcome projective measurement.

Proposition 3 (SWAP test, projector version). *Write Π_{SWAP} for the projector $\frac{1}{2}(I_{AB} + \text{SWAP}_{AB})$. Then the SWAP test implements the projective measurement $\{\Pi_{\text{SWAP}}, I - \Pi_{\text{SWAP}}\}$.*

Proof. Given an input state $|\psi\rangle_{AB}$, the SWAP test acts as follows.

$$\begin{aligned} |\psi\rangle_{AB} &\longrightarrow |+\rangle_{\text{Anc}} \otimes |\psi\rangle_{AB} && \text{(append the ancilla)} \\ &\longrightarrow \frac{1}{\sqrt{2}} \cdot |0\rangle_{\text{Anc}} \otimes |\psi\rangle_{AB} + \frac{1}{\sqrt{2}} \cdot |1\rangle_{\text{Anc}} \otimes (\text{SWAP}_{AB} \cdot |\psi\rangle_{AB}) && \text{(apply the controlled SWAP)} \\ &\longrightarrow \frac{1}{\sqrt{2}} \cdot |+\rangle_{\text{Anc}} \otimes |\psi\rangle_{AB} + \frac{1}{\sqrt{2}} \cdot |-\rangle_{\text{Anc}} \otimes (\text{SWAP}_{AB} \cdot |\psi\rangle_{AB}). && \text{(Hadamard the ancilla)} \end{aligned}$$

This state is equal to

$$\begin{aligned} &\frac{1}{2} \cdot |0\rangle_{\text{Anc}} \otimes (|\psi\rangle_{AB} + \text{SWAP}_{AB} \cdot |\psi\rangle_{AB}) + \frac{1}{2} \cdot |1\rangle_{\text{Anc}} \otimes (|\psi\rangle_{AB} - \text{SWAP}_{AB} \cdot |\psi\rangle_{AB}) \\ &= |0\rangle_{\text{Anc}} \otimes (\Pi_{\text{SWAP}} \cdot |\psi\rangle_{AB}) + |0\rangle_{\text{Anc}} \otimes ((I - \Pi_{\text{SWAP}}) \cdot |\psi\rangle_{AB}), \end{aligned}$$

where here we used the fact that $I - \Pi_{\text{SWAP}} = \frac{1}{2}(I_{AB} - \text{SWAP}_{AB})$. The SWAP test concludes by measuring the ancilla in the standard basis, which concludes the proof. \square

Hence, the probability that the SWAP test accepts on $\rho^{\otimes 2}$ is $\text{Tr}(\Pi_{\text{SWAP}} \cdot \rho^{\otimes 2}) = \frac{1}{2} + \frac{1}{2} \cdot \text{Tr}(\text{SWAP} \cdot \rho^{\otimes 2})$. The next proposition computes the second term.

Proposition 4 (Purity formula).

$$\text{Tr}(\text{SWAP} \cdot \rho^{\otimes 2}) = \text{Tr}(\rho^2).$$

Proof. Write $\rho^{\otimes 2}$ as in (6). Then we have

$$\text{Tr}(\text{SWAP} \cdot |v_i\rangle\langle v_i| \otimes |v_j\rangle\langle v_j|) = \text{Tr}(|v_j\rangle\langle v_i| \otimes |v_i\rangle\langle v_j|) = |\langle v_i | v_j \rangle|^2 = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Extending via linearity,

$$\text{Tr}(\text{SWAP} \cdot \rho \otimes \rho) = \sum_{i=1}^d \sum_{j=1}^d \alpha_i \alpha_j \cdot 1[i = j] = \sum_{i=1}^d \alpha_i^2 = \text{Tr}(\rho^2).$$

□

Putting everything together gives the following formula for the probability the SWAP test accepts.

Corollary 3 (SWAP test acceptance probability). *The probability the SWAP test accepts on $\rho^{\otimes 2}$ is $\frac{1}{2} + \frac{1}{2} \cdot \text{Tr}(\rho^2)$.*

Thus, if ρ is pure, i.e. its purity is 1, then the SWAP test will always accept, but if ρ is very mixed, i.e. its purity is close to 0, then the SWAP test will accept with probability roughly $\frac{1}{2}$.

This also gives an algorithm for testing if a bipartite pure state $|\psi\rangle_{AB}$ is entangled given just two copies $|\psi\rangle_{AB} \otimes |\psi\rangle_{A'B'}$: simply run the SWAP test on the A and A' registers of this two-copy state, which is equivalent to running the SWAP test on $\psi_A^{\otimes 2}$. Doing so will accept with probability $\frac{1}{2} + \frac{1}{2} \cdot \text{Tr}(\psi_A^2)$, which is equal to 1 if $|\psi\rangle_{AB}$ is a product state but is at most $1 - \epsilon^2/2$ if $|\psi\rangle_{AB}$ is ϵ -far from product (via [Proposition 2](#)). This algorithm can be thought of as exploiting the fact that $|\psi\rangle_{AB} \otimes |\psi\rangle_{A'B'}$ is unchanged by applying $\text{SWAP}_{AA'}$ if and only if $|\psi\rangle$ is a product state.

2.4 Multipartite product states

Notation 1 (n -qubit systems). *Much of this paper is about n -qubit systems. Given a subset $S \subseteq [n]$ of the qubits, we will write $\bar{S} \equiv [n] \setminus S$ for the qubits outside of S . We will write \mathcal{H}_S for the Hilbert space consisting of the qubits in S , and so we will often write a state in \mathcal{H}_S as $|\psi\rangle_S$, i.e. with the “ S ” subscript.*

Definition 11 (Multipartite product states). *An n -qubit state $|\psi\rangle$ is a multipartite product state if there exists a subset of the qubits $C \subseteq [n]$ such that $|\psi\rangle$ can be written as $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\bar{C}}$, for some states $|a\rangle_C$ and $|b\rangle_{\bar{C}}$ supported on the qubits in C and \bar{C} , respectively.*

We will often consider the case when $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\bar{C}}$ in which $|a\rangle_C$ and $|b\rangle_{\bar{C}}$ are both ϵ -far from any multipartite product state, and our goal is to determine C . It is natural to pick a subset S and test if $C = S$ by running the product test on the qubits within S . To analyze this, we first show the following proposition.

Proposition 5 (Purity across different cuts). *Let $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\overline{C}}$ in which $|a\rangle_C$ and $|b\rangle_{\overline{C}}$ are both ϵ -far from any multipartite product state. Let $S \subseteq [n]$ be a subset of the qubits. Then the purity of ψ_S is $\text{Tr}(\psi_S^2) = 1$ if $S = C$ or \overline{C} and otherwise $\text{Tr}(\psi_S^2) \leq 1 - \epsilon^2$.*

Proof. If $S = C$ then $\psi_S = |a\rangle\langle a|$, which is a pure state, and so its purity is 1; a similarly argument applies when $S = \overline{C}$. On the other hand, when $S \neq C, \overline{C}$, we have that

$$\psi_S = a_{C \cap S} \otimes b_{\overline{C} \cap S}.$$

Because $S \neq C, \overline{C}$, it must be the case that either $\emptyset \subsetneq C \cap S \subsetneq C$ or $\emptyset \subsetneq \overline{C} \cap S \subsetneq \overline{C}$; let us assume without loss of generality that the former is true. Then because $|a\rangle_C$ is ϵ -far from multiproduct, it is ϵ -far from being a product state on the bipartition $(C \cap S, C \cap \overline{S})$. Hence, by [Proposition 2](#), we can bound its purity by

$$\text{Tr}(a_{C \cap S}^2) \leq 1 - \epsilon^2.$$

As a result, we can bound the purity of the overall state by

$$\text{Tr}(\psi_S^2) = \text{Tr}(a_{C \cap S}^2) \cdot \text{Tr}(b_{\overline{C} \cap S}^2) \leq \text{Tr}(a_{C \cap S}^2) \leq 1 - \epsilon^2.$$

□

Combining this with [Corollary 3](#), we get the following bound on the probability that the SWAP test on the qubits in S accepts.

Corollary 4. *Let $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\overline{C}}$ in which $|a\rangle_C$ and $|b\rangle_{\overline{C}}$ are both ϵ -far from any multipartite product state. Suppose we are given two copies of $|\psi\rangle$ and we run the SWAP test on some subset $S \subseteq [n]$ of the qubits. Then if $S = C$ or $S = \overline{C}$, the SWAP test always accepts. Otherwise, if $S \neq C, \overline{C}$,*

$$\mathbb{P}[\text{SWAP test accepts}] \leq 1 - \epsilon^2/2.$$

We can also amplify the probability of detecting that $|\psi\rangle$ in the $S \neq C, \overline{C}$ case of [Corollary 4](#) by taking additional copies of $|\psi\rangle$. In particular, suppose we have $2m$ copies of $|\psi\rangle$ and we group them up into m pairs. If we run the SWAP test on the qubits in S for each pair and accept only if all n SWAP tests accept, then the probability we accept is at most

$$(1 - \epsilon^2/2)^m \leq e^{-\frac{1}{2}\epsilon^2 m}.$$

This gives us the following proposition.

Proposition 6. *Given an integer k , there is a projective measurement $\{\Pi_S, \overline{\Pi}_S\}$ which acts as follows. Let $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\overline{C}}$ in which $|a\rangle_C$ and $|b\rangle_{\overline{C}}$ are both ϵ -far from any multipartite product state. Suppose we measure $|\psi\rangle^{\otimes 2m}$ with $\{\Pi_S, \overline{\Pi}_S\}$. If $S = C$ or $S = \overline{C}$, this measurement always accepts. Otherwise, if $S \neq C, \overline{C}$, the probability it accepts is at most $\exp(-\epsilon^2 m/2)$.*

That the measurement is projective follows from the fact that it is performing m SWAP tests, and each SWAP test is a projective measurement due to [Proposition 3](#).

2.5 An information theoretic algorithm for the hidden cut problem

Harrow, Lin, and Montanaro [[HLM17](#)] studied the problem of testing whether a given n -qubit state $|\psi\rangle$ is a multipartite product state or is ϵ -far from all multipartite product states. The key intuition

is that if $|\psi\rangle$ is a multipartite product state, then there exists an S such that the measurement $\{\Pi_S, \bar{\Pi}_S\}$ accepts with probability 1. On the other hand, if $|\psi\rangle$ is ϵ -far from all multipartite product states, then the measurement will accept with probability at most $\exp(-\epsilon^2 m/2)$, which can be made smaller than, say, 100^{-n} by taking $m = O(n/\epsilon^2)$. Since this is so small, we can apply their quantum OR bound to combine all $(2^n - 2)$ different $\{\Pi_S, \bar{\Pi}_S\}$ measurements into a single measurement $\{Q, \bar{Q}\}$ which always accepts on multipartite product states and accepts with exponentially small probability on states which are ϵ -far from multipartite product.

We now observe that if instead of combining these measurements with the quantum OR bound, we combine them with Gao's quantum union bound, we get a sample-efficient algorithm for finding the cut C .

Fact 2.1 (Gao's quantum union bound [Gao15]). *Let ρ be a density matrix. For each $1 \leq i \leq k$, let $\{\Pi_i, \bar{\Pi}_i\}$ be a two outcome projective measurement, and write $\text{err}_i := \text{Tr}(\Pi_i \cdot \rho)$. If we measure ρ with each $\{\Pi_i, \bar{\Pi}_i\}$ measurement in sequence from $i = 1$ to k , then the probability that we only observe the $\bar{\Pi}_i$ outcomes is at least $1 - 4 \cdot (\text{err}_1 + \dots + \text{err}_k)$.*

Theorem 3 (Information theoretic algorithm for the hidden cut problem). *Let $|\psi\rangle = |a\rangle_C \otimes |b\rangle_{\bar{C}}$ in which $|a\rangle_C$ and $|b\rangle_{\bar{C}}$ are both ϵ -far from any multipartite product state. There is an algorithm which can identify C with probability at least 99% given $m = O(n/\epsilon^2)$ copies of $|\psi\rangle$.*

Proof. Set $m = O(n/\epsilon^2)$ so that $\exp(-\epsilon^2 m/2) \leq 100^{-n}$. Then for each subset S , Proposition 6 gives us a projective measurement $\{\Pi_S, \bar{\Pi}_S\}$ which accepts with probability 1 if $S = C, \bar{C}$ and with probability at most 100^{-n} if $S \neq C, \bar{C}$. Set $N = 2^n - 2$ and pick an arbitrary order S_1, \dots, S_N on the nontrivial subsets of $[n]$. The algorithm is as follows: given m copies of $|\psi\rangle$, perform the N measurements $\{\Pi_{S_1}, \bar{\Pi}_{S_1}\}$ through $\{\Pi_{S_N}, \bar{\Pi}_{S_N}\}$ in order until the first time observing a Π_{S_i} outcome; when this happens, output " S_i " and terminate.

Suppose without loss of generality that S_i is equal to the true cut C , and none of the previous S_j 's is equal to \bar{C} . The algorithm succeeds if the first $i - 1$ measurements reject and the i -th measurement accepts. To compute the probability that this does *not* happen, we can apply Gao's quantum union bound with $\Pi_j = \Pi_{S_j}$ for each $1 \leq j \leq i - 1$ and $\bar{\Pi}_i = \bar{\Pi}_{S_i}$. Then $\text{err}_j = \text{Tr}(\Pi_{S_j} \cdot \rho) \leq 100^{-n}$ for each $1 \leq j \leq i - 1$ and $\text{err}_i = \text{Tr}(\bar{\Pi}_{S_i} \cdot \rho) = 0$. Then the union bound says that the probability the algorithm does not succeed is at most

$$4 \cdot (\text{err}_1 + \dots + \text{err}_{i-1} + \text{err}_i) \leq 4 \cdot (i - 1) \cdot \frac{1}{100^n} \leq \frac{4N}{100^n} \leq 0.01.$$

This completes the proof. □

Note that the multipartite product state detection algorithm of Harrow, Lin, and Montanaro runs in exponential time, because it involves computing the quantum OR of an exponentially large number of measurements and then implementing that (likely computationally infeasible) measurement on $|\psi\rangle^{\otimes m}$. Similarly, this algorithm for the hidden cut problem also requires exponential time as it performs exponentially many measurements in sequence.

3 The State Hidden Subgroup Problem

To produce our algorithm we introduce a *quantum state* version of the hidden subgroup problem which may be of independent interest. Our definition is motivated by the observation that in the

hidden cut problem, the input states have certain symmetries determined by the secret cut. For example, the Haar measure is invariant under arbitrary unitaries. If we instantiate the hidden cut problem with two $n/2$ -qubit Haar random states separated by a random cut, this means that the input states to the hidden cut problem, when viewed as a density matrix, are invariant under the action of a group isomorphic to $U(2^{n/2}) \times U(2^{n/2})$ — the issue is that we don't know the qubit bipartition which defines the specific symmetry group.

This sounds related to the well-studied Hidden Subgroup Problem (HSP) [NC10, Section 5.4.3], in which one is given a function $f : G \rightarrow \{0, 1\}^n$ invariant on left cosets of a subgroup $H < G$, with the goal of learning H :

Definition 12 (Hidden subgroup problem (HSP)). *A function $f : G \rightarrow L$ from a finite group G to a set of discrete labels L is said to hide a subgroup $H \leq G$ if it is constant on the (left) cosets of H , and different across the cosets, in other words $f(x) = f(y)$ if and only if $x^{-1}y \in H$. The hidden subgroup problem is the task of determining the hidden subgroup H from as few queries to the function f as possible, assuming black-box access to an oracle implementation $O_f : |g\rangle |0\rangle \mapsto |g\rangle |f(g)\rangle$.*

One crucial difference, however, is that the HSP takes as input a *function* respecting certain subgroup symmetries, while the hidden cut problem takes as input *quantum states* invariant under a certain subgroup action. Motivated by this observation, we define a quantum state version of the HSP:

Definition 2 (The state hidden subgroup problem (StateHSP) — restated). *Let G be a finite group with a unitary representation $R : G \rightarrow U(d)$. Let \mathbb{C}^G denote a Hilbert space in the regular representation of G , meaning $\mathbb{C}^G = \text{span}\{|g\rangle\}_{g \in G}$ where $\langle g|h\rangle = \delta_{g,h}$. Assume efficient implementation of the controlled group action $U_R = \sum_{g \in G} |g\rangle\langle g| \otimes R(g)$ acting on $\mathbb{C}^G \otimes \mathbb{C}^d$. Assume access to (copies of) a quantum state $|\psi\rangle \in \mathbb{C}^d$ with the following properties:*

- $|\psi\rangle$ is invariant under the action of a subgroup H , i.e. for all $h \in H$, $R(h)|\psi\rangle = |\psi\rangle$.
- $|\psi\rangle$ is acted on nontrivially by elements outside the subgroup: for any $g \notin H$, $|\langle \psi | R(g) | \psi \rangle| \leq 1 - \epsilon$,

where the parameter ϵ , which we call the “orthogonality allowance”, can depend on the dimension d , the group G , and the representation R . The goal is to identify the hidden subgroup H .

Here we are assuming one has efficient access to the group representation — i.e. given $g \in G$, one can apply $R(g)$ efficiently. In this sense the problem is similar to the black box group model of computing (e.g. employed in [Wat00]), but defined with respect to a representation of the group other than the left regular representation. Our definition can also be viewed as inspired by recent works studying quantum state/unitary variants of complexity classes and cryptography such as [BEM⁺23, RY22, LMW24, Zha24]. We emphasize that the value of ϵ could vary substantially between different representations, so the scaling behavior of ϵ might significantly affect the difficulty of this problem. Additionally, this framework can accommodate problems in which either the parent group G , or the representation R and its dimension d , or possibly both, can depend on the specific underlying parameter of the problem. Multiple variations can be imagined, such as an additional promise that the hidden subgroups are mutually conjugate (a common HSP flavor), or introducing unknowns about the specific group representation R .

We note that a recent set of works has studied the problem of determining whether a given quantum state is preserved by a known, specific group action [LRW23, RLW23]. This particular property

testing task can be seen as a special case of decisional StateHSP, in which the role of the hidden subgroup H is played by the parent group G itself. Finally, we notice a connection to the problem of *state isomorphism* [LG17], which asks whether two input states can be made equal under a permutation of the qubits.

3.1 Coset states and the standard HSP approach

Despite the syntactic differences between the HSP and the StateHSP, there is a sense in which the HSP can be viewed as a special case of the StateHSP. The “standard method” for the HSP [GSVV04] involves preparing a uniform superposition over the elements of G and feeding it into the f oracle, resulting in the state

$$|\psi\rangle = \frac{1}{\sqrt{|G|}} \cdot \sum_{x \in G} |x\rangle_G \otimes |f(x)\rangle_L. \quad (7)$$

Next, one discards the label register, resulting in a uniform mixture of coset states

$$\frac{1}{|G|} \cdot \sum_{g \in G} |gH\rangle\langle gH|, \quad \text{where } |gH\rangle := \frac{1}{\sqrt{|H|}} \cdot \sum_{h \in H} |gh\rangle \text{ is a coset state.}$$

One can then repeatedly run this procedure to generate multiple coset states, and the task is to use them to learn H .

However, even before discarding the label register, the state in Equation (7) is already an instance of the StateHSP. In particular, let $R : G \rightarrow U(d)$ be the *right regular representation* of G , meaning that it acts on \mathbb{C}^G via $R(g) \cdot |x\rangle = |xg^{-1}\rangle$. Suppose f hides the subgroup H . Then $|\psi\rangle$ is invariant under the action of H , because for any $h \in H$,

$$\begin{aligned} R(h)_G \cdot |\psi\rangle &= \frac{1}{\sqrt{|G|}} \cdot \sum_{x \in G} |xh^{-1}\rangle_G \otimes |f(x)\rangle_L \\ &= \frac{1}{\sqrt{|G|}} \cdot \sum_{y \in G} |y\rangle_G \otimes |f(yh)\rangle_L = \frac{1}{\sqrt{|G|}} \cdot \sum_{y \in G} |y\rangle_G \otimes |f(y)\rangle_L = |\psi\rangle, \end{aligned}$$

where we used the fact that $f(y) = f(yh)$ because $y^{-1}yh = h \in H$. On the other hand, for any $g \notin H$,

$$R(g)_G \cdot |\psi\rangle = \frac{1}{\sqrt{|G|}} \cdot \sum_{x \in G} |xg^{-1}\rangle_G \otimes |f(x)\rangle_L = \frac{1}{\sqrt{|G|}} \cdot \sum_{y \in G} |y\rangle_G \otimes |f(yg)\rangle_L.$$

But $f(y) \neq f(yg)$ because $y^{-1}yg = g \notin H$. This means that for all $g \notin H$,

$$\langle \psi | R(g)_G | \psi \rangle = 0,$$

and so this state satisfies the definition of the StateHSP with an orthogonality allowance of $\epsilon = 1$. In general, we will see that instances of the StateHSP where ϵ is close to 1 act like instances of the traditional HSP, which we can sometimes solve efficiently.

3.2 Fourier sampling in HSP vs. StateHSP

We begin by reviewing the Fourier sampling approach to solving the standard hidden subgroup problem, which we will proceed to generalize to the StateHSP setting. The construction at its core is the group Fourier transform:

Definition 13 (Group Fourier transform [Dia88]). *Let G be a finite group (not necessarily Abelian), and let $\{\rho_\lambda\}_{\lambda \in \text{Irr}[G]}$ be a full set of irreducible unitary G -representations (irreps), such that each $\rho_\lambda : G \rightarrow \text{U}(d_\lambda)$ is a unitary irrep of G of dimension d_λ . Then, the group Fourier transform is the $|G| \times |G|$ unitary \mathcal{F}_G which transforms from the regular representation basis $\{|g\rangle\}_{g \in G}$ to a Schur basis $\{|\lambda, i, j\rangle\}_{\substack{\lambda \in \text{Irr}[G] \\ i, j \in [d_\lambda]}}$. Explicitly:*

$$\mathcal{F}_G |g\rangle = \sum_{\substack{\lambda \in \text{Irr}[G] \\ i, j \in [d_\lambda]}} \sqrt{\frac{d_\lambda}{|G|}} \rho_\lambda(g)_{i,j} |\lambda, i, j\rangle, \quad \mathcal{F}_G^{-1} |\lambda, i, j\rangle = \sqrt{\frac{d_\lambda}{|G|}} \sum_{g \in G} \rho_\lambda(g^{-1})_{j,i} |g\rangle. \quad (8)$$

The Fourier sampling approach to the generic HSP becomes possible when there is an efficient circuit for the group Fourier transform. This is usually a safe assumption if the group is Abelian; efficient circuits for the non-Abelian quantum Fourier transform are known for several important groups including the dihedral and symmetric groups, but other groups are conjectured not to admit efficient QFT circuits [MRR06]. Given an HSP with hidden subgroup $H < G$, the measurement outcome of a so-called “weak” Fourier sampling circuit are samples from a distribution over the irrep labels $\lambda \in \text{Irr}[G]$:

Fact 3.1 (Weak Fourier sampling in HSP [HRTS03]). *In an HSP over the parent group G with a hidden subgroup $H < G$, the weak Fourier sampling outcome distribution over the labels $\lambda \in \text{Irr}[G]$ is given by:*

$$\text{P}_{\text{HSP}}[\lambda] = \frac{d_\lambda}{|G|} \sum_{h \in H} \chi_\lambda(h), \quad (9)$$

where χ_λ denotes the corresponding irreducible character of G , i.e. $\chi_\lambda(g) = \text{Tr} \rho_\lambda(g)$.

Proof. In the HSP weak Fourier sampling setting, one starts with an arbitrary *coset state* in the regular representation (which can be efficiently prepared from oracle access to the input function):

$$|gH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |gh\rangle, \quad (10)$$

to which the G -Fourier transform \mathcal{F}_G (8) is applied, leading to the state:

$$\mathcal{F}_G |gH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} \sum_{\substack{\lambda \in \text{Irr}[G] \\ i, j \in [d_\lambda]}} \sqrt{\frac{d_\lambda}{|G|}} \rho_\lambda(gh)_{i,j} |\lambda, i, j\rangle. \quad (11)$$

Finally, only the irrep label register is measured, leading to an output probability:

$$\text{P}_{\text{HSP}}[\lambda] = \frac{d_\lambda}{|G||H|} \sum_{\substack{h, h' \in H \\ i, j \in [d_\lambda]}} \rho_\lambda(gh)_{i,j} \overline{\rho_\lambda(gh')_{i,j}} \quad (\text{via Born rule}) \quad (12)$$

$$= \frac{d_\lambda}{|G||H|} \sum_{\substack{h, h' \in H \\ i, j \in [d_\lambda]}} \rho_\lambda(gh)_{i,j} \rho_\lambda(h'^{-1}g^{-1})_{j,i} \quad (\text{since } \rho_\lambda(gh') \text{ is unitary}) \quad (13)$$

$$= \frac{d_\lambda}{|G||H|} \sum_{h, h' \in H} \text{Tr} \rho_\lambda(hh'^{-1}) \quad (\text{by cyclic property of trace}) \quad (14)$$

$$= \frac{d_\lambda}{|G|} \sum_{h \in H} \chi_\lambda(h) \quad (\text{by double-summing over the subgroup}). \quad (15)$$

See [HRTS03] for more detail. \square

A natural question is understanding how this distribution changes when we generalize Fourier sampling to a StateHSP problem — specifically, by applying the circuit in Figure 4b:

Fact 3.2 (Weak Fourier sampling in StateHSP). *The output of the weak Fourier sampling circuit in Figure 4b for a StateHSP problem in which the input state $|\psi\rangle$ is invariant under the R -action of the hidden subgroup $H < G$ is:*

$$P_{\text{StateHSP}_\psi}[\lambda] = \frac{d_\lambda}{|G|} \sum_{c \in G/H} \sum_{h \in H} \chi_\lambda(ch) \langle \psi | R(c) | \psi \rangle, \quad (16)$$

where $\lambda \in \text{Irr}[G]$ is a G -irrep label, and G/H denotes a set of left coset representatives.

Proof. The result follows from direct calculation along similar lines as Fact 3.1. In the StateHSP setting, there is an ancillary register admitting a regular representation of the group G (initialized in the trivial representation), together with the input state $|\psi\rangle$. The first application of the inverse G -Fourier transform prepares a uniform superposition over group elements in the first register:

$$(\mathcal{F}_G^{-1} \otimes I) |\mathbf{0}\rangle |\psi\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |\psi\rangle. \quad (17)$$

Applying the controlled group action $U_G = \sum_{g \in G} |g\rangle\langle g| \otimes R(g)$ leads to the state:

$$U_G (\mathcal{F}_G^{-1} \otimes I) |\mathbf{0}\rangle |\psi\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle \otimes R(g) |\psi\rangle. \quad (18)$$

Finally, the last application of the G -Fourier transform results in the state:

$$(\mathcal{F}_G \otimes I) U_G (\mathcal{F}_G^{-1} \otimes I) |\mathbf{0}\rangle |\psi\rangle = \frac{1}{|G|} \sum_{\substack{\lambda \in \text{Irr}[G] \\ i, j \in [d_\lambda]}} \sqrt{d_\lambda} |\lambda, i, j\rangle \otimes \sum_{g \in G} \rho_\lambda(g)_{i, j} R(g) |\psi\rangle. \quad (19)$$

On this state, we measure the ancillary register corresponding to the irrep label $\lambda \in \text{Irr}[G]$, leading to the output probability:

$$\begin{aligned} P_{\text{StateHSP}_\psi}[\lambda] &= \frac{d_\lambda}{|G|^2} \sum_{\substack{i, j \in [d_\lambda] \\ g, g' \in G}} \rho_\lambda(g)_{i, j} \overline{\rho_\lambda(g')_{i, j}} \langle \psi | R(g')^\dagger R(g) | \psi \rangle && (\text{via Born rule}) \\ &= \frac{d_\lambda}{|G|^2} \sum_{g, g' \in G} \text{Tr} \rho_\lambda(g'^{-1}g) \langle \psi | R(g'^{-1}g) | \psi \rangle && (\text{by unitarity of } \rho_\lambda \text{ and } R) \\ &= \frac{d_\lambda}{|G|} \sum_{g \in G} \chi_\lambda(g) \langle \psi | R(g) | \psi \rangle && (\text{after double-summing over the group } G). \end{aligned}$$

Splitting the group elements over the right- H -cosets as $g = ch$, where $h \in H$ are subgroup elements

and $c \in G/H$ are coset representatives, we obtain:

$$\begin{aligned}
P_{\text{StateHSP}_\psi}[\lambda] &= \frac{d_\lambda}{|G|} \sum_{\substack{c \in G/H \\ h \in H}} \chi_\lambda(ch) \langle \psi | R(ch) | \psi \rangle \\
&= \frac{d_\lambda}{|G|} \sum_{\substack{c \in G/H \\ h \in H}} \chi_\lambda(ch) \langle \psi | R(c) R(h) | \psi \rangle \quad (\text{since } R \text{ is a } G\text{-representation}) \\
&= \frac{d_\lambda}{|G|} \sum_{\substack{c \in G/H \\ h \in H}} \chi_\lambda(ch) \langle \psi | R(c) | \psi \rangle \quad (\text{by the } H\text{-invariance of } |\psi\rangle),
\end{aligned}$$

where in the last line we used the subgroup invariance assumption about the input state: $R(h) |\psi\rangle = |\psi\rangle$ for all $h \in H$. \square

Notice that the StateHSP outcome distribution (16) and the equivalent HSP outcome distribution (9) are identical when $\langle \psi | R(c) | \psi \rangle = 0$ for all nontrivial coset representatives $c \neq \text{id}$. The analogy is explained by the fact that in StateHSP, the states $\{R(c) |\psi\rangle\}_{c \in G/H}$ play a similar role to the *coset states* $|cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle$ in HSP. The coset states $\{|cH\rangle\}_{c \in G/H}$ are manifestly H -invariant under a right-regular group action, and also mutually orthogonal. Therefore, the usual approach to HSP involving the construction of coset states via the function oracle is a specific instance of StateHSP. However, the more generic StateHSP problem differs in that it allows non-orthogonal coset states $\{R(c) |\psi\rangle\}_{c \in G/H}$. On the other hand, if the coset states are too close to each other, the problem risks becoming information-theoretically intractable: the state would be too close to the symmetric subspace invariant under all group operations, and distinguishing between any nontrivial subgroup symmetry and the full group symmetry would require an inefficient number of measurements. For this reason, it is crucial to introduce an orthogonality allowance ϵ in [Definition 2](#).

A simple but powerful observation is that a large enough orthogonality can be further amplified when one is allowed multiple copies of the input state:

Fact 3.3. *There is a Fourier sampling circuit for an StateHSP problem with orthogonality allowance ϵ which uses t copies of the input state, yielding outcome distribution:*

$$P_{\text{StateHSP}_{\psi,t}}[\lambda] = P_{\text{HSP}}[\lambda] + O(d_\lambda^2 \exp(-\epsilon t)). \quad (20)$$

The depth of the circuit can be as shallow as $O(\log t)$ by using $O(t \log |G|)$ ancillary qubits.

Proof. The construction is natural and it involves promoting the group action $\sum_{g \in G} |g\rangle\langle g| \otimes R(g)$ on $\mathbb{C}^G \otimes \mathbb{C}^d$ to the t -fold version $\sum_{g \in G} |g\rangle\langle g| \otimes R(g)^{\otimes t}$ on $\mathbb{C}^G \otimes (\mathbb{C}^d)^{\otimes t}$, and plugging this action into the StateHSP Fourier sampling circuit of [Fact 3.2](#) with input state $|\psi\rangle^{\otimes t}$. A simple application of triangle inequality gives:

$$|P_{\text{StateHSP}_{\psi,t}}[\lambda] - P_{\text{HSP}}[\lambda]| = \left| \frac{d_\lambda}{|G|} \sum_{\substack{c \in G/H \neq \text{id} \\ h \in H}} \chi_\lambda(ch) \langle \psi^{\otimes t} | R(c)^{\otimes t} | \psi^{\otimes t} \rangle \right| \quad (21)$$

$$\leq \frac{d_\lambda}{|G|} \sum_{\substack{c \in G/H \neq \text{id} \\ h \in H}} |\chi_\lambda(ch)| |\langle \psi | R(c) | \psi \rangle|^t \quad (22)$$

$$\leq d_\lambda^2 \left(1 - \frac{|H|}{|G|}\right) (1 - \epsilon)^t \leq O(d_\lambda^2 \exp(-\epsilon t)), \quad (23)$$

where in the last line we used the simple character bound $|\chi_\lambda(g)| \leq d_\lambda$ for all $g \in G$. The circuit can be implemented by t successive applications of the single group action, once per each copy of the input Hilbert space. To apply this action in depth $O(\log t)$, we make use of t ancillary registers which host regular representations of G . One can copy the group information from the first regular representation register onto all of these additional registers by a $O(\log t)$ -depth binary tree of controlled two-register unitaries; then, each of these t ancillary registers can control the group action on the t input state copies in parallel. Finally, one uncomputes the binary tree of two-register unitaries in depth $O(\log t)$. \square

We note that this simple bound uses no information about the specific group G ; it is often the case that the characters decay rapidly in magnitude from the maximum value $\chi_\lambda(\text{id}) = d_\lambda$ across the group, so even tighter bounds might be possible. Similarly, improvements can be obtained if the specific StateHSP problem presents additional information about the inner products of the coset states³.

The key takeaway is that by increasing the number of copies t , we can naturally enhance the orthogonality of coset states and make the StateHSP instance behave like the equivalent HSP problem from the point of view of Fourier sampling. The number of copies required to make the non-orthogonality correction negligible in this way will depend on the specifics of the problem: for a problem parameter n , as long as $\epsilon \geq \Omega(1/\text{poly}(n))$ and⁴ $|G| \leq O(\exp(\text{poly}(n)))$, then there is a choice of $t = \text{poly}(n)$ which would ensure the corrections are negligibly small in n .

As we will show in the next section, this framework applies to the hidden cut problem, and we will be able to use an efficient number of copies to enhance orthogonality such that a standard HSP can be applied to the hidden cut problem via Fourier sampling. We leave it as an open question for future work to find meaningful instances of StateHSP which cannot be efficiently amplified by a polynomial number of copies for purposes of Fourier sampling.

We end this section by mentioning an immediate corollary which follow from the above connection between HSP and StateHSP when multiple state copies are available. Namely, the general non-Abelian StateHSP is information-theoretically solvable using few copies of the state, just as the standard HSP is information-theoretically solvable with few queries to the function [EHK04].

Corollary 5. *The general non-Abelian HSP over a group G with can be information-theoretically determined with a number of copies $O(\text{poly}(\log |G|, \epsilon^{-1}))$ of the input state $|\psi\rangle$.*

Of course, whether computationally efficient algorithms exist for non-Abelian groups is an open problem, and the non-existence of such algorithms underlies hardness of post-quantum cryptographic schemes [Reg04]. This corollary follows immediately from the information-theoretic feasibility of HSP [EHK04], combined with Fact 3.3 outlined above. Therefore, we know a StateHSP is information-theoretically solvable with enough copies and a large enough orthogonality allowance; the question is when it is computationally efficiently solvable.

³This comment captures the different behavior of our algorithm for the hidden cut problem with a constant entanglement promise versus a Haar-random promise, as will be described in later sections.

⁴This is because irrep dimensions cannot be larger than $d_\lambda \leq \sqrt{|G|}$.

4 An algorithm for the hidden cut problem

In the hidden cut problem, the global state is a tn -qubit state of the form $|\psi\rangle^{\otimes t}$, where $|\psi\rangle \in (\mathbb{C}^2)^n$ is separable across an unknown cut $C \subset [n]$, denoted $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$, where $\bar{C} \equiv [n] \setminus C$. Aligning the global state in an imaginary qubit grid with t rows and n columns (see [Figure 2](#)), let us define the relevant column permutation operations:

Definition 14 (Permutation operations). *The state $|\psi\rangle^{\otimes t} \in (\mathbb{C}^{2^n})^{\otimes t} = \mathbb{C}^{2^{tn}}$ is a state on $t \times n$ qubits. Let the standard basis of this space be of the form $\bigotimes_{i \in [t], j \in [n]} |e_{i,j}\rangle$. Let us define the action of \mathbb{S}_t on the k -th column of t qubits as the operators $R_k(\pi)$:*

$$\text{for } \pi \in \mathbb{S}_t, k \in [n]: \quad R_k(\pi) \bigotimes_{i \in [t], j \in [n]} |e_{i,j}\rangle \equiv \bigotimes_{i \in [t], j \in [n] \setminus \{k\}} |e_{i,j}\rangle \bigotimes_{i \in [t]} |e_{\pi^{-1}(i),k}\rangle. \quad (24)$$

In what follows, we will choose appropriate group actions on the global state $|\psi\rangle^{\otimes t}$ expressed by groups of ‘column-wise’ permutation operators $\{R_k(\pi)\}_{k \in [n], \pi \in \mathbb{S}_t}$, i.e. the permutation of the k ’th qubit across the t copies.

4.1 The full permutational symmetries defy Fourier sampling

In order to apply the StateHSP framework from [Section 3](#) to the hidden cut problem, the key first step is to choose an appropriate group action on the state. A first natural choice is to take advantage of all the manifest permutational symmetries of the global state $|\psi\rangle^{\otimes t}$ containing the t copies of the input state. Such a t -fold state is trivially symmetric under the \mathbb{S}_t group which permutes the copies. When additionally the input state is separable across a cut $C \subset [n]$, then the t -fold global state $|\psi\rangle^{\otimes t} = |\phi_1\rangle_C^{\otimes t} \otimes |\phi_2\rangle_{\bar{C}}^{\otimes t}$ has a larger $\mathbb{S}_t \times \mathbb{S}_t$ permutational symmetry group which acts by permuting the individual factor state copies (see [Figure 2](#)). The natural parent group which accommodates all of these cut-specific symmetry subgroups is $G = \mathbb{S}_t^{\times n}$ acting as single-column permutations. In the language of StateHSP, the corresponding group action is $R : (\sigma_1, \dots, \sigma_n) \mapsto R_1(\sigma_1) \otimes \dots \otimes R_n(\sigma_n)$, and the hidden cut subgroup $H_C < G$ preserving the state under this action is⁵ $H_C = \{\sigma^C \mu^{\bar{C}}\}_{\sigma, \mu \in \mathbb{S}_t}$.

In [Section 3](#) we imported the Fourier sampling approach from HSP as a possible algorithm to solve StateHSP. We now briefly outline an obstacle to applying Fourier sampling to find the hidden cut with the permutation group action introduced above. While the Fourier sampling circuit can be implemented efficiently,⁶ the difficulty comes from the information-theoretic properties of the equivalent standard HSP:

Fact 4.1. *Assume $t \geq \Omega(\text{poly}(n))$. Given a cut $C \subset [n]$ which determines a symmetry subgroup H_C isomorphic to $\mathbb{S}_t \times \mathbb{S}_t$ inside $\mathbb{S}_t^{\times n}$ as defined above, then performing non-Abelian weak Fourier sampling over $\mathbb{S}_t^{\times n}$ gives the output probabilities:*

$$P_{\text{HSP}}[\lambda_1, \dots, \lambda_n] = \frac{d_{\lambda_1}^2 \dots d_{\lambda_n}^2}{t!^n} (1 + O(t^2 b^n)) \quad \text{for } \frac{\lambda_{i,1}}{t}, \frac{\lambda'_{i,1}}{t} < o(1), i \in [n], \quad (25)$$

for some constant $0 < b < 1$, with a negligible probability mass outside of this regime. Here, λ_i are irreducible representations of \mathbb{S}_t of dimensions d_{λ_i} , and $\lambda_{i,1}, \lambda'_{i,1}$ are the lengths of the first row and first column of the Young diagram λ_i . As a consequence, the probability of observing irreps

⁵As elsewhere in this paper, the notation $\sigma^C \mu^{\bar{C}}$ means the vector in $\mathbb{S}_t^{\times n}$ with σ in the C positions and μ elsewhere.

⁶This is because of the known efficient quantum Fourier transform constructions for \mathbb{S}_t [[Bea97](#)].

λ_i outside of the range $\frac{\lambda_{i,1}}{t}, \frac{\lambda'_{i,1}}{t} < o(1)$ is negligibly small. Inside the typical observable range, the result means that all cuts result in the same “Plancherel distribution” of outcomes to within negligible relative corrections $O(t^2 b^n)$.

Proof sketch. We merely outline the argument, which relies on technical aspects of the representation theory of the symmetric group; a similar proof is detailed in [MRS08] to show that Fourier sampling cannot solve the generic HSP for the symmetric group. The goal is to estimate the benchmark HSP outcome distribution (9). Let $\text{cyc}(\sigma)$ denote the number of cycles in the permutation σ . The starting observation is that a subgroup element $\sigma_1^C \sigma_2^{\bar{C}}$ has a number of cycles equal to $n \text{cyc}(\sigma_1)/2 + n \text{cyc}(\sigma_2)/2$ as a member of $\mathbb{S}_t^{\times n}$. Using Roichman’s bounds on the characters of the symmetric group [Roi96], this gives us that there exists some $0 < b < 1$ such that:

$$\left| \chi_{\lambda_1, \dots, \lambda_n}(\sigma_1^C \sigma_2^{\bar{C}}) \right| \leq d_{\lambda_1} \dots d_{\lambda_n} b^{n(t - \text{cyc}(\sigma_1) - \text{cyc}(\sigma_2))/2}. \quad (26)$$

This holds true as long as the Young diagrams $\lambda_1, \dots, \lambda_n \vdash t$ have a first row or column shorter than a $o(1)$ fraction of t . This condition is satisfied with high probability due to arguments such as [BDJ99] about the typical Young diagrams concentrating towards $\Theta(\sqrt{t})$ rows and columns. Therefore, the chance of ever seeing any diagram outside the scope of this typical regime is negligibly small when $t \geq \Omega(\text{poly}(n))$. Within the typical regime, the bound above is enough to control the non-identity terms in the sum over the subgroup H in (9), which results in the claim. \square

The message of the above claim is that the corresponding HSP problem becomes information-theoretically harder with increasing number of copies t , which is counter-productive if we expect to use the number of copies t as a method of signal amplification. Additionally, it is unclear whether in the low- t regime Fourier samples can be efficiently analyzed to detect the cut. This choice of group action has the disadvantage that the same parameter t defines both the accuracy of the HSP approximation via orthogonality amplification, as well as the complexity of the HSP problem. In the next section, we find a much simpler permutation action such that the parent group depends on the state size n but not on t , with the added benefit of the group being Abelian.

4.2 A first Abelian HSP algorithm for the hidden cut problem

It turns out we can restrict the full permutation group $\mathbb{S}_t^{\times n}$ to a smaller group which still supports the mapping of cuts to subgroups, but whose size grows only with the number of qubits n , but not with the number of copies t . Crucially, the group is the simple Abelian group $G = \mathbb{Z}_2^n$; since HSP is known to allow efficient Fourier sampling algorithms in the Abelian case, this opens up the possibility of an efficient algorithm for the hidden cut problem by the technique in Fact 3.3.

Assume t is even and define the following permutation in \mathbb{S}_t :

$$\xi \equiv (1 \ 2)(3 \ 4) \dots (t-1 \ t). \quad (27)$$

Let us allow $G = \mathbb{Z}_2^n$ to act on the space of $t \times n$ qubits as:

$$(x_1, \dots, x_n) \circ |\psi\rangle^{\otimes t} \equiv R_1(\xi)^{x_1} R_2(\xi)^{x_2} \dots R_n(\xi)^{x_n} |\psi\rangle^{\otimes t}. \quad (28)$$

This is a well-defined \mathbb{Z}_2^n group action since the $R_i(\xi)$ operators mutually commute (Abelian) and $R_i(\xi)^2 = I$ (order two since $\xi^2 = \text{id}$). Unless specified otherwise, we can simply denote $R_k \equiv R_k(\xi)$ and $R(\mathbf{x}) \equiv R_1^{x_1} \dots R_n^{x_n}$ to simplify notation from here onwards. See Figure 3 for an illustration.

A first algorithm to find the hidden cut $C \subset [n]$ can be laid out as follows:

Algorithm 1: Non-adaptive hidden cut algorithm

- Requirements:** n additional qubits, implementation of the \mathbb{Z}_2^n group action $U_{\mathbb{Z}_2^n}$.
- 1 **if** the factor states are promised to be ϵ -far from separable ([Theorem 1](#)):
 - 2 Let $t = O(n/\epsilon^2)$.
 - 3 **if** the factor states are promised to be Haar-random ([Theorem 2](#)):
 - 4 Let $t = 2$.
 - 5 **for** sample count $k \in \{1, \dots, p\}$, where $p = O(n)$:
 - 6 Prepare t copies of the state $|\psi\rangle$.
 - 7 Run the Fourier sampling circuit ([Figure 4c](#)): $(H^{\otimes n} \otimes I) U_{\mathbb{Z}_2^n} (H^{\otimes n} \otimes I) |0^n\rangle \otimes |\psi\rangle^{\otimes t}$.
 - 8 Measure the group register to get sample $\mathbf{y}^{(k)} \in \mathbb{Z}_2^n$.
 - 9 Classically solve for the nullspace of $Y = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)})^T \in \mathbb{Z}_2^{p \times n}$ which is $\text{span}\{1^C 0^{\bar{C}}, 0^C 1^{\bar{C}}\}$.
-

The above algorithm succeeds in efficiently finding the hidden cut:

Theorem 4 (Non-adaptive hidden cut algorithm). *For $\epsilon > 0$ and an n -qubit input state $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$ separable across a cut $C \in \binom{[n]}{n/2}$, assume that the factor states $|\phi_{1,2}\rangle$ are at least ϵ -far from all separable $(n/2)$ -qubit states. Then, [Algorithm 1](#) succeeds in finding the hidden cut C with high probability using $O(n^2/\epsilon^2)$ copies of the input state $|\psi\rangle$. The algorithm requires coherent access to $O(n/\epsilon^2)$ copies at a time, on which it acts with circuits of depth $O(\log(n/\epsilon^2))$ given $O(n^2/\epsilon^2)$ ancillary qubits.*

As outlined above, [Algorithm 1](#) addresses both the generic hidden cut problem, which promises that the factor states are ϵ -far from separable, as well as the case of Haar-random factor states. The difference between the two is in the choice of the number of state copies t used to produce each Fourier sample. This section focuses on the first, generic case. In the later [Section 5](#), we will prove why the stronger promise of Haar-random factor states improves the requirements of [Algorithm 1](#), such that useful Fourier samples can be produced by constant-depth circuits acting on only two state copies at a time (see [Theorem 2](#)).

As stated in [Theorem 4](#), the above [Algorithm 1](#) succeeds in finding the cut with $O(n^2/\epsilon^2)$ total state copies. Crucially, at the end of this section we will introduce an adaptive modification of [Algorithm 1](#), which will allow us to find the hidden cut with only $O(n/\epsilon^2)$ state copies, thus further reducing the requirement by a factor of n . This second, adaptive hidden cut algorithm will be described in [Section 4.4](#) and will build upon the analysis of [Algorithm 1](#), which follows below.

4.3 Analyzing the hidden cut algorithm: proof of [Theorem 4](#)

To show that the algorithm succeeds in finding the cut with high probability, we organize the analysis across the following facts. The first two facts are straightforward:

Fact 4.2 (Circuit size). *With n ancillary qubits to represent the group \mathbb{Z}_2^n , the Fourier sampling circuit for the \mathbb{Z}_2^n action defined above on the t -fold input state $|\psi\rangle^{\otimes t} \in (\mathbb{C}^{2^n})^{\otimes t}$ can be implemented efficiently with a circuit of depth $O(\log t)$ involving $nt/2$ ancillary qubits.*

Proof. The circuit is sketched in [Figure 4c](#). The group quantum Fourier transform over \mathbb{Z}_2^n is simply n parallel Hadamard gates, while the controlled group action $U_{\mathbb{Z}_2^n} = \sum_{\mathbf{x} \in \mathbb{Z}_2^n} |\mathbf{x}\rangle\langle\mathbf{x}| \otimes R(\mathbf{x})$ is efficient to implement as n parallel sequences of $t/2$ controlled-SWAPs. With a single ancillary

register of n qubits, this can be implemented in depth $O(t)$, by having each of the n ancillary qubits control a sequence of $t/2$ SWAPs in parallel. The depth can be lowered to $O(\log t)$ by a fan-out construction at the expense of using $t/2$ ancillary n -qubit registers; one copies the information from the first ancillary register onto all t registers by a $O(\log t)$ -depth binary tree of CNOT gates, after which each of the $nt/2$ ancillary qubits controls a SWAP in parallel, followed by uncomputing the CNOT fan-out in depth $O(\log t)$. \square

Fact 4.3 (Hidden cut as Abelian StateHSP). *To each cut $C \subset [n]$ corresponds a hidden subgroup $H_C < \mathbb{Z}_2^n$ isomorphic to \mathbb{Z}_2^2 which preserves the state under the group action, given by:*

$$H_C = \{0^n, 1^C 0^{\bar{C}}, 0^C 1^{\bar{C}}, 1^n\}, \quad (29)$$

where by $1^C 0^{\bar{C}}$ we mean the n -bit string with 1's on the positions in C and 0's everywhere else. Similarly, denote by $\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}$ the n -bit string whose restriction to the positions in C is the sub-string $\mathbf{y}_1 \in \mathbb{Z}_2^{|C|}$, and whose restriction to the positions in $\bar{C} = [n] \setminus C$ is the sub-string $\mathbf{y}_2 \in \mathbb{Z}_2^{n-|C|}$. Then, performing Fourier sampling on the standard HSP over \mathbb{Z}_2^n with the hidden subgroup H_C produces n -bit string samples from the probability distribution:

$$P_{\text{HSP}}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] = \frac{\delta_{|\mathbf{y}_1| \text{ even}} \delta_{|\mathbf{y}_2| \text{ even}}}{2^{n-2}} = \begin{cases} 2^{-n+2} & \text{if } \mathbf{y}_1 \cdot 1^{n/2} = \mathbf{y}_2 \cdot 1^{n/2} = 0 \pmod{2} \\ 0 & \text{otherwise.} \end{cases} \quad (30)$$

This distribution is uniform over the $(n-2)$ -dimensional Boolean subspace in \mathbb{Z}_2^n defined by the cut C :

$$H_C^\perp \equiv \left\{ \mathbf{y} \in \mathbb{Z}_2^n : \mathbf{y} \cdot 1^C 0^{\bar{C}} = \mathbf{y} \cdot 0^C 1^{\bar{C}} = 0 \pmod{2} \right\}. \quad (31)$$

This formalizes the hidden cut problem as an instance of StateHSP, and provides the benchmark Fourier sampling probability of the equivalent standard HSP. The proof of the above fact is straightforward, given the manifest permutational symmetries of the t -fold input state (see Figure 3). We remark that the resulting HSP is a minor variation of the well-known Simon's algorithm [Sim97], which provides uniform samples from the $(n-1)$ -dimensional subspace of \mathbb{Z}_2^n orthogonal to the secret string \mathbf{s} : $H_s^\perp \equiv \{\mathbf{y} \in \mathbb{Z}_2^n : \mathbf{y} \cdot \mathbf{s} = 0\}$. Just like in Simon's algorithm, learning the orthogonal subspace based on samples from this distribution is the same as learning the secret, which in this case means learning the n -bit string $1^C 0^{\bar{C}}$ which describes the hidden cut C :

Fact 4.4 (Linear system). *A complete spanning set $(\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)})$ for the cut subspace H_C^\perp can be obtained with $p = O(n)$ independent samples from the HSP distribution (30) with high probability. Given such a spanning set, the string encoding the cut $1^C 0^{\bar{C}}$ or its equivalent mirror opposite $0^C 1^{\bar{C}}$ can be determined by solving for the nullspace of the matrix $Y = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)})^T$ with the samples as rows, which can be done in $\text{poly}(n)$ time, e.g. by Gaussian elimination.*

Proof. The proof of this fact is straightforward and follows the same logic as Simon's algorithm. As an alternative to solving for the nullspace of the matrix Y , we mention here a slightly slower, but more illustrative equivalent procedure of analyzing the collected samples. The idea is to iteratively learn the members of each side of the cut by solving a number of $n-1$ linear equations involving the matrix Y , in the following way: first, ask whether positions 1 and 2 (out of n) are on the same side of the cut (i.e. whether they are both in C or both in $\bar{C} = [n] \setminus C$). This is answered by solving for \mathbf{x} in the linear system $Y\mathbf{x} = (1, 1, 0, \dots, 0)^T$; there is a solution \mathbf{x} if 1 and 2 are on the same

side, otherwise the system is infeasible. Continue in this way for each of the remaining positions $3, \dots, n$ by solving the same pairwise membership check of each position against position 1, thus determining the cut allocation of all coordinates in $\text{poly}(n)$ time. \square

To understand when this Simon-like HSP algorithm can be applied to the Fourier samples coming from the associated StateHSP hidden cut problem, we need to apply the framework from [Section 3](#) to bound the difference between the two distributions. This will inform the number of copies necessary for orthogonality amplification such that the two probability distributions become negligibly close at the level of each outcome.

Fact 4.5 (Output distribution over \mathbb{Z}_2^n). *Assume $C \in \binom{[n]}{n/2}$ is the true cut and the factorization of the input state is $|\psi\rangle = |\phi_1\rangle^C \otimes |\phi_2\rangle^{\bar{C}}$. Given t state copies for each sample, [Algorithm 1](#) returns strings in $\{0, 1\}^n$ according to the probability distribution $\text{P}_{\text{StateHSP}_{\psi,t}}$, which respects:*

$$\left| \text{P}_{\text{StateHSP}_{\psi,t}}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] - \text{P}_{\text{HSP}}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] \right| \leq \frac{1}{2^{n-2}} \left[\prod_{k \in \{1,2\}} (1 + \Delta_{\phi_k,t}) - 1 \right], \quad (32)$$

where:

$$\Delta_{\phi,t} \equiv \sum_{\substack{S \subset [n/2] \\ 1 \notin S \\ S \neq \emptyset}} \text{Tr}[\phi_S^2]^{t/2}. \quad (33)$$

Proof. The starting point is the observation that powers of the purity enter naturally as the inner products of the coset states from [Section 3](#):

$$\langle \psi^{\otimes t} | R(\mathbf{x}) | \psi^{\otimes t} \rangle = \text{Tr}[\psi_{\mathbf{x}}^2]^{t/2}, \quad (34)$$

where the $\text{Tr}[\psi_{\mathbf{x}}^2]$ denotes the purity across the cut $C_{\mathbf{x}} \equiv \{i \in [n] : x_i = 1\}$ represented by the bit-string $\mathbf{x} \in \mathbb{Z}_2^n$. Importing the StateHSP analysis ([16](#)) from [Section 3](#), it follows that our Fourier sampling circuit effectively performs a Boolean Fourier transform on this set of amplified purities. Specifically, the output distribution is:

$$\text{P}_{\text{StateHSP}_{\psi,t}}[\mathbf{y}] = \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} (-1)^{\mathbf{x} \cdot \mathbf{y}} \text{Tr}[\psi_{\mathbf{x}}^2]^{t/2}, \quad \text{where } \mathbf{y} \in \mathbb{Z}_2^n. \quad (35)$$

Operations in the hidden subgroup H_C ([29](#)) preserve the state, so we can split the above sum over the group \mathbb{Z}_2^n into a sum over the subgroup H_C and a sum over the coset representatives \mathbb{Z}_2^n/H_C :

$$\text{P}_{\text{StateHSP}_{\psi,t}}[\mathbf{y}] = \frac{1}{2^n} \sum_{\mathbf{h} \in H_C} (-1)^{\mathbf{h} \cdot \mathbf{y}} \times \sum_{\mathbf{x} \in \mathbb{Z}_2^n/H_C} (-1)^{\mathbf{x} \cdot \mathbf{y}} \text{Tr}[\psi_{\mathbf{x}}^2]^{t/2}. \quad (36)$$

The first term is precisely the $\text{P}_{\text{HSP}}[\mathbf{y}]$ distribution corresponding to the standard HSP problem with the same specifications. Given the internal structure of the state, the purity factors into the two separate contributions from each substate:

$$\text{P}_{\text{StateHSP}_{\psi,t}}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] = \text{P}_{\text{HSP}}[\mathbf{y}] \times \prod_{k \in \{1,2\}} \left[\sum_{\substack{\mathbf{z} \in \mathbb{Z}_2^{n/2} \\ z_1=0}} (-1)^{\mathbf{z} \cdot \mathbf{y}_k} \text{Tr}[\phi_{k,\mathbf{z}}^2]^{t/2} \right]. \quad (37)$$

Here, we have chosen coset representatives $\mathbb{Z}_2^n/H_C = \{\mathbf{z}\bar{\mathbf{z}}^C : \mathbf{z}, \bar{\mathbf{z}} \in \mathbb{Z}_2^{n/2}, z_1 = \bar{z}_1 = 0\}$. The conclusion follows from a triangle inequality (i.e. ignoring the ± 1 phases) on all the terms on the right hand side except the leading term from $\mathbf{z} = 0^{n/2}$, which corresponds to a trivial void cut with purity one. \square

A sufficient condition for the algorithm to work is to ensure that the two distributions are negligibly close in a relative sense at the level of each outcome, meaning $P_{\text{StateHSP}_{\psi,t}}[\mathbf{y}] = P_{\text{HSP}}[\mathbf{y}] (1 + \text{negl}(n))$. Given the above fact, it is enough to choose the number of copies per sample t such that $\Delta_{\phi_{1,2},t} = \text{negl}(n)$.

Fact 4.6. *If a state $|\phi\rangle$ on $n/2$ qubits is at least ϵ -far from any separable state along any internal cut, then: $\Delta_{\phi,t} \leq 2^{n/2}(1 - \epsilon^2)^{t/2}$. Therefore, a choice of $t = O(n/\epsilon^2)$ copies per sample makes the relative correction in [Fact 4.5](#) negligible in n .*

Proof. The result follows from a straightforward binomial sum argument. Using [Proposition 5](#), we have that each nontrivial purity is upper-bounded by $\text{Tr}[\phi_S^2] \leq 1 - \epsilon^2$, therefore:

$$\Delta_{\phi,t} = \sum_{\substack{S \subset [n/2] \\ 1 \notin S \\ S \neq \emptyset}} \text{Tr}[\phi_S^2]^{t/2} \tag{38}$$

$$\leq 2^{n/2}(1 - \epsilon^2)^{t/2} \tag{39}$$

$$\leq e^{\frac{\ln 2}{2}n - \frac{t\epsilon^2}{2}}. \tag{40}$$

Therefore, a choice of $t = O(n/\epsilon^2)$ is enough to make this quantity negligible in n . \square

This completes the proof of the main theorem.

4.4 Improving the Abelian HSP algorithm by adaptive subspace preparations

In this section, we describe an adaptive modification of [Algorithm 1](#) which improves the number of state copies required to determine the hidden cut by a factor of n , from $O(n^2/\epsilon^2)$ down to $O(n/\epsilon^2)$. This achieves an optimal asymptotic in terms of the number of state copies (up to logarithmic factors) as announced in the introduction, given the related decision lower bound of Jones and Montanaro [[JM24](#)]. The adaptive algorithm operates as follows:

Algorithm 2: Hidden cut algorithm with adaptive subspaces

Parameters: n qubits, factor states promised to be ϵ away from product states.

Requirements: $2n$ additional qubits, implementation of the \mathbb{Z}_2^n group action $U_{\mathbb{Z}_2^n}$.

- 1 **for** *sample count* $k \in \{1, \dots, n-2\}$:
 - 2 Define the Boolean subspaces $\Sigma_k^\perp \equiv \text{span}\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)}\}$ and $\Sigma_k \equiv (\Sigma_k^\perp)^\perp$, defined as $\Sigma_k = \{\mathbf{z} \in \mathbb{Z}_2^n : \mathbf{z} \cdot \mathbf{y}^{(j)} = 0 \pmod{2}, \forall j \in [k-1]\}$. If $k = 1$, then set $\Sigma_k = \mathbb{Z}_2^n$.
 - 3 In the ancillary group register, prepare the superposition $|\Sigma_k\rangle = \frac{1}{\sqrt{2^{n-k+1}}} \sum_{\mathbf{z} \in \Sigma_k} |\mathbf{z}\rangle$.
 - 4 Prepare $t = O(1/\epsilon^2)$ copies of the state $|\psi\rangle$.
 - 5 Run the Fourier sampling circuit: $(H^{\otimes n} \otimes I) U_{\mathbb{Z}_2^n} |\Sigma_k\rangle \otimes |\psi\rangle^{\otimes t}$.
 - 6 Measure the group register to obtain a new sample $\mathbf{y}^{(k)} \in \mathbb{Z}_2^n$.
 - 7 Keep the sample if it is nonzero and outside Σ_k^\perp , otherwise repeat.
 - 8 Classically solve for the nullspace of $Y = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n-2)})^T \in \mathbb{Z}_2^{p \times n}$ which is $\text{span}\{1^C 0^{\bar{C}}, 0^C 1^{\bar{C}}\}$.
-

Compared to [Algorithm 1](#), the key difference is a different initial state in the ancillary register which hosts the regular representation of the parent group $G = \mathbb{Z}_2^n$. The previous [Algorithm 1](#) followed a standard Fourier sampling procedure which initialized the group register in a uniform superposition over all group elements, i.e. over all of \mathbb{Z}_2^n . By comparison, the adaptive [Algorithm 2](#) introduced here will instead initialize the group register in a uniform superposition over the Boolean subspace which is orthogonal to previously collected samples. We will show that this serves to boost the probability that new samples will be linearly independent, such that a smaller number of copies is needed at every step for amplification purposes. Our main result is the analysis of this algorithm, showing that it succeeds in finding the hidden cut with constant probability:

Theorem 1 (Hidden cut algorithm — restated). *For $\epsilon > 0$ and an n -qubit input state $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$ separable across a cut $C \in \binom{[n]}{n/2}$, assume that the factor states $|\phi_{1,2}\rangle$ are at least ϵ -far from all separable $(n/2)$ -qubit states. Then, [Algorithm 2](#) succeeds in finding the hidden cut C with constant probability using $O(n/\epsilon^2)$ copies of the input state $|\psi\rangle$. The algorithm requires coherent access to $O(1/\epsilon^2)$ copies at a time, on which it acts with circuits of depth $O(n^2) + O(\log \epsilon^{-1})$, and polynomial-time classical processing.*

Proof. Since [Algorithm 2](#) is a direct modification of the Fourier sampling approach of [Algorithm 1](#), the proof of this theorem proceeds along similar lines. Three key technical points need to be added to the analysis, which we prove in the rest of this section. First, we show that the ‘subspace states’ $|\Sigma_k\rangle$ can indeed be efficiently prepared on the group register at the beginning of each sampling round (this is shown in [Fact 4.7](#) below). The efficient circuits involved in preparing these states rely on finding basis vectors for the corresponding subspaces, which can be efficiently obtained classically. Second, we show that all samples lie inside the cut subspace H_C^\perp with probability one, which is a consequence of the hidden cut StateHSP instance admitting the subgroup H_C as the hidden symmetry subgroup; we show this in [Fact 4.8](#) below. Finally, we show that consuming a number of $t = O(1/\epsilon^2)$ state copies per sample results in a constant probability of the new sample being linearly independent with respect to previous samples (see [Fact 4.9](#) below). This suffices for an overall constant probability of success of [Algorithm 2](#) due to the rejection sampling procedure on [Line 7](#), since at every sampling round we reject new outcomes until they are linearly independent. \square

Fact 4.7. *If Σ is a d -dimensional subspace of \mathbb{Z}_2^n , then the n -qubit subspace state $|\Sigma\rangle \equiv \frac{1}{2^{d/2}} \sum_{\mathbf{z} \in \Sigma} |\mathbf{z}\rangle$*

can be efficiently prepared with circuits of size $O(nd) \leq O(n^2)$.

Proof. Given an n -bit string $\mathbf{z} \in \mathbb{Z}_2^n$, one can easily implement the $(1+n)$ -qubit controlled addition unitary $U_{\mathbf{z}} : |a\rangle |\mathbf{x}\rangle \mapsto |a\rangle |\mathbf{x} \oplus a\mathbf{z}\rangle$ for any $a \in \{0, 1\}$, $\mathbf{x} \in \mathbb{Z}_2^n$. Specifically, this can be implemented with a number $|\mathbf{z}| = O(n)$ of sequential CNOT gates controlled on the a register, which act on the \mathbf{x} registers in the locations on which the string \mathbf{z} has entries equal to one.

Let $\mathbf{z}_1, \dots, \mathbf{z}_d \in \mathbb{Z}_2^n$ be a basis of the subspace Σ . Then, by a sequence of unitaries $U_{\mathbf{z}_1}, \dots, U_{\mathbf{z}_d}$ of the kind described above, one can efficiently implement the $(d+n)$ -qubit unitary:

$$U_{\Sigma} : |\mathbf{a}\rangle |\mathbf{x}\rangle \mapsto |\mathbf{a}\rangle |\mathbf{x} \oplus a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle, \quad (41)$$

with a circuit of total depth $O(nd)$. Similarly, with the same gate count one can implement the ‘inverse’ $(d+n)$ -qubit unitary which acts as:

$$V_{\Sigma} : |\mathbf{b}\rangle |a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle \mapsto |\mathbf{b} \oplus \mathbf{a}\rangle |a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle, \quad (42)$$

for any $\mathbf{b}, \mathbf{a} = (a_1, \dots, a_d) \in \mathbb{Z}_2^d$.

Then, starting from the zero state on $d+n$ qubits, the substate state can be prepared as:

$$\begin{aligned} |\mathbf{0}\rangle \otimes |\mathbf{0}\rangle &\longrightarrow \frac{1}{2^{d/2}} \sum_{\mathbf{a} \in \mathbb{Z}_2^d} |\mathbf{a}\rangle \otimes |\mathbf{0}\rangle && \text{(applying } d \text{ Hadamard gates on the first } d \text{ qubits)} \\ &\longrightarrow \frac{1}{2^{d/2}} \sum_{\mathbf{a} \in \mathbb{Z}_2^d} |\mathbf{a}\rangle \otimes |a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle && \text{(applying the } U_{\Sigma} \text{ circuit defined above)} \\ &\longrightarrow |\mathbf{0}\rangle \otimes \frac{1}{2^{d/2}} \sum_{\mathbf{a} \in \mathbb{Z}_2^d} |a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle && \text{(applying the } V_{\Sigma} \text{ circuit defined above)} \\ &= |\mathbf{0}\rangle \otimes |\Sigma\rangle. \end{aligned}$$

This procedure prepares the substate state $|\Sigma\rangle = 2^{-d/2} \sum_{\mathbf{a} \in \mathbb{Z}_2^d} |a_1\mathbf{z}_1 \oplus \dots \oplus a_d\mathbf{z}_d\rangle$ on the last n qubits, with a circuit of overall size $O(nd)$. \square

Fact 4.8. Each new sample $\mathbf{y}^{(k)}$ is always in the cut subspace H_C^{\perp} .

Proof. The state prepared at the k -th round of Algorithm 2 is of the form:

$$\frac{1}{2^{n-(k-1)/2}} \sum_{\mathbf{y} \in \mathbb{Z}_2^n} |\mathbf{y}\rangle \otimes \sum_{\mathbf{z} \in \Sigma_k} (-1)^{\mathbf{z} \cdot \mathbf{y}} R(\mathbf{z}) |\psi\rangle^{\otimes t}, \quad (43)$$

on which measuring the first register returns an outcome $\mathbf{y} \in \mathbb{Z}_2^n$ with probability:

$$\begin{aligned} \mathbb{P}[\mathbf{y}] &= \frac{1}{2^{2n-k+1}} \sum_{\mathbf{z}, \mathbf{z}' \in \Sigma_k} (-1)^{\mathbf{y} \cdot (\mathbf{z} \oplus \mathbf{z}')} \langle \psi^{\otimes t} | R(\mathbf{z}')^{\dagger} R(\mathbf{z}) | \psi^{\otimes t} \rangle \\ &= \frac{1}{2^{2n-k+1}} \sum_{\mathbf{z}, \mathbf{z}' \in \Sigma_k} (-1)^{\mathbf{y} \cdot (\mathbf{z} \oplus \mathbf{z}')} \langle \psi^{\otimes t} | R(\mathbf{z}' \oplus \mathbf{z}) | \psi^{\otimes t} \rangle && \text{(as } R \text{ is } \mathbb{Z}_2^n \text{-representation)} \\ &= \frac{1}{2^n} \sum_{\mathbf{z} \in \Sigma_k} (-1)^{\mathbf{y} \cdot \mathbf{z}} \langle \psi^{\otimes t} | R(\mathbf{z}) | \psi^{\otimes t} \rangle && \text{(by summing over } \Sigma_k) \end{aligned}$$

where we used the fact that Σ_k is a $(n - k + 1)$ -dimensional subspace of \mathbb{Z}_2^n , so it also operates as a subgroup of \mathbb{Z}_2^n under addition.

We notice that if a string \mathbf{z} is in the subspace Σ_k , then we must have that all elements in the associated hidden coset are also in Σ_k . The argument proceeds by induction. Specifically, $\mathbf{z} \in \Sigma_k$ if it is orthogonal to previous samples: $\mathbf{z} \cdot \mathbf{y}^{(j)} = 0 \pmod{2}$, for $j \in [k - 1]$. Assume that previous samples are in the cut subspace H_C^\perp , meaning that $1^C 0^{\bar{C}} \cdot \mathbf{y}^{(j)} = 0^C 1^{\bar{C}} \cdot \mathbf{y}^{(j)} = 0$ for $j \in [k - 1]$. Then, we also have that $(\mathbf{z} \oplus 1^C 0^{\bar{C}}) \cdot \mathbf{y}^{(j)} = (\mathbf{z} \oplus 0^C 1^{\bar{C}}) \cdot \mathbf{y}^{(j)} = (\mathbf{z} \oplus 1^n) \cdot \mathbf{y}^{(j)} = 0 \pmod{2}$ for $j \in [k - 1]$ — in other words, if $\mathbf{z} \in \Sigma_k$, then also the rest of the coset $\mathbf{z} \oplus 1^C 0^{\bar{C}}$, $\mathbf{z} \oplus 0^C 1^{\bar{C}}$, $\mathbf{z} \oplus 1^n$ are in Σ_k . The base case for the induction is true due to the argument of the previous section which underlies [Algorithm 1](#). Another way of stating this fact is that H_C remains a subgroup of all intermediate subspaces Σ_k , when viewing Σ_k as subgroups of \mathbb{Z}_2^n .

Finally, we use the fact that H_C is the hidden subgroup defining this StateHSP, which means that the inner product $\langle \psi^{\otimes t} | R(\mathbf{z}) | \psi^{\otimes t} \rangle$ remains invariant when taking $\mathbf{z} \rightarrow \mathbf{z} + \mathbf{h}$, for $\mathbf{h} \in H_C = \{0^n, 1^C 0^{\bar{C}}, 0^C 1^{\bar{C}}, 1^n\}$. Therefore we can reformulate the outcome distribution derived above in terms of the cosets of Σ_k by the hidden subgroup H_C :

$$\begin{aligned} \mathbb{P}[\mathbf{y}] &= \frac{1}{2^n} \sum_{\mathbf{h} \in H_C} (-1)^{\mathbf{y} \cdot \mathbf{h}} \sum_{\mathbf{z} \in \Sigma_k / H_C} (-1)^{\mathbf{y} \cdot \mathbf{z}} \langle \psi^{\otimes t} | R(\mathbf{z}) | \psi^{\otimes t} \rangle \\ &= \frac{1 + (-1)^{\mathbf{y} \cdot 1^C 0^{\bar{C}}}}{2} \frac{1 + (-1)^{\mathbf{y} \cdot 0^C 1^{\bar{C}}}}{2} \frac{1}{2^{n-2}} \sum_{\mathbf{z} \in \Sigma_k / H_C} (-1)^{\mathbf{y} \cdot \mathbf{z}} \langle \psi^{\otimes t} | R(\mathbf{z}) | \psi^{\otimes t} \rangle \\ &= \frac{\delta_{\mathbf{y} \in H_C^\perp}}{2^{n-2}} \sum_{\mathbf{z} \in \Sigma_k / H_C} (-1)^{\mathbf{y} \cdot \mathbf{z}} \langle \psi^{\otimes t} | R(\mathbf{z}) | \psi^{\otimes t} \rangle, \end{aligned}$$

such that all measurement outcomes lie in the cut subspace $H_C^\perp = \{\mathbf{y} \in \mathbb{Z}_2^n : \mathbf{y} \cdot 1^C 0^{\bar{C}} = \mathbf{y} \cdot 0^C 1^{\bar{C}} = 0 \pmod{2}\}$ by a similar mechanism as in the previous algorithm. \square

Fact 4.9. *When using $t = O(1/\epsilon^2)$ state copies per sample, each new sample $\mathbf{y}^{(k)}$ is outside of the subspace Σ_k^\perp with constant probability.*

Proof. We can use the derived outcome distribution from the previous [Fact 4.8](#) to express the probability that a new sample is not in the subspace Σ_k^\perp by its complement:

$$\begin{aligned} \mathbb{P}[\mathbf{y} \notin \Sigma_k^\perp] &= 1 - \sum_{\mathbf{y} \in \Sigma_k^\perp} \mathbb{P}[\mathbf{y}] \\ &= 1 - \frac{1}{2^n} \sum_{\mathbf{y} \in \Sigma_k^\perp} \sum_{\mathbf{z} \in \Sigma_k} (-1)^{\mathbf{y} \cdot \mathbf{z}} \text{Tr}[\psi_{\mathbf{z}}]^{t/2} && \text{(In terms of purities, as in [Fact 4.5](#))} \\ &= 1 - \frac{1}{2^{n-k+1}} \sum_{\mathbf{z} \in \Sigma_k} \text{Tr}[\psi_{\mathbf{z}}]^{t/2} && \text{(Since } \mathbf{y} \cdot \mathbf{z} = 0\text{)} \\ &= 1 - \frac{1}{2^{n-k-1}} \sum_{\mathbf{z} \in \Sigma_k / H_C} \text{Tr}[\psi_{\mathbf{z}}]^{t/2}, && \text{(Organizing the sum by cosets)} \end{aligned}$$

where in the last line we split the sum over the H_C cosets of Σ_k , using the findings from the proof

of [Fact 4.8](#) outlined above. Next, using [Proposition 5](#) to bound all nontrivial purities leads to:

$$\begin{aligned}
\mathbb{P}[\mathbf{y} \notin \Sigma_k^\perp] &= 1 - \frac{1}{2^{n-k-1}} \left(1 + \sum_{0^n \neq \mathbf{z} \in \Sigma_k/H_C} \text{Tr}[\psi_{\mathbf{z}}]^{t/2} \right) && \text{(Separating the zero term)} \\
&\geq 1 - \frac{1}{2^{n-k-1}} \left(1 + (2^{n-k-1} - 1)(1 - \epsilon^2)^{t/2} \right) && \text{(Using Proposition 5)} \\
&\geq \frac{1}{2} \left(1 - (1 - \epsilon^2)^{t/2} \right). && \text{(Since } k \in [n - 2])
\end{aligned}$$

Therefore, a choice of $t = O(1/\epsilon^2)$ makes this probability at least a constant, which suffices for the purpose of [Algorithm 2](#). \square

5 The special case of Haar-random states: proof of [Theorem 2](#)

In this section, we will study the hidden cut problem when the factor states are promised to be sampled independently from the Haar measure. Intuitively, Haar-random states would be at least a constant distance away from product states with high probability, such that [Algorithm 2](#) of the previous section can be applied to find the cut given $O(n)$ state copies. While we do not improve on the number of state copies required (and beyond a possible factor of $\log n$, no improvement should be possible at all, given the decision lower bound of [\[JM24\]](#)), in this section we show how a careful analysis can reduce the other algorithmic requirements. Specifically, instead of running our adaptive algorithm ([Algorithm 2](#)), we show that our first, conceptually simpler non-adaptive algorithm ([Algorithm 1](#)) suffices in this case. By taking advantage of the properties of the Haar measure, we show that in this case [Algorithm 1](#) finds the cut with minimal requirements, involving circuits of constant depth (as opposed to depth $O(n^2) + O(\log \epsilon^{-1})$ as required by [Algorithm 2](#)) acting on only two state copies at a time:

Theorem 2 (Hidden cut algorithm with Haar-random states — restated). *Under the stronger promise of Haar-random factor states, the hidden cut can be found by the version of [Algorithm 1](#) using only $O(n)$ copies of the input state, by running circuits of constant depth which coherently access only $t = 2$ state copies at a time.*

To prove this result, we will have to further analyze the details of the StateHSP Fourier sampling distribution. In particular, we will relax the strong requirement of negligible relative error between the StateHSP and HSP Fourier sampling distributions ([2](#)) used in the previous sections. The Haar measure toolkit will nonetheless provide enough analytic control over the resulting distributions. Many of the technical details will be delegated to [Appendices A](#) and [B](#), but this section will contain the main workflow behind the proof of [Theorem 2](#).

To start, define the Fourier purity probabilities generated by an n -qubit state $|\phi\rangle$ as:

$$\mathbb{P}(\mathbf{y}; \phi) \equiv \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} (-1)^{\mathbf{y} \cdot \mathbf{x}} \text{Tr}[\phi_{\mathbf{x}}^2], \quad \text{where } \mathbf{y} \in \mathbb{Z}_2^n, \quad (44)$$

where $\text{Tr}[\phi_{\mathbf{x}}^2]$ is the purity of $|\phi\rangle$ across the cut represented by the binary vector $\mathbf{x} \in \mathbb{Z}_2^n$, i.e. when tracing out the qubits in the set $S_{\mathbf{x}} = \{i \in [n] : x_i = 1\}$. Notice that this is similar to the probabilities studied in the previous section, except that the number of copies is fixed to $t = 2$.

The central fact is that these quantities self-average in a strong sense under Haar-random states, as formalized in the following lemma:

Lemma 1 (Self-averaging of Fourier sampling distribution). *With high probability over the choice of a Haar-random state $|\phi\rangle$, the Fourier probabilities self-concentrate:*

$$P(\mathbf{y}; \phi) = \mathbb{E}_{\psi \sim \text{Haar}[(\mathbb{C}^2)^{\otimes n}]} P(\mathbf{y}; \psi) (1 + \text{negl}(n)) \quad (45)$$

$$= \delta_{\mathbf{y} \cdot \mathbf{1}^n = 0 \pmod{2}} \frac{2 \cdot 3^{n-|\mathbf{y}|}}{2^n(2^n + 1)} (1 + \text{negl}(n)). \quad (46)$$

Specifically, with probability at least $1 - \delta$:

$$\forall \mathbf{y} \in \mathbb{Z}_2^n : \left| \frac{P(\mathbf{y}; \phi)}{\mathbb{E}_{\psi \sim \text{Haar}[(\mathbb{C}^2)^{\otimes n}]} P(\mathbf{y}; \psi)} - 1 \right| \leq \delta^{-1/2} \cdot 3^{-n/2+o(1)}. \quad (47)$$

Therefore, we can choose $\delta = \kappa^{-n}$ for any constant $\kappa \in (1, 3)$ to satisfy the conclusion.

Proof. The proof follows from second-order tail bounds applied to the covariance of the internal purities of a Haar-random state. We delegate the proof details to [Appendix A](#). \square

Applying this fact to the factor states making up the separable input state immediately leads to the following modification of [Fact 4.5](#) in the case of Haar-random states, when we restrict the number of copies to $t = 2$:

Fact 5.1. *Consider the hidden cut problem with input state $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$ separable across cut $C \in \binom{[n]}{n/2}$, and assume the factor states $|\phi_{1,2}\rangle$ are independent Haar-random states on $n/2$ qubits. Then, with high probability over the Haar-random samples $|\phi_{1,2}\rangle$, the Fourier sampling probability distribution with $t = 2$ copies of the state is:*

$$P_{\text{StateHSP}_{\psi,2}}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] = P(\mathbf{y}_1; \phi_1) P(\mathbf{y}_2; \phi_2) \quad (48)$$

$$= \frac{\delta_{|\mathbf{y}_1| \text{ even}} \delta_{|\mathbf{y}_2| \text{ even}}}{2^{n-2}} \times \frac{3^{n-|\mathbf{y}_1|-|\mathbf{y}_2|}}{(2^{n/2} + 1)^2} (1 + \text{negl}(n)), \quad (49)$$

where the first factor is the associated HSP Fourier distribution $P_{\text{HSP}}[\mathbf{y}]$ defined in [\(30\)](#), which is uniform over the $(n - 2)$ -dimensional Boolean subspace H_C^\perp induced by the cut.

This distribution is equivalent (up to the negligible relative correction) to producing outcomes \mathbf{y} by the following rejection sampling protocol: produce a sample $\mathbf{y} \in \mathbb{Z}_2^n$ by independently sampling each bit from a Bernoulli distribution $y_1, \dots, y_n \sim \text{Ber}(1/4)$; keep the sample \mathbf{y} if it lies inside the cut subspace H_C^\perp , and sample again otherwise.

We observe that, while the HSP distribution is uniform over the cut subspace H_C^\perp , the distribution [\(48\)](#) derived above is still supported inside the cut subspace H_C^\perp , however it is non-uniform since it skews towards smaller-weight outcomes. It remains to show that this modification does not significantly impact the number of samples required to accumulate a complete basis of the cut subspace H_C^\perp :

Theorem 2 (Hidden cut algorithm with Haar-random states — restated). *When the input state $|\psi\rangle = |\phi_1\rangle_C \otimes |\phi_2\rangle_{\bar{C}}$ is a product of two $(n/2)$ -qubit Haar-random factor states $|\phi_1\rangle, |\phi_2\rangle$, a variation*

of the hidden cut algorithm finds the hidden cut with constant probability and using only $O(n)$ copies of the state, involving circuits of constant depth which coherently access only two state copies at a time.

Proof outline. We delegate the full proof to [Appendix B](#), but we outline the argument here. The standard form of Simon’s algorithm, including the modification relevant for [Theorem 1](#) above, relies on output distributions which are uniformly supported inside the hidden subspace. This uniformity condition makes it easy to compute the probabilities involved in the “basis coupon collection” process, showing that $n - k$ independent random samples can form a complete basis for an $(n - k)$ -dimensional hidden subspace with constant probability. However, we are interested in the special case of purity Fourier sampling with a hidden cut state made of Haar-random factor states, which produces n -bit string outputs from the $P_{\text{StateHSP}_{\psi,2}}$ distribution defined above in [\(48\)](#). As mentioned in [Fact 5.1](#), the resulting samples are not uniformly distributed inside the subspace, but they are skewed towards shorter-weight strings; the distribution is equivalent to rejection-sampling from an entrywise Bernoulli with $1/4$ probability of returning one, and keeping the sample if it lies inside the cut subspace H_C^\perp . The simple mathematics of Simon’s coupon collection does not work anymore since the uniform assumption is violated. Our goal is to show that nonetheless, a similar conclusion still holds in this non-uniform case, such that a basis for the cut subspace can be collected in $O(n)$ samples.

In fact, we will prove a slightly weaker form of the necessary basis coupon collection, but one which places us within an $O(1)$ distance to the full answer. Recall that the hidden cut subspace $H_C^\perp \subset \mathbb{Z}_2^n$ is of dimension $n - 2$. Specifically, we will show that $n - 3$ independent samples from the desired distribution [\(48\)](#) are linearly independent with a probability of at least one half. This is only one false cut direction away from finding the true cut. Specifically, the nullspace of the matrix $Y \in \mathbb{Z}_2^{(n-3) \times n}$ whose rows are the linearly independent samples has dimension three, and will contain eight vectors; two of them are the trivial cuts 0^n and 1^n , leaving a number of six non-trivial candidate cuts. Two of the six candidate cuts are the two equivalent cut strings $1^C 0^{\bar{C}}$ and $1^{\bar{C}} 0^C$. Checking the six candidate cuts can be done by cut-specific SWAP tests, each SWAP test requiring a constant number of copies for a constant success probability guarantee. This is enough to show that the true cut can be found with a constant probability using $O(n)$ state copies.

We remark that numerical evidence strongly suggests that $n - 2$ i.i.d. samples from the distribution [\(48\)](#) form a complete basis for the cut subspace with constant probability, such that in practice the standard Simon’s basis coupon collection routine succeeds in this case as well without the need to explicitly find the last basis vector via SWAP tests. We leave a formal proof of this technical conjecture about Boolean random matrix theory to future work. \square

6 The many-cut case

The hidden cut algorithms from the previous section also apply naturally to the more general setting in which the input state is separable into two unequal subsystems, or indeed into more than two subsystems across an arbitrary set partition of the qubits, which we will refer to as the ‘hidden many-cut problem’. Finding the cut means identifying the set partition $C_1 \sqcup \dots \sqcup C_m = [n]$ (where the \sqcup symbol stands for disjoint union). Just as before, the number of copies used coherently to generate each Fourier sample will be chosen such that the corresponding StateHSP problem produces a similar outcome as the benchmark HSP distribution, up to negligible relative corrections in n .

The Simon-like target distribution will now be supported on an $(n - m)$ -dimensional subspace of \mathbb{Z}_2^n orthogonal to the cut strings, where m is the number of parts in the cut:

$$P_{\text{HSP}}[\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}] = \begin{cases} 2^{-n+m} & \text{if } \mathbf{y}_k \cdot \mathbf{1}^{|C_k|} = 0 \pmod{2}, \text{ for all } k \in [m] \\ 0 & \text{otherwise.} \end{cases} \quad (50)$$

Here, the notation $\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}$ represents the n -bit string with $\mathbf{y}_1 \in \mathbb{Z}_2^{|C_1|}$ in the positions in C_1 , $\mathbf{y}_2 \in \mathbb{Z}_2^{|C_2|}$ in the positions in C_2 , etc. A similar analysis bounding the contributions from possible false cuts applies at the level of the m factor states.

Additionally, we note that the case of interest for our algorithm involves many-cuts for which all the parts are bigger than a constant. Otherwise, a naïve brute-force approach involving sequential SWAP tests of all possible qubit combinations of constant size can be performed to iteratively discover the parts of the cut in polynomial time.

Corollary 1 (Algorithm for the many-cut problem — restated). *Assume an n -qubit state $|\psi\rangle$ is separable across an unknown set partition into m parts $C_1 \sqcup \dots \sqcup C_m = [n]$:*

$$|\psi\rangle = |\phi_1\rangle_{C_1} \otimes \dots \otimes |\phi_m\rangle_{C_m}. \quad (51)$$

Then the set partition/‘many-cut’ $\{C_1, \dots, C_m\}$ can be identified in polynomial time:

- (a) *If the factor states are promised to be ϵ -far in trace distance from separable, then [Algorithm 2](#) can identify the many-cut with $O(n/\epsilon^2)$ total number of state copies, with constant success probability. The algorithm runs circuits of depth $O(n^2) + O(\log \epsilon^{-1})$ on $O(1/\epsilon^2)$ state copies at a time.*
- (b) *If the factor states are Haar-random, then [Algorithm 1](#) can identify the many-cut with $O(n)$ state copies with high probability, provided the additional constraint that the cut parts are superlogarithmic in size: $\min_{k \in [m]} |C_k| > \omega(\log n)$. The algorithm runs circuits of constant depth, acting on two state copies at a time.*

Proof. The proof follows the same logic as [Theorem 1](#) (see [Section 4.4](#)) and [Theorem 2](#) (see [Section 5](#)). The outcome probability of the non-adaptive Fourier sampling circuit (i.e. the updated version of (37) in the case of m partitions) becomes:

$$P_{\text{StateHSP}_{\psi,t}}[\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}] = P_{\text{HSP}}[\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}] \prod_{k \in [m]} \left[\sum_{\substack{\mathbf{x}_k \in \mathbb{Z}_2^{|C_k|} \\ x_{k,1}=0}} (-1)^{\mathbf{y}_k \cdot \mathbf{x}_k} \text{Tr}[\phi_{k,\mathbf{x}_k}^2]^{t/2} \right], \quad (52)$$

where the benchmark HSP distribution is the one defined in (50). A similar triangle inequality as in [Fact 4.5](#) gives us that:

$$\left| P_{\text{StateHSP}_{\psi,t}}[\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}] - P_{\text{HSP}}[\mathbf{y}_1^{C_1} \dots \mathbf{y}_m^{C_m}] \right| \leq \frac{1}{2^{n-m}} \left[\prod_{k \in [m]} (1 + \Delta_{\phi_k,t}) - 1 \right], \quad (53)$$

where we define, as before:

$$\Delta_{\phi_k,t} \equiv \sum_{\substack{\mathbf{x}_k \in \mathbb{Z}_2^{|C_k|} \\ \mathbf{x}_k \neq \mathbf{0}^{|C_k|} \\ x_{k,1}=0}} \text{Tr}[\phi_{k,\mathbf{x}_k}^2]^{t/2}. \quad (54)$$

For the non-adaptive [Algorithm 1](#) to efficiently find the cut, it is sufficient to choose t such that all $\Delta_{\phi_k, t}$ are negligible in n .

- (a) If $|\phi\rangle_k$ is at least ϵ -far from all separable states, then [Proposition 5](#) gives us that $\text{Tr}[\phi_{k, \mathbf{x}}^2] \leq 1 - \epsilon^2$ for any nontrivial cut \mathbf{x} , leading to an upper bound via triangle inequality:

$$\Delta_{\phi_k, t} \leq (2^{|C_k|} - 1)(1 - \epsilon^2)^{t/2} \quad (55)$$

$$\leq 2^{|C_*| - O(t\epsilon^2)}. \quad (56)$$

Here, we defined the largest cut component size as $|C_*| \equiv \max_{k \in [m]} |C_k|$. The above can be made negligible in n if $t = O(|C_*|/\epsilon^2) + O(\log^2 n)$. This suffices to show how the non-adaptive [Algorithm 1](#) applies to finding the many-cut with high probability with $O(n/\epsilon^2)$.

Applying the adaptive [Algorithm 2](#) to the many-cut case is similarly straightforward. Just as in [Section 4.4](#), the adaptive algorithm will accumulate $n - m$ linearly independent vectors in the $(n - m)$ -dimensional cut subspace H_C^\perp by the adaptive Fourier sampling method. New samples are always in the cut subspace by the StateHSP subgroup symmetry; a number of copies $t = O(1/\epsilon^2)$ suffices to lower-bound the probability that a new sample is linearly independent with respect to previously collected samples by a constant. This in turn is enough to make sure that the algorithm terminates and succeeds to find the cut with a constant overall success probability.

- (b) The results of [Section 5](#) apply in this case as well, since each part C_k of the partition incurs a relative error of size $O(2^{-|C_k|})$. Since there are at most n parts, the overall corrections remain negligible as long as each individual correction remains negligible, i.e. if all parts are more than logarithmic in size, i.e. $\min_{k \in [m]} |C_k| > \omega(\log n)$. \square

We conjecture that the stricter requirement of superlogarithmic part size in the case of Haar-random factor states can be removed with a more careful accounting of the concentration properties of Haar-random states.

Finally, we remark that one does not necessarily need to know the number of unentangled parts which make up the input state a priori, since this is not a parameter in our algorithms. In fact, one can efficiently infer the number of parts in the many-cut by analyzing the linear independence of the obtained Fourier samples: if the input state is a product of m factor states, then the rank of the accumulated samples will plateau at a value of $n - m$.

7 Discussion, applications, and open questions

7.1 Applications: cryptography and pseudorandomness

As mentioned in the introduction, our hidden cut algorithm provides a no-go result for certain recursive constructions of pseudorandom states. In particular, our algorithm shows that a product of pseudorandom states across a random cut is not itself pseudorandom. More broadly, our hidden cut algorithm prohibits pseudorandom state constructions with zero entanglement across any partition of the qubits. However, certain generalizations of these constructions are not ruled out by our algorithm. For example, consider a nonzero-entropy hidden cut state, for example a rank-two state

of the form:

$$|\psi\rangle \cong \frac{1}{\sqrt{2}} (|\alpha_1\rangle_C \otimes |\beta_1\rangle_{\bar{C}} + |\alpha_2\rangle_C \otimes |\beta_2\rangle_{\bar{C}}), \quad (57)$$

where the tensor products are taken across a random cut $C \subset [n]$, and the sub-states $|\alpha_{1,2}\rangle$ and $|\beta_{1,2}\rangle$ are pairs of orthogonal pseudorandom states. Such a state would have constant, but nonzero entanglement entropy across the cut; the cut remains information-theoretically detectable, and verifying the cut can still be achieved with only a constant number of copies via a standard SWAP test. Interestingly, running our algorithm on such an input state would fail to identify the hidden cut C . Specifically, the resulting Fourier sampling distribution obtained becomes a noisy version of the Simon’s problem, such that the samples now have a constant, nonzero probability of lying outside of the cut subspace H_C^\perp . In other words, solving for the cut subspace would require solving a noisy system of linear equations over \mathbb{Z}_2^n with a constant noise rate, i.e. it would require solving a version of the learning parity with noise (LPN) problem, which is conjectured to be cryptographically hard. We leave open the question of designing an algorithm for finding the hidden cut in the nonzero cut entropy scenario, or alternatively of producing further evidence that the problem is computationally hard.

We also note the hidden cut problem might be useful for constructions of quantum money. Here, the goal is to produce ‘banknote’ states which are difficult to copy but easy to verify. One could imagine a quantum money scheme based on the hidden cut problem, in which the cut serves as the secret key, and the banknote would be composed of only a constant number of copies of the separable state. This means our algorithm cannot be used to find the cut, since it would require a linear number of copies. Therefore, it is possible the cut could be cryptographically protected. On the other hand, verifying the cut only requires only a constant number of copies by a standard fixed-cut SWAP test. It remains an open question whether such a quantum money construction can be made public-key compatible. We remark that the factor states themselves, not just the location of the cut, could potentially serve a cryptographic function.

7.2 The StateHSP framework

Motivated by the hidden cut problem, in [Section 3](#) we introduced a state version of the hidden subgroup problem as a flexible framework for problems with state input which feature a hidden symmetry subgroup. A natural question is whether the StateHSP framework can be used to derive quantum algorithms for other quantum information tasks. One source of inspiration could be tasks in unitary complexity theory [[RY22](#), [MY23](#), [BEM⁺23](#)]. Alternatively, in the other direction there is the question as to whether StateHSP can give rise to cryptographic primitives via information-computation gaps. As a corollary of our work is that StateHSP is information-theoretically solvable with enough copies and orthogonality allowance (see [Corollary 5](#)), which opens the possibility of information-computation gaps in the general case. Another potential avenue is to consider cases where the mechanism for orthogonality amplification by preparing poly-many state copies is not available — in such cases, the StateHSP might become hard, while also resisting the reduction to HSP via orthogonality amplification.

7.3 Entanglement features

We note that a combinatorial view considering all internal purities of a given state ψ organized as a so-called *entanglement feature* vector $|W_\psi\rangle \propto \sum_{\mathbf{x} \in \mathbb{Z}_2^n} \text{Tr}[\psi_{\mathbf{x}}^2] |\mathbf{x}\rangle$ appears in the condensed matter

literature⁷ [YYQ18, YG18, FVVY21]. While constructing the entanglement feature state $|W_\psi\rangle$ from the input state $|\psi\rangle$ seems to be generally hard, our hidden cut algorithm is able to indirectly manipulate this quantity. In particular, we remark that our [Algorithm 1](#) is able to perform Fourier sampling on the moments of the entanglement feature vector (see equation (35)). We leave it as an open question to further explore this connection.

Acknowledgments

We thank Roozbeh Bassirian, Bill Fefferman, Soumik Ghosh, Patrick Hayden, Matteo Ippoliti, Fernando Jeronimo, Benjamin Jones, Ashley Montanaro, Henry Yuen, Chenyi Zhang, and Jack Zhou for insightful discussions. A.B. and T.G.T. were supported in part by the U.S. DOE Office of Science under Award Number DE-SC0020377. A.B. was supported in part by the DOE QuantISED grant DE-SC0020360 and by the AFOSR under grants FA9550-21-1-0392 and FA9550-24-1-0089. J.W. is supported in part by NSF CAREER award CCF-2339711.

⁷We thank Matteo Ippoliti for pointing out this connection.

References

- [ABF⁺24] Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum Pseudoentanglement. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, 2024. [arXiv:2211.00747](#).
- [BDJ99] Jinho Baik, Percy Deift, and Kurt Johansson. On the distribution of the length of the longest increasing subsequence of random permutations. *Journal of the American Mathematical Society*, 12(4):1119–1178, 1999. [arXiv:math/9810105](#).
- [Bea97] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 48–53, 1997.
- [BEM⁺23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the Uhlmann transformation problem. *arXiv preprint arXiv: 2306. 13073*, 2023.
- [BO20] Costin Bădescu and Ryan O’Donnell. Lower bounds for testing complete positivity and quantum separability. In *Latin American Symposium on Theoretical Informatics*, pages 375–386. Springer, 2020. [arXiv:1905.01542](#).
- [CVD10] Andrew M Childs and Wim Van Dam. Quantum algorithms for algebraic problems. *Reviews of Modern Physics*, 82(1):1–52, 2010. [arXiv:0812.0380](#).
- [Dia88] Persi Diaconis. *Group Representations in Probability and Statistics*, volume 11. Institute of Mathematical Statistics, 1988.
- [EHK04] Mark Ettinger, Peter Høyer, and Emanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, 91(1):43–48, 2004. [arXiv:quant-ph/0401083](#).
- [FO24] Steven T Flammia and Ryan O’Donnell. Quantum chi-squared tomography and mutual information testing. *Quantum*, 8:1381, 2024. [arXiv:2305.18519](#).
- [FVVY21] Ruihua Fan, Sagar Vijay, Ashvin Vishwanath, and Yi-Zhuang You. Self-organized error correction in random unitary circuits with measurement. *Physical Review B*, 103(17):174309, 2021. [arXiv:2002.12385](#).
- [Gao15] Jingliang Gao. Quantum union bounds for sequential projective measurements. *Physical Review A*, 92(5):052331, 2015.
- [GC01] Daniel Gottesman and Isaac Chuang. Quantum digital signatures. *arXiv preprint quant-ph/0105032*, 2001.
- [Gha08] Sevag Gharibian. Strong NP-hardness of the quantum separability problem. *arXiv preprint arXiv: 0810. 4507*, 2008.
- [GHMW13] Gus Gutoski, Patrick Hayden, Kevin Milner, and Mark M Wilde. Quantum interactive proofs and the complexity of separability testing. *arXiv preprint arXiv: 1308. 5788*,

- 2013.
- [GSVV04] Michelangelo Grigni, Leonard Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica*, 1(24):137–154, 2004.
- [Har13] Aram W Harrow. The church of the symmetric subspace. *arXiv preprint [arXiv:1308.6595](#)*, 2013.
- [HLM17] Aram W Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1598–1611. SIAM, 2017. [arXiv:1607.03236](#).
- [HM13] Aram W Harrow and Ashley Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *Journal of the ACM (JACM)*, 60(1):1–43, 2013. [arXiv:1001.0017](#).
- [HRTS03] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. The hidden subgroup problem and quantum computation using group representations. *SIAM Journal on Computing*, 32(4):916–934, 2003.
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018. [iacr:2018/544](#).
- [JM24] Benjamin DM Jones and Ashley Montanaro. Testing multipartite productness is easier than testing bipartite productness. *arXiv preprint [arXiv:2406.16827](#)*, 2024.
- [LG17] Joshua Lockhart and Carlos E González Guillén. Quantum state isomorphism. *arXiv preprint [arXiv:1709.09622](#)*, 2017.
- [LMW24] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 979–990, 2024. [arXiv:2310.08870](#).
- [LRW23] Margarite L LaBorde, Soorya Rethinasamy, and Mark M Wilde. Testing symmetry on quantum computers. *Quantum*, 7:1120, 2023. [arXiv:2105.12758](#).
- [MdW13] Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. *arXiv preprint [arXiv:1310.2035](#)*, 2013.
- [MRR06] Cristopher Moore, Daniel Rockmore, and Alexander Russell. Generic quantum Fourier transforms. *ACM Transactions on Algorithms (TALG)*, 2(4):707–723, 2006.
- [MRS08] Cristopher Moore, Alexander Russell, and Leonard J Schulman. The symmetric group defies strong Fourier sampling. *SIAM Journal on Computing*, 37(6):1842–1864, 2008. [quant-ph/0501056](#).
- [MY23] Tony Metger and Henry Yuen. $\text{stateQIP} = \text{statePSPACE}$. In *2023 IEEE 64th Annual*

- Symposium on Foundations of Computer Science (FOCS)*, pages 1349–1356. IEEE, 2023. [arXiv:2301.07730](#).
- [NC10] Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2nd edition, 2010.
- [OW16] Ryan O’Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 899–912, 2016. [arXiv:1508.01907](#).
- [Reg04] Oded Regev. Quantum computation and lattice problems. *SIAM Journal on Computing*, 33(3):738–760, 2004. [arXiv:cs/0304005](#).
- [RLW23] Soorya Rethinasamy, Margarite L LaBorde, and Mark M Wilde. Quantum Computational Complexity and Symmetry. *arXiv preprint [arXiv:2309.10081](#)*, 2023.
- [Roi96] Yuval Roichman. Upper bound on the characters of the symmetric groups. *Inventiones mathematicae*, 125:451–485, 1996.
- [RY22] Gregory Rosenthal and Henry Yuen. Interactive Proofs for Synthesizing Quantum States and Unitaries. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, 2022. [arXiv:2108.07192](#).
- [Sim97] Daniel R Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.
- [SW22] Mehdi Soleimanifar and John Wright. Testing matrix product states. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1679–1701. SIAM, 2022. [arXiv:2201.01824](#).
- [Wat00] John Watrous. Succinct quantum proofs for properties of finite groups. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 537–546. IEEE, 2000. [arXiv:cs/0009002](#).
- [YG18] Yi-Zhuang You and Yingfei Gu. Entanglement features of random Hamiltonian dynamics. *Physical Review B*, 98(1):014309, 2018. [arXiv:1803.10425](#).
- [YYQ18] Yi-Zhuang You, Zhao Yang, and Xiao-Liang Qi. Machine learning spatial geometry from entanglement features. *Physical Review B*, 97(4):045153, 2018. [arXiv:1709.01223](#).
- [Zha24] Mark Zhandry. Quantum Money from Abelian Group Actions. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, 2024. [arXiv:2307.12120](#).

A Internal purity covariance of Haar-random states and the self-averaging of Fourier sampling distributions

A.1 Proof of Lemma 1

The proof of Lemma 1 will make use of the two auxiliary results Fact A.1 and Fact A.2, detailed below. Here, we import these facts to show how they lead to the conclusion of Lemma 1.

The key quantity is the collection of purity Fourier sampling probabilities induced by a state $|\phi\rangle$ (44), defined as:

$$P(\mathbf{y}; \phi) \equiv \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{Z}_2^n} (-1)^{\mathbf{y} \cdot \mathbf{x}} \text{Tr}[\phi_{\mathbf{x}}^2], \quad \text{where } \mathbf{y} \in \mathbb{Z}_2^n. \quad (58)$$

The goal is to study how these probabilities self-average when the state $|\phi\rangle$ is fixed to a typical sample from the Haar measure on n -qubit states, which we will denote by Haar_n to condense notation.

Fact A.2 gives us explicit expression for the mean and variance of a single purity Fourier sampling probability $P(\mathbf{y}; \phi)$, when $|\phi\rangle$ is sampled from the Haar measure on n -qubit states. Chebyshev's inequality then allows us to show how one of these quantities concentrates:

$$\mathbb{P}_{\phi \sim \text{Haar}_n} \left[\left| P(\mathbf{y}; \phi) - \mathbb{E}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi) \right| \geq \beta \mathbb{E}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi) \right] \leq \frac{\text{Var}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi)}{\beta^2 (\mathbb{E}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi))^2} \quad (59)$$

$$\leq \frac{3^{|\mathbf{y}|-n} (1 + 2^{-n}) - 2^{-2n+1}}{\beta^2 (6 + 5 \cdot 2^n + 4^n)}. \quad (60)$$

Here, the states are implicitly understood to be sampled from the Haar measure on n -qubit states. The notation $|\mathbf{y}|$ denotes the Hamming weight of the bit-string $\mathbf{y} \in \mathbb{Z}_2^n$. We want a typicality statement about all of the probabilities $(P_{\phi}(\mathbf{y}))_{\mathbf{y} \in \mathbb{Z}_2^n}$ induced by a single state $|\phi\rangle$ sampled from the Haar measure. This can be obtained by a simple union bound over all the outcomes $\mathbf{y} \in \mathbb{Z}_2^n$, resulting in:

$$\mathbb{P}_{\phi \sim \text{Haar}_n} \left[\exists \mathbf{y} \in \mathbb{Z}_2^n : \left| P(\mathbf{y}; \phi) - \mathbb{E}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi) \right| \geq \beta \mathbb{E}_{\psi \sim \text{Haar}_n} P(\mathbf{y}; \psi) \right] \leq \frac{(2/3)^n + (4/3)^n - 2^{1-n}}{\beta^2 (6 + 5 \cdot 2^n + 4^n)} \quad (61)$$

$$\leq \frac{3^{-n+o(1)}}{\beta^2}, \quad (62)$$

from which the conclusion follows. \square

A.2 The covariance of internal purities of Haar-random states

The proof above invokes the following two facts involving properties of the second and fourth moments of the Haar measure. The first fact involves calculating the mean and covariance of the internal purities of a Haar-random state:

Fact A.1. *Let $|\phi\rangle$ be a state sampled from the n -qubit Haar measure. Let $\text{Tr}[\phi_{\mathbf{x}}^2]$ denote the purity of $|\phi\rangle$ across the cut determined by the bit-string \mathbf{x} , i.e. by tracing out the qubits in $S_{\mathbf{x}} \equiv \{i \in [n] : x_i = 1\}$. Then we have that the average purity across a cut is:*

$$\mathfrak{p}_{\mathbf{x}} \equiv \mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}}^2] \quad (63)$$

$$= \frac{2^{-|\mathbf{x}|} + 2^{-n+|\mathbf{x}|}}{1 + 2^{-n}}. \quad (64)$$

Also, the covariance between pairs of purities is given by:

$$\Sigma_{\mathbf{x}, \mathbf{x}'} \equiv \mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}}^2] \text{Tr}[\phi_{\mathbf{x}'}^2] - \mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}}^2] \mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}'}^2] \quad (65)$$

$$= 2 \frac{2^{|\mathbf{x} \oplus \mathbf{x}'|} + 2^{n-|\mathbf{x} \oplus \mathbf{x}'|}}{(2^n + 1)(2^n + 2)(2^n + 3)} - \frac{2}{(2^n + 2)(2^n + 3)} \mathbf{p}_{\mathbf{x}} \mathbf{p}_{\mathbf{x}'}. \quad (66)$$

Proof. Averages involving state purities can be turned into averages over the corresponding symmetric subspace starting with the following reformulation:

$$\text{Tr}[\phi_{\mathbf{x}}^2] = \text{Tr} \left[\bigotimes_{i: x_i=1} R_i((1\ 2)) \cdot |\phi\rangle\langle\phi|^{\otimes 2} \right], \quad (67)$$

which follows the notation from Section 4. Specifically, $R_i(\sigma)$ applied the permutation $\sigma \in \mathbb{S}_t$ to the i -th qubit across the t copies. The above equality involving the simple purity $\text{Tr}[\phi_{\mathbf{x}}^2]$ corresponds to the special case $t = 2$. Next, we use the well-known equality between the Haar t -fold ensemble and the maximally mixed state over the t -fold symmetric subspace (see for example [Har13]), which in our notation translates to:

$$\mathbb{E}_{\phi \sim \text{Haar}_n} |\phi\rangle\langle\phi|^{\otimes t} = \frac{(2^n - 1)!}{(2^n + t - 1)!} \sum_{\sigma \in \mathbb{S}_t} R_1(\sigma) \otimes \cdots \otimes R_n(\sigma). \quad (68)$$

Applying this fact, we have that the average purity is:

$$\mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}}^2] = \text{Tr} \left[\bigotimes_{i: x_i=1} R_i((1\ 2)) \cdot \mathbb{E}_{\phi \sim \text{Haar}_n} |\phi\rangle\langle\phi|^{\otimes 2} \right] \quad (69)$$

$$= \frac{1}{2^n(2^n + 1)} \sum_{\sigma \in \mathbb{S}_2} \text{Tr} \left[\bigotimes_{i: x_i=1} R_i((1\ 2) \cdot \sigma) \otimes \bigotimes_{j: x_j=0} R_j(\sigma) \right] \quad (70)$$

$$= \frac{1}{2^n(2^n + 1)} \sum_{\sigma \in \mathbb{S}_2} \prod_{i: x_i=1} \text{Tr}[R_i((1\ 2) \cdot \sigma)] \prod_{j: x_j=0} \text{Tr}[R_j(\sigma)] \quad (71)$$

$$= \frac{1}{2^n(2^n + 1)} \sum_{\sigma \in \mathbb{S}_2} 2^{|\mathbf{x}| \text{cyc}((1\ 2) \cdot \sigma) + (n - |\mathbf{x}|) \text{cyc}(\sigma)}, \quad (72)$$

where in the final line we used the fact that $\text{Tr}[R_i(\sigma)] = 2^{\text{cyc}(\sigma)}$. Here, $\text{cyc}(\sigma)$ denotes the number of cycles in the permutation σ . The explicit sum over $\sigma \in \mathbb{S}_2 = \{(1)(2), (1\ 2)\}$ leads to the closed-form result:

$$\mathbf{p}_{\mathbf{x}} \equiv \mathbb{E}_{\phi \sim \text{Haar}_n} \text{Tr}[\phi_{\mathbf{x}}^2] \quad (73)$$

$$= \frac{2^{-|\mathbf{x}|} + 2^{-n+|\mathbf{x}|}}{1 + 2^{-n}}. \quad (74)$$

To calculate the covariance entries we will need a fourth-moment calculation. This is because the same technique used in (67) above can be used to rewrite the product of two purities in terms of

four copies of the state:

$$\mathrm{Tr}[\phi_{\mathbf{x}}^2] \mathrm{Tr}[\phi_{\mathbf{x}'}^2] = \mathrm{Tr} \left[\bigotimes_{i: x_i=1} R_i((1\ 2)) \bigotimes_{j: x'_j=1} R_j((3\ 4)) \cdot |\phi\rangle\langle\phi|^{\otimes 4} \right], \quad (75)$$

where in this case $t = 4$ because the column-wise permutations $R_i(\sigma)$ represent $\sigma \in \mathbb{S}_4$. We apply the same workflow from above to translate the average over 4 copies of the Haar-random state $|\phi\rangle$ to a combinatorial sum over the symmetric group \mathbb{S}_4 , leading to:

$$\mathbb{E}_{\phi \sim \mathrm{Haar}_n} \mathrm{Tr}[\phi_{\mathbf{x}}^2] \mathrm{Tr}[\phi_{\mathbf{x}'}^2] = \frac{1}{2^n(2^n+1)(2^n+2)(2^n+3)} \sum_{\sigma \in \mathbb{S}_4} 2^{|\mathbf{x} \setminus \mathbf{x}'| \mathrm{cyc}((1\ 2)\sigma) + |\mathbf{x} \cap \mathbf{x}'| \mathrm{cyc}((1\ 2)(3\ 4)\sigma) + |\mathbf{x}' \setminus \mathbf{x}| \mathrm{cyc}((3\ 4)\sigma) + (n - |\mathbf{x} \cup \mathbf{x}'|) \mathrm{cyc}(\sigma)}. \quad (76)$$

In the above, we introduced set notation for binary strings in a natural sense, meaning that $\mathbf{x} \setminus \mathbf{x}' \equiv \{i \in [n] : x_i = 1 \text{ and } x'_i = 0\}$, also $\mathbf{x} \cap \mathbf{x}' \equiv \{i \in [n] : x_i = 1 \text{ and } x'_i = 1\}$, as well as $\mathbf{x} \cup \mathbf{x}' = \{i \in [n] : x_i = 1 \text{ or } x'_i = 1\}$. From here, the closed-form expressions for the covariance matrix entries $\Sigma_{\mathbf{x}, \mathbf{x}'}$ follow from explicit calculation by summing over the permutations $\sigma \in \mathbb{S}_4$. \square

Fact A.2. *The average and variance of the purity Fourier probabilities $P(\mathbf{y}; \phi)$ when the state $|\phi\rangle$ is sampled from the Haar measure are:*

$$\mathbb{E}_{\phi \sim \mathrm{Haar}_n} P(\mathbf{y}; \phi) = \frac{3^{n-|\mathbf{y}|}(1 + (-1)^y)}{2^n(2^n + 1)} = \frac{2 \cdot 3^{n-|\mathbf{y}|}}{2^n(2^n + 1)} \delta_{|\mathbf{y}| \text{ even}} \quad (77)$$

$$\mathrm{Var}_{\phi \sim \mathrm{Haar}_n} P(\mathbf{y}; \phi) = \frac{2^{2-4n} \cdot 3^{n-2|\mathbf{y}|} (2^n(2^n + 1) 3^{|\mathbf{y}|} - 2 \cdot 3^n)}{(2^n + 1)^2 (2^n + 2) (2^n + 3)} \delta_{|\mathbf{y}| \text{ even}}. \quad (78)$$

Proof. The proof follows from explicit calculations with the results of [Fact A.1](#), applied to the definition [\(58\)](#) of the probabilities $P(\mathbf{y}; \phi)$. \square

Remark. While numerical evidence suggests a much stronger self-averaging result applies at the level of each purity $\mathrm{Tr}[\phi_{\mathbf{x}}^2]$ of a Haar-random sample, applying the above Chebyshev tail + naïve union bound to individual purities fails to confirm this. However, one reason why this simple approach succeeds in [Lemma 1](#) to bound the Fourier probabilities $P(\mathbf{y}; \phi)$ is that, in fact, the purity covariance matrix $\Sigma_{\mathbf{x}, \mathbf{x}'} = \mathbf{Cov}_{\phi}[\mathrm{Tr} \phi_{\mathbf{x}}^2, \mathrm{Tr} \phi_{\mathbf{x}'}^2]$ computed in [Fact A.1](#) above is approximately diagonalized by the Boolean Fourier transform. Letting $F = H^{\otimes n}$ be the Boolean Fourier transform unitary:

$$F_{\mathbf{y}, \mathbf{x}} = \frac{1}{2^{n/2}} (-1)^{\mathbf{y} \cdot \mathbf{x}}, \quad (79)$$

then we have that in the Fourier basis the purity covariance is almost diagonal:

$$(F \Sigma F^\dagger)_{\mathbf{y}, \mathbf{y}'} = \frac{4 \cdot 3^{n-|\mathbf{y}|}}{(2^n + 1)(2^n + 2)(2^n + 3)} \left(\delta_{\mathbf{y}, \mathbf{y}'} - \frac{2}{2^n(2^n + 1)} 3^{n-|\mathbf{y}'|} \right) \delta_{|\mathbf{y}| \text{ even}} \delta_{|\mathbf{y}'| \text{ even}}. \quad (80)$$

In other words, the covariance matrix $\Sigma_{\mathbf{x}, \mathbf{x}'}$ is dominated by the first term in equation [\(66\)](#).

B Simon's algorithm with non-uniform samples and the hidden cut algorithm for Haar product states

Recall that sampling $\mathbf{y} \in \mathbb{Z}_2^n$ from the distribution $\mathbb{P}_{\text{StateHSP}_{\psi,2}}$ (which we will denote $\mathbb{P}_{\psi,2}$ for ease of notation) defined in (48) is equivalent to sampling each of the n entries i.i.d. from a Bernoulli with $1/4$ probability of yielding 1 and $3/4$ probability of yielding 0, and afterwards keeping the sample if it has even weight on both sides of the cut. To restate the explicit form of the distribution (48):

$$\mathbb{P}_{\psi,2}[\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}] = \delta_{|\mathbf{y}_1| \text{ even}} \delta_{|\mathbf{y}_2| \text{ even}} \left(\frac{3}{4}\right)^n 3^{-|\mathbf{y}_1| - |\mathbf{y}_2|} (1 + O(2^{-n/2})). \quad (81)$$

Let us first prove two helpful statements which will build towards [Theorem 2](#). The first fact is a convenient simplification:

Fact B.1. *Assume $p \leq O(n)$. Averaging over p samples from the $\mathbb{P}_{\psi,2}$ cut-specific distribution can be replaced with averaging over the simpler distribution in which the entries are independent Bernoullis, up to a negligible relative correction:*

$$\mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim \mathbb{P}_{\psi,2}} \mathbb{P}_{\psi,2}[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)}] = \mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim \text{Ber}(1/4)^n} \mathbb{P}_{\psi,2}[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)}] \left(1 + O(2^{-n/2})\right) \quad (82)$$

$$= q_p \left(1 + O(2^{-n/2})\right), \quad (83)$$

where we define the useful quantity:

$$q_p \equiv 4 \left(\frac{2 + 2^{-p}}{4}\right)^n. \quad (84)$$

Proof. Taking a fixed $\mathbf{z} = \mathbf{z}_1^C \mathbf{z}_2^{\bar{C}} \in \mathbb{Z}_2^n$ of even weight on both sides of the cut (i.e. $|\mathbf{z}_1|, |\mathbf{z}_2|$ even), consider the quantity:

$$\mathbb{E}_{\mathbf{y} \sim \mathbb{P}_{\psi,2}} 3^{-|\mathbf{y} \oplus \mathbf{z}|} = \sum_{0 \leq d_1, d_2 \leq n/2} \sum_{\substack{0 \leq b \leq |\mathbf{z}_1| \\ 0 \leq b \leq |\mathbf{z}_2|}} \mathbb{P}_{\mathbf{y}_1^C \mathbf{y}_2^{\bar{C}}} [|\mathbf{y}_1| = 2b_1 + d_1 - |\mathbf{z}_1| \cap |\mathbf{y}_2| = 2b_2 + d_2 - |\mathbf{z}_2|] 3^{-d_1 - d_2}, \quad (85)$$

where we define $d_1 = |\mathbf{y}_1 \oplus \mathbf{z}_1|$ and $b_1 = |\mathbf{y}_1 \cap \mathbf{z}_1| = |\{i \in [n/2] : y_{1,i} = z_{1,i} = 1\}|$ (and similarly for d_2, b_2). Notice from (81) that the distribution over \mathbf{y} completely factorizes over the two sides of the cut \mathbf{y}_1 and \mathbf{y}_2 , and so does the above sum, so it is enough to study only one side of the cut. We can therefore express explicitly, using (81):

$$\sum_{\substack{0 \leq d_1 \leq n/2 \\ 0 \leq b \leq |\mathbf{z}_1|}} \mathbb{P}_{\mathbf{y}_1} [|\mathbf{y}_1| = 2b_1 + d_1 - |\mathbf{z}_1|] 3^{-d_1} = 2 \frac{3^{n/2 - |\mathbf{z}_1|}}{2^{n/2} (2^{n/2} + 1)} \sum_{\substack{0 \leq d_1 \leq n/2 \text{ even} \\ 0 \leq b_1 \leq |\mathbf{z}_1|}} \binom{|\mathbf{z}_1|}{b_1} \binom{\frac{n}{2} - |\mathbf{z}_1|}{d_1 - |\mathbf{z}_1| + b_1} 3^{-2b_1}. \quad (86)$$

There are two observations. First, since $|\mathbf{y}_1| = 2b_1 + d_1 - |\mathbf{z}_1|$ has to be even, and since we assume $|\mathbf{z}_1|$ is also even, then this restricts d_1 to only take even values. Second, the d_1 in the exponent

cancels out, and the sum over d_1 becomes easy, since it is exactly half of a sum over a full set of binomial coefficients, leading to:

$$\sum_{\substack{0 \leq d_1 \leq n/2 \\ 0 \leq b \leq |z_1|}} \mathbb{P}_{\mathbf{y}^1} [|\mathbf{y}_1| = 2b_1 + d_1 - |z_1|] 3^{-d_1} = \frac{3^{n/2 - |z_1|} 2^{-|z_1|}}{2^{n/2} (2^{n/2} + 1)} \sum_{0 \leq b_1 \leq |z_1|} \binom{|z_1|}{b_1} 3^{-2b_1} \quad (87)$$

$$= \frac{3^{n/2}}{2^{n/2} (2^{n/2} + 1)} \left(\frac{5}{27} \right)^{|z_1|}. \quad (88)$$

It is helpful to notice that this is negligibly close to the result obtained when \mathbf{y} is sampled not from the $\mathbb{P}_{\text{StateHSP}_{\psi,2}}$ distribution, but from the much simpler distribution in which each entry is i.i.d. sampled from a $\text{Ber}(1/4)$ distribution, i.e. removing the parity constraints. In that case, the result is:

$$\sum_{\substack{0 \leq d_1 \leq n/2 \\ 0 \leq b \leq |z_1|}} \mathbb{P}_{\mathbf{y}^1 \sim \text{Ber}(1/4)^{n/2}} [|\mathbf{y}_1| = 2b_1 + d_1 - |z_1|] 3^{-d_1} = \left(\frac{3}{4} \right)^{n/2} \left(\frac{5}{27} \right)^{|z_1|}. \quad (89)$$

This means that:

$$\mathbb{E}_{\mathbf{y} \sim \mathbb{P}_{\text{StateHSP}_{\psi,2}}} 3^{-|\mathbf{y} \oplus \mathbf{z}|} = \mathbb{E}_{\mathbf{y} \sim \text{Ber}(1/4)^n} 3^{-|\mathbf{y} \oplus \mathbf{z}|} \left(1 + O(2^{-n/2}) \right). \quad (90)$$

We can apply this result to estimate the more useful quantity:

$$\mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim \mathbb{P}_{\text{StateHSP}_{\psi,2}}} 3^{-|\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)}|}. \quad (91)$$

Replacing the $\mathbb{P}_{\text{StateHSP}_{\psi,2}}$ distribution with independent entrywise Bernoullis simplifies the calculation significantly, since:

$$\left| \mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right| = \sum_{i=1}^n \mathbf{1} \left[y_i^{(1)} + \dots + y_i^{(p)} = 1 \pmod{2} \right]. \quad (92)$$

This means that, with the unrestricted Bernoulli distribution, the average from above can be factored across the indices $\{1, \dots, n\}$, leading to the closed-form expression:

$$\mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim \text{Ber}(1/4)^n} 3^{-|\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)}|} = \left(\mathbb{E}_{\mathbf{u} \sim \text{Ber}(1/4)^p} 3^{-(\mathbf{u} \cdot 1^p \pmod{2})} \right)^n \quad (93)$$

$$= \left(\frac{3}{4} \right)^{pn} \left(\sum_{\substack{0 \leq u \leq p \\ u \text{ even}}} \binom{p}{u} 3^{-u} + \sum_{\substack{0 \leq u \leq p \\ u \text{ odd}}} \binom{p}{u} 3^{-u-1} \right)^n \quad (94)$$

$$= \left(\frac{2 + 2^{-p}}{3} \right)^n \quad (95)$$

Equivalently, let us define:

$$\mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim \mathbb{P}_{\text{StateHSP}_{\psi,2}}} \mathbb{P}_{\text{StateHSP}_{\psi,2}} \left[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right] = 4 \left(\frac{2 + 2^{-p}}{4} \right)^n \left(1 + O(2^{-n/2}) \right) \quad (96)$$

$$= q_p \left(1 + O(2^{-n/2})\right). \quad (97)$$

□

Second, let us prove a lower bound on the relevant quantity we are aiming to estimate:

Fact B.2. *Let us define the probability that k samples are linearly independent as vectors over \mathbb{Z}_2^n :*

$$\pi_k \equiv \mathbb{P}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k)} \sim P_{\psi,2}} \left[(\mathbf{y}^{(j)})_{j \in [k]} \text{ lin. indep.} \right], \quad (98)$$

where $\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k)}\} \subset H_C^\perp$ are understood to be k independent samples from the $P_{\psi,2}$ distribution (81). Then we have the lower bound:

$$\pi_k \geq 1 - \sum_{p=1}^k \binom{k}{p} q_{p-1} \left(1 + O(2^{-n/2})\right). \quad (99)$$

In particular, it follows that that $\pi_{n-3} \geq 1/2 + O(\text{negl}(n))$.

Proof. Explicitly write the probability π_k as the sum over linear independent combinations of vectors:

$$\pi_k = \sum_{\substack{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k)} \text{ lin.indep.}}} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k)} \right] \quad (100)$$

$$= \sum_{\substack{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \text{ lin.indep.}}} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k-1)} \right] \sum_{\substack{\mathbf{y}^{(k)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(k)} \notin \text{span}\{\mathbf{y}^{(j)}\}_{j=1}^{k-1}}} P_{\psi,2} \left[\mathbf{y}^{(k)} \right] \quad (101)$$

$$= \sum_{\substack{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \text{ lin.indep.}}} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k-1)} \right] \left(1 - \sum_{a \in \{0,1\}^{k-1}} P_{\psi,2} \left[a_1 \mathbf{y}^{(1)} \oplus \dots \oplus a_{k-1} \mathbf{y}^{(k-1)} \right] \right) \quad (102)$$

$$= \pi_{k-1} - \sum_{p=0}^{k-1} \binom{k-1}{p} \sum_{\substack{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \text{ lin.indep.}}} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k-1)} \right] P_{\psi,2} \left[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right] \quad (103)$$

$$\geq \pi_{k-1} - \sum_{p=0}^{k-1} \binom{k-1}{p} q_p \left(1 + O(2^{-n/2})\right). \quad (104)$$

The fourth line comes from symmetry, and the final line comes from bounding:

$$\sum_{\substack{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \in \mathbb{Z}_2^n \\ \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \text{ lin.indep.}}} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k-1)} \right] P_{\psi,2} \left[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right] \quad (105)$$

$$\leq \sum_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)} \in \mathbb{Z}_2^n} P_{\psi,2} \left[\mathbf{y}^{(1)} \right] \dots P_{\psi,2} \left[\mathbf{y}^{(k-1)} \right] P_{\psi,2} \left[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right] \quad (106)$$

$$= \mathbb{E}_{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(p)} \sim P_{\psi, 2}} \mathbb{P}_{\psi, 2} \left[\mathbf{y}^{(1)} \oplus \dots \oplus \mathbf{y}^{(p)} \right] \quad (107)$$

$$= q_p \left(1 + O(2^{-n/2}) \right). \quad (\text{using Fact B.1}) \quad (108)$$

Therefore, unfolding the recursion (104) down to the base $\pi_0 = 1$, we get the desired lower bound:

$$\pi_k \geq 1 - \sum_{\ell=0}^{k-1} \sum_{p=0}^{\ell} \binom{\ell}{p} q_p \left(1 + O(2^{-n/2}) \right) \quad (109)$$

$$= 1 - \sum_{p=0}^{k-1} \binom{k}{p+1} q_p \left(1 + O(2^{-n/2}) \right). \quad (110)$$

Given the explicit form for the q_p quantities derived previously in (84), we have that the choice of $k = \Theta(n)$ means the above lower bound for π_k is well-approximated by taking $q_p \approx 2^{2-p}$. Specifically, assume $k = \Theta(n)$, and upper bound the relevant sum as:

$$\sum_{p=0}^{k-1} \binom{k}{p+1} q_p = 2^{2-n} \sum_{p=1}^k \binom{k}{p} (1 + 2^{-p})^n \quad (111)$$

$$\leq 2^{2-n} \left(\frac{3}{2} \right)^n \sum_{p=0}^{\lfloor \log^2(n) \rfloor} \binom{k}{p} + 2^{2-n} \left(1 + \frac{1}{2^{\lfloor \log^2(n) \rfloor}} \right)^n \sum_{p=\lfloor \log^2(n) \rfloor + 1}^k \binom{k}{p}, \quad (112)$$

where we have simply split the sum at an appropriate term of order polylogarithmic in n . The first contribution can be upper bounded by a typical Höfding tail bound of the binomial distribution, and the second contribution can be bounded by a standard binomial sum. This results in:

$$\sum_{p=0}^{k-1} \binom{k}{p+1} q_p \leq 2^{2+k-n} \left[\exp\left(-\frac{k}{2} + n \ln \frac{3}{2} + O(\text{polylog}(n))\right) + \left(1 + \frac{1}{2^{\lfloor \log^2(n) \rfloor}} \right)^n \right]. \quad (113)$$

The second term above is negligibly close to one. Also, when $k = n - b$ for a constant $b > 0$, the first term is exponentially decaying in n , since $\frac{1}{2} - \ln \frac{3}{2} \approx 0.0945 > 0$. This means that in this case, the lower bound derived above (110) becomes:

$$\pi_{n-b} \geq 1 - 2^{2-b} (1 + \text{negl}(n)), \quad (114)$$

which in particular means that $\pi_{n-3} \geq 1/2$ up to a negligible correction. \square

We have the necessary ingredients to assemble the proof of [Theorem 2](#):

Proof of Theorem 2. In the main text, we have shown in [Fact 5.1](#) that, with high probability over the Haar-random factor states, purity Fourier sampling with two copies of the input state with a hidden cut yields samples from the $P_{\psi, 2}$ probability distribution (81). This distribution is supported (non-uniformly) inside the $(n-2)$ -dimensional hidden cut subspace H_C^\perp , which is defined in (31) as the subspace orthogonal to the equivalent cut strings $1^C 0^{\bar{C}}$ and $0^C 1^{\bar{C}}$:

$$H_C^\perp = \left\{ \mathbf{y} \in \mathbb{Z}_2^n : \mathbf{y} \cdot 1^C 0^{\bar{C}} = \mathbf{y} \cdot 0^C 1^{\bar{C}} = 0 \pmod{2} \right\}. \quad (115)$$

[Fact B.2](#) proven above gives us that, with probability at least one half (up to negligible corrections), $n-3$ i.i.d. samples $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n-3)}$ from the $P_{\psi, 2}$ distribution will be linearly independent as

vectors in \mathbb{Z}_2^n . Assembling these samples as the rows of a Boolean matrix $Y = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n-3)})^T \in \mathbb{Z}_2^{(n-3) \times n}$, we have that with probability at least one half the nullspace of Y is the three-dimensional Boolean subspace:

$$\text{nullspace } Y = \text{span} \left\{ 1^C 0^{\bar{C}}, 1^n, \mathbf{b} \right\}, \quad (116)$$

spanned by the two equivalent cut strings and an additional arbitrary vector $\mathbf{b} \in \mathbb{Z}_2^n$. Given the samples, the nullspace can be determined in polynomial time by simple Boolean linear algebra methods. Excluding the trivial vectors 0^n and 1^n , finding the correct nullspace therefore is equivalent to narrowing down the possible hidden cuts to six non-trivial candidates. Each of these candidates can be individually tested up to a constant confidence interval with a constant number of copies of the input state by standard single-cut SWAP tests. In other words, a constant number of copies is required at the end to find the final vector which completes the basis for the cut subspace. This suffices to show that the correct hidden cut can be found with $O(n)$ total number of copies with constant probability, with only two copies at a time consumed either by the purity Fourier sampling or by the final standard SWAP tests. \square

Remark. Numerical evidence strongly indicates that a complete basis for the cut subspace H_C^\perp can be obtained with constant probability directly from $n - 2$ i.i.d. samples from the $P_{\psi,2}$ distribution (81). This means that, in practice, the standard Simon's protocol (for example, as used in Theorem 1) will still return a correct answer when applied to this non-uniform distribution. This does not change the $O(n)$ requirement in terms of number of copies, but would remove the need for additional SWAP tests. However, proving this version of the result would likely require a more involved Boolean random matrix analysis, which we leave to future work.