Efficient Quantum Pseudorandomness from Hamiltonian Phase States

John Bostanci[1], Jonas Haferkamp[2], Dominik Hangleiter[3,4], and Alexander Poremba[5]

[1]Columbia University
[2]Harvard University
[3]QuICS, University of Maryland & NIST
[4]Simons Institute for the Theory of Computing, UC Berkeley
[5]Massachusetts Institute of Technology

Abstract

Quantum pseudorandomness has found applications in many areas of quantum information, ranging from entanglement theory, to models of scrambling phenomena in chaotic quantum systems, and, more recently, in the foundations of quantum cryptography. Kretschmer (TQC '21) showed that both pseudorandom states and pseudorandom unitaries exist even in a world without classical one-way functions. To this day, however, all known constructions require classical cryptographic building blocks which are themselves synonymous with the existence of one-way functions, and which are also challenging to implement on realistic quantum hardware.

In this work, we seek to make progress on both of these fronts simultaneously—by decoupling quantum pseudorandomness from classical cryptography altogether. We introduce a quantum hardness assumption called the *Hamiltonian Phase State* (HPS) problem, which is the task of decoding output states of a random instantaneous quantum polynomial-time (IQP) circuit. Hamiltonian phase states can be generated very efficiently using only Hadamard gates, single-qubit Z rotations and CNOT circuits. We show that the hardness of our problem reduces to a worst-case version of the problem, and we provide evidence that our assumption is plausibly *fully quantum*; meaning, it cannot be used to construct one-way functions. We also show information-theoretic hardness when only few copies of HPS are available by proving an approximate t-design property of our ensemble. Finally, we show that our HPS assumption and its variants allow us to *efficiently* construct many pseudorandom quantum primitives, ranging from pseudorandom states, to quantum pseudoentanglement, to pseudorandom unitaries, and even primitives such as public-key encryption with quantum keys. Along the way, we analyze a natural iterative construction of pseudorandom unitaries which resembles a candidate of Ji, Liu, and Song (CRYPTO'18).

1

Contents

1 Introduction

Pseudorandomness [Gol06, Vad12] is ubiquitous in theoretical computer science and has found applications in many areas, ranging from cryptography, to computational complexity, to the study of randomized algorithms, and even to combinatorics. The celebrated result of Håstad, Impagliazzo, Levin, and Luby [HILL99] shows that one can construct a *pseudorandom generator* from any one-way function—a function that is easy to evaluate but computationally hard to invert. Pseudorandom generators can then in turn be used to construct more advanced cryptographic primitives, such as *pseudorandom functions* [GGM86], i.e., keyed families of functions that appear random to any computationally bounded observer. This fact has elevated the notion of a one-way function as the minimal assumption in all of theoretical cryptography. One-way functions are typically built from well-studied mathematical conjectures, such as the hardness of factoring [RSA78] and discrete logarithms [Mer78], decoding error correcting codes [BFKL94, Ale03], or finding short vectors in high-dimensional lattices [Reg09]. More advanced cryptographic primitives (which are believed to lie beyond what is generically possible to construct from any one-way function), such as public-key encryption, tend to require highly structured assumptions which are more susceptible to algorithmic attacks—particularly by quantum computers [Sho97], which has led to the design of *post-quantum assumptions* [ACD+22].

In quantum cryptography, there has recently been a significant interest in so-called "fully quantum" cryptographic primitives (occasionally referred to as *MicroCrypt* primitives) which are potentially *weaker* than the conventional minimal assumptions used in classical cryptography. Here, the notion of *quantum pseudorandomness* has emerged as the natural quantum analogue of pseudorandomness in the classical world [JLS18, Kre21, ABF+24]. In particular, Ji, Liu and Song [JLS18] proposed the notion of pseudorandom states [JLS18] and pseudorandom unitaries as the natural quantum analogues of pseudorandom generators [HILL99] and pseudorandom functions [GGM86], respectively. The work of Kretschmer [Kre21, KQST23] has shown that such fully quantum cryptographic primitives can exist in a world in which no classical cryptography exists—including one-way functions. At the same time, quantum pseudorandomness has applications in many areas of quantum information, ranging from entanglement theory [ABF+24, BFG+24, FI24], quantum learning theory [ZLK+23], to models of scrambling phenomena in chaotic quantum systems [KP23, EFL+24], and, more generally, even in the foundations of quantum cryptography [JLS18, Kre21, KQST23, MY22, AGQY22, BCQ22, BEM+23, KT24a, BJ24].

Limitations of existing constructions. Despite strong evidence that MicroCrypt primitives such as pseudorandom states and pseudorandom unitaries lie "below" one-way functions [Kre21, KQST23], all known constructions implicitly make use of one-way functions (or other assumptions which are themselves synonymous with the existence of one-way functions) [JLS18, BS19, MPSY24]. This begs the question:

Is it possible to construct fully quantum primitives, including quantum pseudorandomness, from quantum rather than classical hardness assumptions?

Instantiating fully quantum primitives from a concrete and well-founded quantum hardness assumption (rather than from the existence of one-way functions) has remained a long standing open problem [AQY22, MY22].

Moreover, the fact that quantum pseudorandom states and unitaries are built from classical one-way functions makes them nearly impossible to realize on realistic quantum hardware. In some sense, this is inherent because cryptographic pseudorandom functions are highly complex by design [BMM24], and therefore require a massive computational overhead to implement coherently. As a result, this severely limits the potential of using quantum pseudorandomness in practical applications; for example in the context of entanglement theory [ABF+24, BFG+24], or when studying the emergence of thermal equilibria in isolated many-body systems [FI24], or when modeling scrambling phenomena in chaotic quantum systems [KP23]. A second limitation of existing pseudorandom constructions is therefore also the notion of quantum efficiency, which begs the question:

Are there more efficient constructions of quantum pseudorandomness which can be implemented on realistic quantum hardware?

Making progress on both of these questions would not only lead to new insights in the foundations of quantum cryptography and the study of quantum hardness assumptions more generally, but also make quantum pseudorandomness more useful in practice. To this day, however, no concrete fully quantum hardness assumption has been explored in an attempt to answer this question.

Towards a fully quantum assumption. In order to plausibly claim that quantum pseudorandomness and other fully quantum cryptographic primitives exist in a world in which classical cryptography does not, we must construct these primitives from new assumptions that do not themselves imply classical cryptography.

The history of cryptography has taught us that finding good and well-founded cryptographic assumptions is not at all an easy task—even entirely plausible assumptions have often found surprising attacks [Sho94, CD22, Beu22]. What makes a new cryptographic assumption reasonable? While no widely agreed upon standards exist [GK15], the conventional belief is to use assumptions

- which are rooted in a well-studied problem (ideally, a problem that has already been analyzed for many years) and which seems intractable in the worst case;

- for which there is a natural notion of what constitutes a "random instance" of the problem; moreover, such an instance can always be efficiently generated;

- for which there is evidence of average-case hardness, ideally in the form of a worst-case to average-case reduction;

- which can be connected to other assumptions or computational tasks that have been studied over the years, and

- which have enough structure to enable interesting cryptographic primitives.

A natural candidate for constructing quantum pseudorandomness (and other fully quantum cryptographic primitives) is via *random quantum circuits*. In fact, the computational pseudorandomness of random quantum circuits appears to be a folklore conjecture and is widely believed among many quantum computer scientists. As we are unaware of a concrete technical conjecture, we provide such a formulation here.

Conjecture 1 (Random quantum circuits give rise to pseudorandom unitaries)**.**
Consider n-qubit random quantum circuits with m gates defined by repeating the following process m times independently at random: Draw a random pair (i, j) of qubits and apply a gate from a universal gate set $\mathsf{G} \subset SU(4)$ to the qubits i and j. Then, there exist univeral constants $c > 0$ and $C_{\mathsf{G}} > 0$ (depending on the gate set G) such that random quantum circuits with $m \geq C_{\mathsf{G}} n^c$ gates form ensembles of pseudorandom unitaries.

We note that many other possible formulations (e.g. with specific geometric architectures or only regarding pseudorandom states) are also possible. Indeed, if Conjecture 1 holds even with exponential security, then Ref. [SHH24] implies that a simple ensemble of random quantum circuits in a 1D architecture of depth $\mathrm{polylog}(n)$ is also pseudorandom.

Conjecture 1 can be seen as a direct quantum analogue of a claim that was first proposed by Gowers [Gow96] who conjectured that random reversible quantum circuits form pseudorandom permutations on bitstrings. This conjecture has inspired multiple recent works in classical cryptography. For instance, it was recently proven by He and O'Donnell [HO24] that the Luby-Rackoff [LR88] construction of pseudorandom permutations from pseudorandom functions can be implemented with reversible permutations. Random reversible circuits have recently also inspired entirely new approaches for constructing program obfuscation schemes [CCMR24].

Gowers originally conjectured the emergence of pseudorandomness when attempting to prove that random quantum circuits converge quickly to ensembles of *t*-wise independent permutations [Gow96] (this bound was further improved later on towards an optimal scaling [HMMR05, BH08, CHH⁺24]. In fact, this property can itself be viewed as evidence for pseudorandomness as we further discuss in Section 5. It turns out that random quantum circuits satisfy an analogous property by converging nearly optimally towards approximate *t*-designs [CHH⁺24, BHH16, Haf22, SHH24].

However, we currently do not have rigorous evidence for Conjecture 1; for example, in terms of a worst-to-average reduction for a corresponding learning problem. Moreover, and maybe more importantly, it is unclear how one would use unstructured random quantum circuits to construct more advanced quantum cryptographic primitives (e.g., as discussed in Section 2.3). A similar situation arises for general one-way functions, which require additional structure to build more advanced cryptographic applications, such as public-key encryption. It could very well be the case that random quantum circuits are simply *too mixing* to be a useful in the context of quantum cryptography. A natural way forward is to search for a sweet spot—an ensemble of random quantum circuits that is sufficiently structured to permit the construction of interesting cryptographic primitives but which, at the same time, is sufficiently mixing to guarantee security.

2 Our contributions

In this work, we simultaneously address the two major open problems in the field of quantum pseudorandomness and propose the first well-founded and fully quantum hardness assumption. To this end, we follow the strategy sketched above, and propose a family of quantum states which we call *Hamiltonian Phase States*. These states are a family of quantum states which are "maximally quantum" in the sense that the state has support on all bitstrings with amplitudes equal in magnitude, but varying phases. Hamiltonian Phase States are generated by a family of commuting *instantaneous quantum polynomial-time* (IQP) circuits which generalize the X programs proposed by Shepherd and Bremner [SB09]. The corresponding circuits are highly structured in that they are generated by a Hamiltonian with only Z-type terms applied to the all-$|+\rangle$ state. This structure makes them amenable to rigorous analysis [Kah23, BCJ23, GH23]. At the same time, these circuits are also believed to be sufficiently mixing and hard to simulate classically [SB09, BJS10, BMS16, HE23]. Moreover, since Hamiltonian phase states can be generated by a commuting Hamiltonian, they admit a highly efficient implementation in practice.

Phase states are a natural direction to look at in the search for a fully-quantum cryptographic assumption with sufficient amounts of structure. On the one hand, this is because of their quantum advantage properties. On the other hand, the (quantum) learnability of different ensembles of phase states has been studied extensively in recent work [ABDY23]. There, the authors give optimal bounds for the sample complexity of learning many families of phase states from quantum samples, as well as upper bounds on the time complexity. Importantly, there are families of phase states generated using a small number of (long-range) gates, which cannot be learned from polynomially many samples. Following this, our proposed cryptographic assumption is that Hamiltonian Phase States are hard to learn, given quantum samples and classical side information.

Moreover, the known constructions for pseudorandom states with useful cryptographic applications are based on phase states [JLS18]. These are generated using a single-bit output quantum-secure pseudorandom function family $\{f_k\}_k$ with

$$|\phi_k\rangle \propto \sum_x \omega_q^{f_k(x)} |x\rangle, \tag{2.1}$$

where ω_q is a q-th root of unity, for example $q = 2$ [BS19]. Because these states are based on a classical assumption, they require the reversible implementation of a classical PRF which requires a large number of Toffoli gates. These are extremely expensive in standard fault-tolerant constructions. However, the results of Refs. [JLS18, ABDY23] suggest that a more natural family of phase states which is generated by a quantum circuit with a small number of expensive gates can also yield quantum pseudorandomness. This would require gates affecting a large number of qubits, since low-degree phase states can be learned efficiently. As we show below, in spite of having terms with high support, the Hamiltonian Phase States can be generated highly efficiently using only local Z-rotations and CNOT gates.
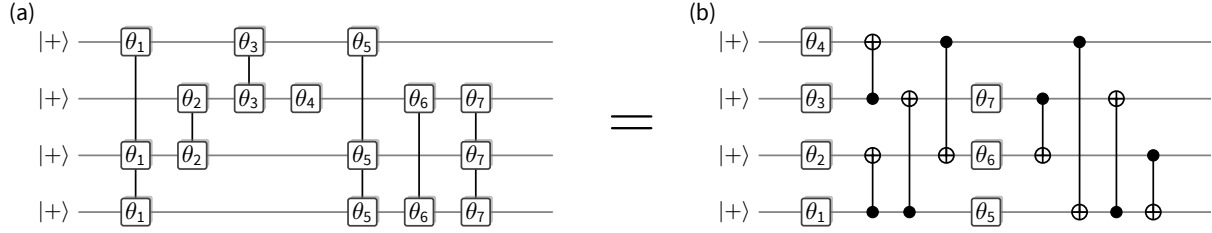
Figure 1: Hamiltonian Phase States (HPS) are generated by sequentially applying Ising-type rotations around angles θ_i to the state $|+^n\rangle = H^{\otimes n}|0^n\rangle$. (a) Example of a HPS on 4 qubits. Connected boxes at sites i, j, k with angle θ represent the unitary $\exp(i\theta Z_i Z_j Z_k)$. (b) HPS can be implemented using only single-qubit Z rotations interlaced with CNOT circuits.

2.1 Hamiltonian Phase States

Let $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ be a binary matrix and let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$ be a set of uniformly random angles in the interval $[0, 2\pi)$ according to some discretization into $q = \text{poly}(n)$ parts. We consider phase states of the form

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}}\right) H^{\otimes n}|0^n\rangle. \tag{2.2}$$

where, for $i \in [m]$, we denote the i-th row of $\mathbf{A}$ by $(\mathbf{A}_{i1}, \ldots, \mathbf{A}_{in})$ and let

$$\bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} = Z^{\mathbf{A}_{i1}} \otimes \cdots \otimes Z^{\mathbf{A}_{in}} \qquad \text{for} \qquad Z^0 = \mathbb{I}, \ Z^1 = Z.$$

We call these states *Hamiltonian Phase States* since they can naturally be prepared as the result of a time evolution under a sparse Ising Hamiltonian. We also call the matrix $\mathbf{A}$ the *architecture* of the states, as it specifies the overall structure/location of the Ising terms. Hamiltonian Phase States with a single fixed angle $\theta_i \equiv \theta$ have been studied as a means to demonstrate verified quantum advantage, when measured in the X basis, under the name X-programs [SB09]. A Hamiltonian Phase State is therefore a generalized version of an X program parameterized by the pair $(\mathbf{A}, \boldsymbol{\theta})$. X programs with $\theta = \pi/8$ have the interesting property that its Fourier coefficients can be computed efficiently classically, but at the same time the simulation of such X-programs is believed to be classically intractable [SB09, BJS10, BMS16, HE23].

Our cryptographic assumption rests on the apparent hardness of *learning* Hamiltonian Phase States (or generalized X programs), which was highlighted in recent work [ABDY23]. Concretely, our quantum computational assumption amounts to the conjecture that our ensemble of Hamiltonian phase states satisfies the following two properties:

- Random Hamiltonian Phase States are *hard to invert* in the following sense: given $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}$, for any $t = \mathsf{poly}(n)$, it is computationally difficult to reverse-engineer the angles $\boldsymbol{\theta}$ and architecture $\mathbf{A}$. This means that the ensemble $\{|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle\}_{\boldsymbol{\theta},\mathbf{A}}$ gives rise to a so-called *one-way state generator* (OWSG) as in Definition 8.

- Random Hamiltonian Phase States are *hard to distinguish from Haar random states* in the following sense: given $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}$, for any $t = \mathsf{poly}(n)$, it is computationally difficult to distinguish $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}$ from $|\Psi\rangle^{\otimes t}$, where $|\Psi\rangle$ is a Haar random state. This means that the ensemble $\{|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle\}_{\boldsymbol{\theta},\mathbf{A}}$ gives rise to a so-called *pseudorandom state generator* (PRSG) as in Definition 9.

We can call the two assumptions above the search (respectively, decision) variant of *Hamiltonian Phase State* assumption ($\mathsf{HPS}_{n,m,q,\chi}$). Here, $n, m \in \mathbb{N}$ are circuit parameters, q is a discretization parameter for the interval $[0, 2\pi)$, and χ is a distribution over the choice of matrix $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$; typically, χ is chosen to be the uniform distribution.

There is some evidence that $\mathsf{HPS}_{n,m,q,\chi}$ is a reasonable assumption for constructing pseudorandom states. Brakerski and Shmueli [BS19] show that the states $DH^{\otimes n}|0^n\rangle$ form a state t-design when the diagonal operator D consists of a $2t$-wise independent binary phase operator. Previously, Nakata, Koashi and Murao [NKM14] also showed that the states $DH^{\otimes n}|0^n\rangle$, where D is a diagonal operator composed of appropriate diagonal gates with random phases, form a t-design. Starting from this intuition, we now provide rigorous evidence for the hardness of the $\mathsf{HPS}_{n,m,q,\chi}$ assumption.

2.2 Overview of our Results

In this work, we establish $\mathsf{HPS}_{n,m,q,\chi}$ as a well-founded quantum computational assumption. Specifically, we address each of the meta-criteria we mentioned before:

- (Evidence of worst-case hardness:) The learnability of ensembles of phase states has been studied extensively in recent work [ABDY23], and has been found to have exponential time complexity in the worst case (despite only having polynomial sample complexity). We give a detailed discussion in Section 5.3.

- (Notion of a random instance:) A random Hamiltonian phase state $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$, e.g., as in Eq. (2.2), is naturally defined in terms of a random binary matrix $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ and a random set of angles $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$. Hence, it can be efficiently generated by a simple quantum circuit comprising $O(m/n \cdot n^2)$ CNOT gates, n Hadamard gates, and m single-qubit Z rotations.

- (Evidence of average-case hardness:) Our learning task admits a worst-case to average-case reduction. In Section 5.1, we separately show how to re-randomize the architecture and the set of angles. Therefore, the hardness of our problem reduces to a worst-case version of the problem.

8

- (Relation to other problems) In Section 5.1, we draw a connection between the task of learning Hamiltonian Phase states and the security of classical *Goppa codes* and the well-known McEliece cryptosystem.

- (Cryptographic applications:) Hamiltonian Phase states have a sufficient amount of structure which suffices to construct a number of interesting cryptographic primitives which we sketch in detail in Section 2.3.

In Section 5.4, we also provide evidence that $\text{HPS}_{n,m,q,\chi}$ is plausibly fully-quantum and does not allow one to construct one-way functions. In particular, we note that the result of [Kre21] indicates that the idealized versions of any assumption that yields *only* pseudorandom states can not be used to build one-way functions in a black-box way. We further discuss the implications of the fact that HPS states are state t-designs on this reduction, noting that the resulting concentration properties by themselves rule out one-way function constructions that do not simulateneously measure many copies of the HPS state. For all of the primitives we construct, in addition to just constructing these primitives from our hardness assumption, we argue that constructing them from our hardness assumption yields more efficient and practical implementations of these primitives (if and when fault-tolerant quantum computers become widely available).

2.3 Applications

In this section, we give an overview of all the applications which are enabled by the HPS assumption. Besides the natural application of constructing efficient one-way state generators (see Definition 8) and pseudorandom state generators (see Definition 9), which essentially follow by definition of our assumption, we also construct a number of other interesting applications that are relevant in quantum information science more broadly.

Quantum Trapdoor Functions and Public-Key Encryption with Quantum Public Keys.
Recent work of Coladangelo [Col23] introduced the notion of a quantum trapdoor function (QTF). This primitive is essentially a variant of a one-way state generator that also features a secret trapdoor which makes inversion possible. QTFs are interesting in the sense that they *almost* enable public-key encryption: two parties can communicate classical messages over a quantum channel without ever exchanging a shared key in advance—the only caveat being that this requires the public keys to be quantum states [Col23]. Using a construction based on binary-phase states, Coladangelo [Col23] showed that quantum trapdoor functions exist, if post-quantum one-way functions exist. However, to this day, it remains unclear how to construct QTFs from assumptions which are potentially weaker than one-way functions, such as the existence of pseudorandom states.

In Section 6.3, we show how to construct QTFs from our (decisional) HPS assumption, which yields the first construction of QTFs from an assumption which is plausibly weaker than that of one-way functions. We believe that this application strongly highlights the

versatility of Hamiltonian Phase states in the context of quantum cryptography; for example, it is far less clear how to construct QTFs from other, less structured, assumptions such as genuinely random quantum circuits via Conjecture 1.

Quantum Pseudoentanglement. The notion of pseudoentanglement [ABF+24, BFG+24] has found many applications in quantum physics, for example to study the emergence of thermal equilibria in isolated many-body systems [FI24]. Pseudoentangled states have also been viewed as a potential tool for probing computational aspects of the AdS/CFT correspondence, which physicists believe may shed insight onto the behavior of black holes in certain simplified models of the universe. We note that it is currently not known how to construct these from *any* assumption other than one-way functions.

In Section 6.5 we give a construction of pseudoentangled states from our HPS assumption, which yields the first construction of pseudoentanglement from an assumption which is plausibly weaker than that of one-way functions. Our proof sheds new light on the entanglement properties of random IQP circuits more generally.[1] Therefore, we belief that this contribution is of independent interest. Moreover, as we point out in the next section, our construction is also highly efficient and could enable implementations of quantum pseudoentanglement in practical scenarios.

Pseudorandom Unitaries. Pseudorandom unitaries are families of unitaries that are indistinguishable from Haar random unitaries in the presence of computationally bounded adversaries. They are widely considered the most powerful fully-quantum primitive, and there has been a long line of work towards constructing them from the existence of one-way functions [AGQY22, BM24, MPSY24, CBB+24], eventually resulting in the most recent breakthrough result by Ma and Huang [MH24].

The result of [MH24] show that the ensemble of unitaries, colloquially known as the PFC-ensemble [MPSY24], form an approximation to a Haar random unitary. However, this construction is not well suited for the HPS assumption, which, in some sense, provides a pseudorandom *diagonal* unitary. Using the techniques introduced in [MH24], we show that another pseudorandom construction from [JLS18], the FHFHFC-ensemble is also an approximation to a Haar random unitary when both functions are taken to be uniformly random binary phase oracles. This resolves a conjecture proposed by Ji, Liu and Song [JLS18]. We further show how diagonal matrices can be used to replace the Clifford operation with more applications of a random diagonal matrix and Hadamard gates. With this result, we provide a plausible construction of efficient pseudorandom unitaries from the HPS assumption: alternating applications of HPS unitaries and Hadamards. See Section 6.6 for more details.

---

[1]To the best of our knowledge, such bounds for random IQP circuits were previously not known.

2.4 Physical Implementations

Hamiltonian Phase States with m terms on n qubits can be generated very efficiently compared to phase states constructed from pseudorandom functions: to prepare a HPS, we require only a layer of Hadamard gates, followed by $\lceil m/n \rceil$ alternating layers of single-qubit Z rotations and CNOT circuits. To see this, we observe two facts. First,

$$\mathsf{CNOT}_{k,l} e^{i\theta Z_l} = e^{i\theta Z_k Z_l} \mathsf{CNOT}_{k,l}, \tag{2.3}$$

where $\mathsf{CNOT}_{k,l}$ is controlled on qubit k and targeted on qubit l. Second,

$$\mathsf{CNOT}_{k,l} |+^n\rangle = |+^n\rangle. \tag{2.4}$$

More generally, if C is a circuit comprised of CNOT gates, then $C|x\rangle = |\mathbf{C} \cdot x\rangle$, where $\mathbf{C} \in \mathrm{GL}(n, \mathbb{Z}_2)$ is an invertible binary matrix. Given an HPS, decompose its architecture matrix $\mathbf{A} \in \mathbb{Z}_2^{n \times m}$ into $\lceil m/n \rceil$ submatrices of n rows (except for the last one), and suppose each of those submatrices has full rank. The HPS $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ can then be prepared as

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \left(C_{\lceil m/n \rceil} \prod_{i=(\lceil m/n \rceil - 1)n}^{m} e^{i\theta_i Z_i} \right) \cdots \left(C_1 \prod_{i=1}^{n} e^{i\theta_i Z_i} \right) |+^n\rangle \tag{2.5}$$

where the CNOT circuits C_k are chosen such that the first n rows of $\mathbf{A}$ are given by $\mathbf{C}_{\lceil m/n \rceil} \cdots \mathbf{C}_1$, the second n rows by $\mathbf{C}_{\lceil m/n \rceil} \cdots \mathbf{C}_2$, and so on, see Fig. 1 for an example. If the rank condition above is not satisfied, decompose $\mathbf{A}$ into the minimal number ℓ of submatrices with full rank, and proceed as above. The smallest meaningful example of such states—with n random, linearly independent terms—can thus be prepared using a single layer of rotations and a CNOT circuit. By the fact that $\mathrm{GL}(n, \mathbb{Z}_2)$ is a group, uniformly random architecture matrices $\mathbf{C}_i$ and phases θ_i generate a uniformly random HPS.

This protocol for the implementation of HPS is also interesting from an early fault-tolerance perspective, since there are quantum codes in which all required operations are transversal, yielding a highly efficient fault-tolerant implementation. To see this, consider a $q = 2^d$-fold discretization of the unit circle. Now, we observe that there are d-dimensional CSS codes with a transversal $Z^{1/2^{d-1}}$ gate such as the $[[2^d - 1, 1, 3]]$ simplex code [ZCCC07]. By the fact that they are CSS codes, they also admit a transversal CNOT gate between code blocks. This means that HPS can be prepared using transversal in-block $Z^{1/2^{d-1}}$ gates as well as inter-block CNOT gates, making them amenable to implementations in early fault-tolerant architectures such as reconfigurable atom arrays [BEG+24], or trapped ion processors [RBB+24, RAC+24] in which arbitrary inter-block connectivities can be achieved.

2.5 Related Work

Minimal assumptions in cryptography. The study of minimal assumptions in quantum cryptography has a long history, dating back to the exploration of the possibility

of unconditional cryptography in the 1980ies and 90ies [BBBV97, LC97, May01, LC99]. Modern research into minimal assumptions largely began with the work of [JLS18], who first proposed the idea of pseudo-random states and unitaries, which are classically samplable ensembles of pure states and unitaries that are indistinguishable from states and unitaries drawn from the Haar measure. They also gave a construction of pseudo-random states from post-quantum one-way functions. Since then, a number of works have proposed quantum analogs of minimal classical assumptions. [MY22] proposed the one-way state generator, a quantum analog of one-way functions. [BCQ22] then proposed the EFI pair as a minimal quantum assumption, showing that it was required for many "useful" primitives, and posing the question of finding a concrete assumption that implies EFI pairs. Since then, a number of works have provided oracle separations and reductions between these primitives [KT24a, BJ24, BN24, BMM+24], showing that unlike in classical cryptography, quantum cryptographic primitives seem to have distinct powers relative to each other.

Whille almost all quantum cryptographic primitives are implied by some form of classical cryptography, it was not clear whether quantum primitives were strictly weaker than their classical counterparts. This question was resolved in [Kre21], which gave a quantum oracle relative to which one-way functions do not exist, but pseudo-random unitaries do. Later this was extended to a classical oracle [KQST23], although relative to this oracle only EFI pairs exist.

Recently, there has been a large body of research constructing pseudo-random unitaries from one-way functions. [JLS18] conjectured several constructions, but were unable to prove their security from one-way functions. [MPSY24] proposed a new construction, the PFC ensemble, which was later proven to be a secure pseudo-random unitary family by [MH24]. The recent work of [MH24] uses a new path recording framework, which we use heavily in our work.

Even more recently, a work has come out motivating the existence of quantum cryptography from conjectures related to quantum supremacy. Namely, the work of [KT24b] argues that a sampling problem for which there exists a polynomial-time quantum sampler but no classical sampler that can approximate the output distribution (even to a very high error) implies one-way puzzles.

Concrete assumptions for quantum cryptography. A number of prior works have proposed assumptions that could be candidate concrete assumptions for building quantum cryptography. The most common approach to building quantum crptography is to assume that random polynomial-depth quantum circuits become pseudo-random unitary or state ensembles after a certain depth. The motivation to use random quantum circuits comes from their extensive study as related to quantum supremacy [BJS10, AC17, HE23], as well as from the theory of black holes [BS18, BCH+21], which posits that random circuits should have "scrambling" properties similar to those expected in black holes. [BFV19] proposes an even more bold foundation for quantum cryptography. By formalizing assumptions and conjectures from physics, such as "complexity = volume" and the

AdS-CFT correspondence, they argue that time evolution by a the Hamiltonian of a black hole in the CFT gives rise to pseudo-random states.

Another interesting approach to quantum cryptography is the result of [KKNY05]. The authors propose the hardness of distinguishing signed coset states of the symmetric group as a basis for cryptography. They show a number of interesting properties of these states, firstly they prove a worst-case to average case reduction. They reduce the hardness of this distinguishing task to the hardness of graph automorphism, which in turn is related to the hardness of the hidden subgroup problem on the symmetric group. Finally, they build a public-key cryptography system with quantum keys (both secret and private), but for which trusted set-up is needed.

t-designs and classical reversible circuits.   A common practice in cryptography used to justify the security of primitives is to prove an approximate t-wise independence property. In particular, a long line of works have steadily improved the bounds on the number of rounds of a block cipher are needed to be applied in order to get certain t-wise independence [LTV21, LPTV23]. This is a popular justification for cryptographic security because it rules out a large class of natural attacks, such as linear attacks (2-wise independence rules this out), and differential attacks (t-wise independence rules out $\log_2(t)$ differential attacks).

Recently, the study of random reversible circuits has led to a number of exciting results pertaining to pseudorandomness [HO24] and obfucation [CCMR24]. These results, motivated by early work on the t-wise independence of bounded depth reversible circuite [Gow96, HMMR05, BH08, CHH+24, GHP24, HO24], have used the structure of random circuits, combined with an assumption that they are pseudo-random functions when given random inputs, to build indistinguishability obfuscators.

Here, we show the approximate state t-design property for HPS in Section 5.2 with $O(nt^2)$ random Z-rotations. Enormous effort went into proving the design property over various ensembles of circuits. Notably, improving over previous bounds [BHH16, Haf22], a near optimal bound (in t) of $\mathrm{poly}(n)t$ was found in Refs. [CHH+24]. Previously, other constructions in depth $\mathrm{poly}(n)t$ where found in Refs. [CBB+24, MPSY24] for the weaker notion of additive error approximate designs. We note that, while we do not achieve a linear scaling in t, our bound comes with very realistic constants and a surprisingly simple proof based on the analysis of Boolean functions. Surprisingly, quantum circuits allow for a further, exponential improvement in the n-dependence, which is provably impossible for reversible circuits [SHH24]. More concretely, Ref. [SHH24] shows that approximate unitary t-designs can be generated in depth $\tilde{O}(t \log(n))$, inheriting the t-dependence from earlier work [CHH+24]. Last, very similar bounds of $O(nt^2)$ "Pauli rotations" was proven for the approximate unitary t-design property in Ref. [HLT24] by a non-trivial generalization of the argument presented here.

Note: The results on the approximate state t-design property presented in Section 5.2 are contained in an earlier preprint [Haf23], which this paper supersedes.

2.6 Concurrent Work

Here we discuss concurrent work on building fully-quantum cryptography and the connections to this work.

The Hardness of Quantum Learning and its Cryptographic Applications. Ref. [FGSY24] proposes a pair of computational assumptions related to random circuits and construct quantum cryptography from them. Specifically, they conjecture that the problem of learning the description of a random circuit of depth $\log^2(n)$, from copies of its output state, and the problem of cloning the output of a random circuit, are computationally hard problems. The authors motivate this problem in a idealized model where the output state is truly Haar random, and subsequently build a number of cryptographic primitives from it. We note that while our suite of results seem similar, there are large differences in both the assumptions themselves: random brickwork circuits versus Hamiltonian phase states, how we motivate the hardness of our problems: a query lower bound in an idealize model versus worst-to-average case reduction, analysis of design properties and analysis of best known algorithms, and the kinds of cryptography we build from it: OWSG and primitives implied by OWSG as well as noise-resistant versions of these primitives, compared our to public key cryptography and pseudorandom unitaries. Both works are interesting explorations into how fully-quantum cryptography might be realized in practice.

Quantum cryptography from #P-hard problems. The recent and concurrent work of Khurana and Tomer [KT24b] argues that one-way puzzles are equivalent to the hardness of average case sampling problems. In particular, they point out that conjectures regarding the hardness of sampling the output distribution of random IQP circuits should imply quantum bit commitments and one-way state synthesis problems. We believe that because our computational assumption takes a quantum input, and because it implies pseudo-random states instead of a classical communication protocol, our assumption is more plausibly "fully-quantum". Additionally, we are able to build a wide plethora of quantum cryptographic primitives, as opposed to just building commitments and one-way puzzles.

Quantum cryptography from meta-complexity. The recent works of [HM24, CGGH24] also constructs one-way puzzles (and therefore quantum bit commitments) from a different classical complexity theoretic assumption, namely the average case hardness of the gapped Kolmogorov complexity problem. They do not use IQP circuit sampling or random circuit sampling to instantiate their one-way puzzles, and also do not construct primitives like quantum pseudo-random states.

Computational complexity of learning pure states. The recent work of [HH24] constructs EFI pairs and one-way state generators from the existance of learning problems that exhibit average-case hardness. We note that this work does not propose a concrete

family of states for which learning might be hard, and by exploiting the properties of our state family, we are able to create quantum cryptographic primitives beyond one-way state generators and EFI pairs.

Quantum group actions. The recent work of [MX24] propose another way to build quantum cryptography from so called "one-way quantum group actions", which generalize group actions to efficiently describable group representations (and starting states). One candidate quantum group action they propose is a random Z-diagonal circuit, which bares resemblance to the HPS assumption. The authors are also able to construct pseudorandom function-like state generators, but provide no evidence of the one-wayness of these random Z-diagonal circuits.

Acknowledgements

3 Background

3.1 Notation

The notation $x \sim X$ describes that an element x is drawn uniformly at random from the set X. Similarly, if $\mathcal{D}$ is a distribution, we let $x \sim \mathcal{D}$ denote sampling x according to $\mathcal{D}$. We denote the expectation value of a random variable X by $\mathbb{E}[X] = \sum_x x \Pr[X = x]$.

3.2 Quantum Preliminaries

A quantum *register* R corresponds to a finite-dimensional complex Hilbert space. We write $L(\mathsf{R})$ to denote the set of linear transformations on R, $\mathsf{GL}(\mathsf{R})$ to denote the set of invertible linear transformations on R, $\mathsf{S}(\mathsf{R})$ to denote the Hermitian, positive semi-definite matrices over R, and $U(\mathsf{R})$ to denote the unitary group over R. When the register is clear from

context, we sometimes write $U(d)$, where d is the dimension of R. A linear transformation is called a quantum channel if it is completely positive and trace-preserving (i.e. for every positive input with trace t, the output of the map is a positive matrix with trace t).

For a vector $|\psi\rangle \in \mathsf{R}$, we write ψ to denote the density matrix $|\psi\rangle\langle\psi|$, and for vectors $|\psi\rangle, |\phi\rangle \in \mathsf{R}$, we write $\langle\psi|\phi\rangle$ to denote the inner product. For a vector $|\psi\rangle \in \mathsf{R}$, we write $\||\psi\rangle\|$ to denote the standard norm over R, i.e. $\langle\psi|\psi\rangle$. We write $\mathrm{Tr}(\cdot)$ to denote the trace and Tr_R to denote the partial trace over a register R.

For a linear operator $X \in L(\mathsf{R})$, let $\|X\|_\infty$ be its operator norm and $\|X\|_1 = \mathrm{Tr}(\sqrt{X^\dagger X})$ be its trace norm. For two density matrices $\rho, \sigma \in \mathsf{S}(\mathsf{R})$, let $\mathrm{td}(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ be the trace distance between the two. We sometimes write X_R to indicate that X acts on R. All un-labeled operators act on all registers that do not have an operator acting on them, and if an operator is associated with specific registers, we drop the register subscripts for brevity.

For two quantum channels $\mathcal{M}$ and $\mathcal{N}$, the diamond distance between them is

$$\|\mathcal{M} - \mathcal{N}\|_\diamond = \max_{|\psi\rangle \in \mathsf{R}^{\otimes 2}} \|(\mathcal{M} \otimes \mathbb{I}_\mathsf{R})(|\psi\rangle\langle\psi|) - (\mathcal{N} \otimes \mathbb{I}_\mathsf{R})(|\psi\rangle\langle\psi|)\|_1 \, ,$$

where $\mathbb{I}_\mathsf{R}$ is the identity channel on R.

3.3 Useful Lemmas

Here we write lemmas that will be important in the rest of the paper, but should be familiar to someone well-versed in quantum computation.

Lemma 1 (Gentle measurement lemma [Win99])**.** *Given a pure state ρ and a projector Λ, let*

$$\rho' = \frac{\Lambda\rho\Lambda}{\mathrm{Tr}(\Lambda\rho)}$$

be the post-measurement state. Then $\mathrm{td}(\rho', \rho) \le \sqrt{1 - \mathrm{Tr}(\Lambda\rho)}$.

We will also need a similar lemma, but where the projector acts on a register being traced out, and normalization is not applied. A similar lemma appears in [MH24], we state the proof here for completeness.

Lemma 2. *Let Λ_B be a projector and ρ_AB be a quantum state, then the following holds*

$$\|\mathrm{Tr}_\mathsf{B}(\rho_\mathsf{AB}) - \mathrm{Tr}_\mathsf{B}(\Lambda_\mathsf{B}\rho_\mathsf{AB}\Lambda_\mathsf{B})\|_1 = 1 - \mathrm{Tr}(\Lambda_\mathsf{B}\rho_\mathsf{AB})\,.$$

Proof. As B is being traced out, we can write the following.

$$\mathrm{Tr}_\mathsf{B}(\rho_\mathsf{AB}) = \mathrm{Tr}_\mathsf{B}(\Lambda_\mathsf{B}\rho_\mathsf{AB}\Lambda_\mathsf{B}) + \mathrm{Tr}_\mathsf{B}((\mathbb{I} - \Lambda_\mathsf{B})\rho_\mathsf{AB}(\mathbb{I} - \Lambda_\mathsf{B}))\,.$$

Then we can write the left hand side of the theorem statement as follows.

$$\begin{aligned}
\|\mathrm{Tr}_B(\rho_{AB}) - \mathrm{Tr}_B(\Lambda_B \rho_{AB} \Lambda_B)\|_1 &= \|\mathrm{Tr}_B((\mathbb{I} - \Lambda_B)\rho_{AB}(\mathbb{I} - \Lambda_B))\|_1 \\
&= \mathrm{Tr}((\mathbb{I} - \Lambda_B)\rho_{AB}(\mathbb{I} - \Lambda_B)) \\
&= \mathrm{Tr}((\mathbb{I} - \Lambda_B)\rho_{AB}) \\
&= 1 - \mathrm{Tr}(\Lambda_B \rho_{AB}).
\end{aligned}$$

Here the first line uses the previous observation. Then we use the fact that $\|X\|_1 = \mathrm{Tr}(X)$ whenever X is positive-semidefinite. Then we use the cyclic property of the trace and the fact that $\mathbb{I} - \Lambda$ is a projector (and thus squares to itself). Finally, we use the fact that ρ is trace 1. $\qquad\square$

3.4 The Haar Measure and t-designs

Here we define the Haar measure and quantum t-designs.

Definition 3 (Haar measure and Haar random states). *The Haar measure is the unique left- and right- invariant probability measure on the unitary group $U(d)$. A Haar random state is a state sampled by first sampling a unitary from the Haar measure and then applying it to $|0\rangle$ (or any fixed state). We use the notation $\mu(d)$ to denote the Haar measure on $d \times d$ unitary matrices, and $\mathrm{Haar}(d)$ to denote the distribution over Haar random states. Similarly, we denote by $\mu_D(d)$ the unique left- and right- invariant probability measure on the diagonal subgroup $D(d) \subset U(d)$. Moreover, the Haar measure induces a unique unitarily invariant probability measure on states in $\mathbb{C}^d$, which we will also denote by $\mu(d)$.*

An approximate t-design is a discrete distribution over unitaries that looks identical to t-calls to a unitary sampled from the Haar measure.

Definition 4 (t'th moment operator). *Let ν be a probability distribution over unitaries, then the t'th moment operator is defined as*

$$\mathcal{M}_\nu^{(t)}(\cdot) = \mathbb{E}_{U \sim \nu}(\cdot)\left[U^{\otimes t}(\cdot)U^{\dagger, \otimes t}\right].$$

Definition 5 (Approximate unitary and t-design). *An ensemble of $d \times d$ unitaries ν in is called an ϵ-approximate t-design if*

$$\left\|\mathcal{M}_\nu^{(t)} - \mathcal{M}_{\mu(d)}^{(t)}\right\|_\diamond \leq \epsilon.$$

Similarly, an ensemble ν of states in $\mathbb{C}^d$ is a ε-approximate state t-design if

$$\left\|\mathbb{E}_{\psi \sim \nu}(|\psi\rangle\langle\psi|)^{\otimes t} - \mathbb{E}_{|\psi\rangle \sim \mu(d)}(|\psi\rangle\langle\psi|)^{\otimes t}\right\|_1 \leq \varepsilon. \tag{3.6}$$

We will also need a notion of approximate designs on the diagonal subgroup.

Definition 7 (Approximate diagonal t-design)**.** *An ensemble of $d \times d$ diagonal unitaries ν in is called an ϵ-approximate diagonal t-design if*

$$\left\| \mathcal{M}_\nu^{(t)} - \mathcal{M}_{\mu_D(d)}^{(t)} \right\|_\diamond \leq \epsilon \,.$$

Here μ_D is the Haar measure (unique left- and right- invariant measure) over diagonal unitary matrices.

3.5   Cryptographic Primitives

Here we define a number of cryptographic primitives.

One-Way State Generator.   First, we define a one-way state generator [MY22]. These are quantum equivalents of one-way functions, with classical input and quantum output.

Definition 8 (One-Way State Generator)**.** *Let $n \in \mathbb{N}$ denote the security parameter. A one-way state generatator is a tuple $(\mathsf{KeyGen}, \mathsf{StateGen}, \mathsf{Ver})$ consisting of the following alorithms:*

- $\mathsf{KeyGen}(1^n) \to \mathsf{k}$*: on input 1^n, it outputs a classical key k.*

- $\mathsf{StateGen}(k) \to |\phi_k\rangle$*: on input k, it outputs a quantum state $|\phi_k\rangle$.*

- $\mathsf{Ver}(k, \sigma) \to \top/\bot$*: on input a key k and a state σ, it outputs $\top$ or $\bot$.*

We require that the algorithms satisfy the following properties:

(a) ***Correctness:*** *For any $n \in \mathbb{N}$, it holds that*

$$\Pr\left[\top \leftarrow \mathsf{Ver}(k, |\phi_k\rangle) \;:\; \begin{smallmatrix} k \leftarrow \mathsf{KeyGen}(1^n) \\ |\phi_k\rangle \leftarrow \mathsf{StateGen}(k) \end{smallmatrix} \right] \geq 1 - \mathsf{negl}(n).$$

(b) ***Security:*** *For any $n \in \mathbb{N}$, any $t = \mathsf{poly}(n)$, and for any QPT adversary $\mathcal{A}$,*

$$\Pr\left[\top \leftarrow \mathsf{Ver}(k', |\phi_k\rangle) \;:\; \begin{smallmatrix} k \leftarrow \mathsf{KeyGen}(1^n) \\ |\phi_k\rangle \leftarrow \mathsf{StateGen}(k) \\ k' \leftarrow \mathcal{A}(|\phi_k\rangle^{\otimes t}) \end{smallmatrix} \right] \leq \mathsf{negl}(n) \,,$$

where $\mathsf{Ver}(k', |\phi_k\rangle)$ denotes the algorithm which applies the projective measurement

$$\{ |\phi_{k'}\rangle\langle\phi_{k'}|, I - |\phi_{k'}\rangle\langle\phi_{k'}| \}$$

onto $|\phi_k\rangle$ and outputs $\top$, if the measurement succeeds, and outputs $\bot$ otherwise.

Pseudorandom State Generator. Pseudorandom states are collections of states that are indistinguishable from a Haar random state to bounded adversaries. These were first defined in [JLS18].

Definition 9 (Pseudorandom State Generator). *Let $n \in \mathbb{N}$ be the security parameter. A pseudorandom state generator (PRSG) is a tuple* $(\mathsf{KeyGen}, \mathsf{StateGen})$ *of QPT algorithms:*

- $\mathsf{KeyGen}(1^n) \to k$: *on input* 1^n, *it outputs a classical key k.*

- $\mathsf{StateGen}(k) \to |\phi_k\rangle$: *on input k, it outputs a quantum state* $|\phi_k\rangle$.

We require that the algorithms satisfy the following security property: For all $n \in \mathbb{N}$ and every polynomial function $t = \mathsf{poly}(n)$, and every QPT adversary $\mathcal{A}$,

$$\left| \Pr_{k \leftarrow \mathsf{KeyGen}(1^n)}[\mathcal{A}(|\phi_k\rangle^{\otimes t}) = 1] - \Pr_{|\phi\rangle \sim \mathrm{Haar}(2^n)}[\mathcal{A}(|\phi\rangle^{\otimes t}) = 1] \right| \leq \mathsf{negl}(n)$$

4 Hamiltonian Phase State Assumption

In this section, we give a formal definition of our hardness assumption. Recall that an n-qubit Hamiltonian Phase State is of the form

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i} \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}} \right) H^{\otimes n} |0^n\rangle$$

where $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ is a binary matrix and $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$ is a set of angles in the interval $[0, 2\pi)$. To avoid matters of precision, we introduce a discretization parameter $q \in \mathbb{N}$ with $q = \mathsf{poly}(n)$ and partition the interval $[0, 2\pi)$ into q parts via the set

$$\Theta_q := \left\{ \frac{2\pi k}{q} : k \in \{0, 1, \ldots, q-1\} \right\}.$$

We now introduce two variants of our hardness assumption.

4.1 Search Variant

Our first variant considers a search problem. Roughly speaking, it says that given many copies of a random Hamiltonian phase state, it is computationally difficult to reverse-engineer its architecture and its angles. Therefore, our assumption says that an ensemble of Hamiltonian Phase states forms a one-way state generator [MY22]. We provide a formal definition of a one-way state generator in Definition 8.

We now give a formal definition.

Definition 1 (Search HPS). *Let $n \in \mathbb{N}$ denote the security parameter, and let m and q be integers (possibly depending on n). Let χ be a distribution with support over matrices in $\mathbb{Z}_2^{m \times n}$. Then, the (search) Hamiltonian Phase State assumption ($\mathsf{HPS}_{n,m,q,\chi}$) states that, for any number of copies $t = \mathsf{poly}(n)$ and for any efficient quantum algorithm $\mathcal{A}$,*

$$\Pr\left[1 \leftarrow \mathsf{Ver}(\mathbf{A}', \boldsymbol{\theta}', |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle) : \begin{matrix} \mathbf{A} \sim \chi, \boldsymbol{\theta} \sim \Theta_q^m \\ (\mathbf{A}', \boldsymbol{\theta}') \leftarrow \mathcal{A}(|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}) \end{matrix}\right] \leq \mathsf{negl}(n),$$

where $\mathsf{Ver}(\mathbf{A}', \boldsymbol{\theta}', |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle)$ denotes the algorithm which applies the projective measurement

$$\{|\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle\langle\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}|, I - |\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle\langle\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}|\}$$

onto $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ and outputs 1, if the measurement succeeds, and outputs 0 otherwise. We say that a quantum algorithm solves the (search) $\mathsf{HPS}_{n,m,q,\chi}$ problem if it runs in time $\mathsf{poly}(n, m, \log q)$ and succeeds with probability at least $1/\mathsf{poly}(n, m, \log q)$.

An alternative but equivalent formulation of the security property is to say that it is computationally difficult to find a state $|\Phi_{\boldsymbol{\theta}'}^{\mathbf{A}'}\rangle$ which has non-vanishing fidelity with the input state, on average over the choice of architecture and set of angles.

4.2 Decision Variant

Our second variant considers a decision problem. Roughly speaking, it says that given many copies of a random Hamiltonian phase state, it is computationally difficult to distinguish it from many copies of a Haar random state. Therefore, our (decision) assumption says that an ensemble of Hamiltonian Phase states forms a so-called pseudorandom state generator [JLS18, MY22]. We provide a formal definition of a pseudorandom state generator in Definition 9.

Definition 2 (Decision HPS). *Let $n \in \mathbb{N}$ denote the security parameter, and let m and q be integers (possibly depending on n). Let χ be a distribution with support over matrices in $\mathbb{Z}_2^{m \times n}$. Then, the (decision) Hamiltonian Phase State assumption ($\mathsf{HPS}_{n,m,q,\chi}$) states that, for any number of copies $t = \mathsf{poly}(n)$ and for any efficient quantum distinguisher $\mathcal{D}$,*

$$\left|\Pr\left[1 \leftarrow \mathcal{D}(|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle^{\otimes t}) : \begin{matrix} \mathbf{A} \sim \chi \\ \boldsymbol{\theta} \sim \Theta_q^m \end{matrix}\right] - \Pr\left[1 \leftarrow \mathcal{D}(|\Psi\rangle^{\otimes t}) : |\Psi\rangle \sim \mathrm{Haar}(2^n)\right]\right| \leq \mathsf{negl}(n),$$

We say that a quantum algorithm solves the (decision) $\mathsf{HPS}_{n,m,q,\chi}$ problem if it runs in time $\mathsf{poly}(n, m, \log q)$ and succeeds with probability at least $1/\mathsf{poly}(n, m, \log q)$.

5 Evidence for Average-Case Hardness and Full Quantumness

In this section, we give several pieces of evidence for the security of the $\mathsf{HPS}_{n,m,q,\chi}$ assumption, as well as evidence that it is a fully quantum assumption.

First, in Section 5.1, we show two worst-to-average-case reductions for the $\text{HPS}_{n,m,q,\chi}$ problem, and also discuss the limitations of those reductions. To this end, we first show that if the Hamiltonian architecture matrix $\mathbf{A}$ is publicly known, then there is a worst-to-average-case reduction for the angles $\theta \in \Theta_q$. Second, we show that for $m = n$, and any fixed set of angles, there is a worst-to-average-case reduction over the architecture matrices $\mathbf{A}$, and discuss limitations of the reduction. We argue that the HPS assumption is related to the security of the McEliece cryptographic system.

Then, we show that if χ is the uniform distribution of $m \times n$ binary matrices and $q > 2t$, Hamiltonian Phase States with $m \gtrsim nt^2$ random terms form approximate state t-designs in Section 5.2. This shows that given less than $\Omega(\sqrt{m/n})$ many copies, HPS are information-theoretically indistinguishable from Haar-random states. It also implies that the Hamiltonian Phase States contain an exponentially large set of almost orthogonal states. This implies that Hamiltonian phase states are fast mixing, giving additional evidence that the learning problem is computationally hard.

Then, in Section 5.3, we discuss algorithms for learning phase states with public and secret architecture matrices. In particular, we give a sample-optimal (but exponential-time) algorithm for solving the HPS problem using pretty good measurements [BK00, Mon19] and a simple algorithm that uses classical shadows [Aar18, HKP20].

Finally, we discuss why the HPS assumption is fully quantum. To this end, we give evidence against the possibility of building one-way functions from HPS.

5.1 Worst-Case to Average-Case Reduction

We begin providing evidence for the security of our assumption by showing how in different regimes learning the parameters of the HPS problem of a fixed (worst-case) instance can be reduced to learning a random instance. Our evidence will treat the angles θ and the Hamiltonian architecture matrix $\mathbf{A}$ separately. Specifically, we will show two types of worst-to-average-case reductions. First, we will show that in a certain regime of m, n, given a copy of a HPS instance, a quantum algorithm can efficiently generate a random HPS with the same angles and architecture dimensions. Second, we will fix the Hamiltonian architecture $\mathbf{A}$ and show that given a copy of a HPS instance and its architecture $\mathbf{A}$, a quantum algorithm can generate a random HPS with the same architecture but uniformly random angles. Our worst-to-average-case reductions are therefore similar to those for, say, the Learning with Errors (LWE) problem [Reg09], with different levels of public knowledge.

Reduction for the architecture for $m \leq n$ First, we observe that for any fixed choice of angles θ, the Hamiltonian architecture can be re-randomized if $m \leq n$ and χ is the uniform distribution over full-rank matrices $\mathcal{R}(m, n) := \{\mathbf{A} \in \mathbb{Z}_2^{m \times n} \mid \text{rank}(\mathbf{A}) = \min(m, n)\}$. Notice that the restriction to Hamiltonian architectures with full rank is not too significant, since the probability that a uniformly random $\mathbb{Z}_2^{m \times n}$ matrix has full rank with probabil-

ity[2] $\prod_{k=1}^{\min(m,n)}(1 - 2^{-k}) \geq 0.288$ [SF]. The basic idea of the reduction is to apply a circuit composed of uniformly random CNOT gates to the given HPS instance. In the parameter regime we consider, this will have the effect of completely scrambling the Hamiltonian architecture to a uniformly random one with the same choices of m, n and subject to the full-rank constraint.

Lemma 1 (Worst-to-average-case reduction for the architecture). *Suppose there exists an algorithm $\mathcal{A}$ that runs in time T and solves the (search) $\mathsf{HPS}_{n,m,q,\chi}$ problem with probability ϵ in the average case, where χ is the uniform distribution over $\mathcal{R}(m, n)$ and $m \leq n$. Then, there exists an algorithm which runs in time $T \cdot \mathrm{poly}(n)$ and inverts Hamiltonian phase states $|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle^{\otimes t}$ with probability ϵ for a worst-case choice of architecture $\mathbf{C} \in \mathcal{R}(m, n)$, uniformly random angles $\boldsymbol{\theta}$, and for any number of copies $t = \mathrm{poly}(n)$. Here, $\mathcal{R}(m, n) = \{\mathbf{A} \in \mathbb{Z}_2^{m \times n} | \mathrm{rank}(\mathbf{A}) = \min(m, n)\}$ is the set of full-rank binary $m \times n$ matrices.*

Proof. Consider the reduction $\mathcal{B}$ which, on input $|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle^{\otimes t}$, does the following:

1. $\mathcal{B}$ samples a uniformly random invertible matrix $\mathbf{R} \sim \mathrm{GL}(n, \mathbb{Z}_2)$.

2. $\mathcal{B}$ runs the average-case solver $\mathcal{A}$ on the input

$$(U_{\mathbf{R}} |\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle)^{\otimes t}.$$

where $U_{\mathbf{R}}$ is the n-qubit unitary transformation given by $U_{\mathbf{R}} : |x\rangle \mapsto |\mathbf{R}^{-1} \cdot x\rangle$, for $x \in \{0, 1\}^n$. Finally, $\mathcal{B}$ outputs whatever $\mathcal{A}$ outputs.

Note that $U_{\mathbf{R}}$ is a quantum circuit composed just of CNOT gates and therefore efficiently implementable. Because the average-case solver $\mathcal{A}$ runs in time T, it follows that the reduction $\mathcal{B}$ runs in time $T \cdot \mathrm{poly}(n)$.

Next, we show that $\mathcal{B}$ also succeeds with probability ϵ. By assumption, the worst-case instance $|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle^{\otimes t}$ consists of structured phase states states

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) H^{\otimes n} |0^n\rangle,$$

where $\mathbf{C} \in \mathcal{R}(m, n)$ and $\boldsymbol{\theta}$ is a tuple of random angles $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m) \in \Theta_q^m$. To complete the proof, it suffices to show that $U_{\mathbf{R}} |\psi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle)^{\otimes t}$ is distributed exactly as in the $\mathsf{HPS}_{n,n,q,\chi}$ problem, where χ is the uniform distribution over $\mathcal{R}(m, n)$. First, we make the following key observation: it follows from unitarity of $U_{\mathbf{R}}$ that

$$U_{\mathbf{R}} |\Phi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle = \left(U_{\mathbf{R}} \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) U_{\mathbf{R}}^{\dagger} \right) U_{\mathbf{R}} H^{\otimes n} |0^n\rangle.$$

---

[2]See https://math.mit.edu/~dav/genlin.pdf, and this Stackexchange post for a proof.

Because $U_{\mathbf{R}}$ is an invertible matrix, it leaves the state $H^{\otimes n} \ket{0^n}$ invariant, and thus we have $U_{\mathbf{R}} H^{\otimes n} \ket{0^n} = H^{\otimes n} \ket{0^n}$. Next, we study the action of $U_{\mathbf{R}}$ onto tensor products of Pauli operators. We find that for any index $i \in [n]$:

$$
U_{\mathbf{R}} \left(\bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) U_{\mathbf{R}}^{\dagger} = \sum_{x \in \{0,1\}^n} \bra{x} U_{\mathbf{R}} \left(\bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) U_{\mathbf{R}}^{\dagger} \ket{x} \cdot \ket{x}\bra{x}
$$

$$
= \sum_{x \in \{0,1\}^n} \bra{\mathbf{R}x} \left(\bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) \ket{\mathbf{R}x} \cdot \ket{x}\bra{x}
$$

$$
= \sum_{x \in \{0,1\}^n} (-1)^{\sum_{j=1}^{n} \mathbf{C}_{ij}(\mathbf{R}x)_j} \ket{x}\bra{x}
$$

$$
= \sum_{x \in \{0,1\}^n} (-1)^{\sum_{j=1}^{n} (\mathbf{C} \cdot \mathbf{R})_{ij} x_j} \ket{x}\bra{x}
$$

$$
= \sum_{x \in \{0,1\}^n} \bra{x} \left(\bigotimes_{j=1}^{n} Z^{(\mathbf{C} \cdot \mathbf{R})_{ij}} \right) \ket{x} \cdot \ket{x}\bra{x}
$$

$$
= \bigotimes_{j=1}^{n} Z^{(\mathbf{C} \cdot \mathbf{R})_{ij}}.
$$

Because $U_{\mathbf{R}}$ is acting on a matrix exponential of a diagonal matrix, it follows that

$$
U_{\mathbf{R}} \exp \left(\mathrm{i} \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) U_{\mathbf{R}}^{\dagger} = \exp \left(\mathrm{i} \sum_{i=1}^{m} \theta_i U_{\mathbf{R}} \left(\bigotimes_{j=1}^{n} Z^{\mathbf{C}_{ij}} \right) U_{\mathbf{R}}^{\dagger} \right)
$$

$$
= \exp \left(\mathrm{i} \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{(\mathbf{C} \cdot \mathbf{R})_{ij}} \right).
$$

Finally, we observe that for $m = n$, $\mathcal{R}(m, n) = \mathrm{GL}(n, \mathbb{Z}_2)$, which is a group. Because $\mathbf{C} \in \mathrm{GL}(n, \mathbb{Z}_2)$ it follows that $\mathbf{C} \cdot \mathbf{R}$ is uniformly distributed whenever $\mathbf{R} \sim \mathrm{GL}(n, \mathbb{Z}_2)$. Putting everything together, it follows that $U_{\mathbf{R}} \ket{\psi_{\boldsymbol{\theta}}^{\mathbf{C}}})^{\otimes t}$ is distributed precisely as in the $\mathrm{HPS}_{n,n,q,\chi}$ problem, and thus $\mathcal{B}$ succeeds with probability ϵ. The claim for $m \leq n$ follows from the fact that in that case $\mathbf{C}$ is a submatrix of a $\mathrm{GL}(n, \mathbb{Z}_2)$ matrix. $\qquad \square$

Lemma 1 shows a worst-to-average-case reduction in the regime $m \leq n$ using only a CNOT circuit for randomization. We can also obtain a meaningful worst-case to average-case reduction when $m > n$ using additional Hamiltonian terms on a larger system, and a CNOT circuit on the larger system. However, we currently do not know how to analyze the resulting distribution and therefore state it as a conjecture:

Conjecture 2. *There is a regime of $m > n$ and choices of $k, l = \mathsf{poly}(n)$ such that the following is true. The uniform distribution with $C \sim \mathbb{Z}_2^{(m+k) \times (n+l)}$ is $1/\mathsf{poly}(n)$ close in total variation distance to the distribution over matrices of the form*

$$\mathbf{M} = \Pi \cdot \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{B} & \mathbf{C} \end{pmatrix} \cdot \mathbf{R}, \tag{5.3}$$

where $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ is a fixed matrix, and all other matrices are uniformly random, i.e., $(\mathbf{B}|\mathbf{C}) \sim \mathbb{Z}_2^{m \times (n+l)}$, $\mathbf{R} \sim \mathrm{GL}(n+l, \mathbb{Z}_2)$, Π is a random row-permutation operator which we identify with an element in S_{m+k}, where S_N denotes the symmetric group over N elements.

The Hamiltonian $\mathbf{M}$ characterizes a HPS which is obtained from a phase state $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ via the following re-randomization:

$$U_{\mathbf{R}} \exp\left(i \sum_{i=1}^{k} \varphi_i \sum_{j=1}^{n+l} Z_j^{(\mathbf{B}|\mathbf{C})_{ij}} \right) |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle \otimes |+\rangle^l = |\Phi_{(\boldsymbol{\theta}, \boldsymbol{\varphi})}^{\mathbf{M}}\rangle, \tag{5.4}$$

using additional angles $\boldsymbol{\varphi} \in \Theta_q^k$. Note that the permutation Π of the rows of $\mathbf{M}$ does not affect the final state, but is a freedom we have in its description since all Hamiltonian terms commute. Interestingly, these are precisely the freedoms exploited in the proposal for verified quantum advantage by Shepherd and Bremner [SB09], and studied in more detail in Refs. [Kah23, BCJ23, GH23].

Worst-case hardness and relation with the McEliece cryptosystem   Conjecture 2 makes the connection to a particular classical cryptosystem, the McEliece cryptosystem, explicit. Using the relation to this cryptosystem, we argue for the worst-case hardness of the HPS problem.

The McEliece cryptosystem is a code-based candidates for post-quantum cryptography [McE78]. The scheme is based on Goppa codes:

Definition 5 (Goppa code). *Let $g(z) = \sum_i g_i z^i \in \mathbb{F}_{q^m}[z]$ be a polynomial and $L = \{\alpha_1, \dots, \alpha_n\} \subseteq \mathbb{F}_{q^m}$ be values that do not evaluate to 0, i.e. $g(\alpha_i) \neq 0$. Then the Goppa code with parameters $g(z)$ and L, $\Gamma(L, g(z))$ has the codewords*

$$\left\{ c = (c_1, \dots, c_n) \in \mathbb{F}_{q^m}^n : \sum_{i=1}^{n} \frac{c_i}{z - \alpha_i} \equiv 0 \bmod g(z) \right\}.$$

Note that the sum in the condition describes a function of z, not an element of $\mathbb{F}_{q^m}$.

The definition of the Goppa code will not be important, but it is important to note that instantiations of the McEliece cryptosystem with some other kinds of codes have been shown to be insecure, while the instantiation with Goppa codes is still believed to be secure against quantum adversaries. The cryptosystem is usually defined in terms of the following public-key encryption scheme.

1. Alice samples a $[n, k, 2t+1]$-Goppa code, specified by its $k \times n$ sized generator matrix G, a random $k \times k$ non-singular matrix R, and a permutation $\Pi \in S_n$.

2. Alice releases $\hat{G} = R \cdot G \cdot \Pi$ and the rate t as the public key, and keeps the tuple (R, G, Π) as her private key.

3. To encrypt a message m of k bits, Bob computes a n-bit error vector e with Hamming weight t and computes $c = m \cdot \hat{G} + e$.

A related problem to the security of the primitive is the so called "Goppa code distinguishing" problem, which asks if the public key $\hat{G} = R \cdot G \cdot \Pi$ can be distinguished from a uniformly random matrix. In certain regimes, specifically the high rate regime, it is known that there is an efficient distinguisher [FGO+13], however for the parameter range used in practice the problem is believed to be hard.

Consider the following variant of the HPS problem

Definition 6 (Goppa code phase learning). *Let G be a $[n, k, t]$-Goppa code, and consider the state $|\Phi_\theta^{G^\mathsf{T}}\rangle$. The search variant of the Goppa code phase learning problem is the problem of learning the phases θ given a desciption of G and samples of $|\Phi_\theta^{G^\mathsf{T}}\rangle$.*

The decision variant is to distinguish the states $|\Phi_\theta^{G^\mathsf{T}}\rangle$ from a Haar random state, given samples of the state and the architecture matrix.

We can relate this problem to the search and decision variants of HPS using the security of the McEleice cryptosystem.

Lemma 7. *If the Goppa code distinguishing problem is hard for quantum adversaries, then the Goppa code phase learning problem is as hard as the* HPS *problem.*

Proof. Given an adversary for the HPS problem, which works for a uniformly random architecture $\mathbf{A}$, we can sample a random permutation and non-singular matrix, and apply the architecture re-randomization from the previous section to get $\hat{G}^\mathsf{T} = \Pi G^\mathsf{T} R$. We can further re-randomize the phase by applying another IQP circuit with the same architecture. By the hardness of the Goppa code distinguishing problem, this architecture is indistinguishable from a uniformly random one for the HPS solver, so the solver will output the correct phases with high probability. $\qquad\square$

Since Lemma 7 holds for all Goppa codes, the Goppa code phase learning problem (Definition 6) is a worst-case problem over a restricted family of architectures that come from Goppa codes. It shows that the $\text{HPS}_{m,n,q,\chi}$ problem is at least as hard as the Goppa code distinguishing problem, thereby showing a worst-case reduction for it. Note however, that the parameter regime of the worst-case instances may not be the same regime in which we can show a worst-to-average-case reduction.

Limitations of the architecture reduction. Let us now discuss some limitations of the worst-to-average-case reductions we have given for the Hamiltonian architectures with fixed angles. Specifically, we provide evidence that randomizing with CNOT circuits

and additional Hamiltonian terms is not sufficient for a worst-to-average-case reduction whenever $m \gg n$, but that this might still be possible in the regime in which $m = n + O(\log n)$. We note however that while the *information-theoretic* worst-to-average-case reduction might not work in this regime, we still believe that the *computational* hardness persists, and therefore Conjecture 2 to be true.

To give a lower bound on the total-variation distance between the uniform distribution and the distribution with a fixed Hamiltonian architecture $\mathbf{A}$, it is sufficient to describe a distinguishing test and give a lower bound on its success probability. Our test will distinguish two matrices M_1, M_2 drawn from the hidden ensemble μ with architecture matrices chosen as in Eq. (5.3) of Conjecture 2 (with $l = 0$) versus the uniform ensemble $\mathcal{U}$ over $(m + k) \times n$ binary matrices, respectively. To this end, let us write $\mathbf{A} = \mathbf{ZE}$ in row echelon form $\mathbf{E}$ via an invertible matrix $\mathbf{Z} \in \mathrm{GL}(m, \mathbb{Z}_2)$. We ask the question: what is the probability that there exist an invertible matrix $\mathbf{Z} \in \mathrm{GL}(m, \mathbb{Z}_2)$ and a matrix $\mathbf{D} \in \mathbb{Z}_2^{m \times n}$ in reduced row-echelon form with the property that

$$\exists \mathbf{P}_1, \mathbf{P}_2 \in S(m + k) : (\mathbf{Z} \oplus \mathbb{I}_k)\mathbf{P}_1\mathbf{M}_1 = (\mathbf{Z} \oplus \mathbb{I}_k)\mathbf{P}_2\mathbf{M}_2 = \begin{pmatrix} \mathbf{E} \\ \vdots \end{pmatrix}? \qquad (5.8)$$

In particular, $\mathbf{E}$ has $m - n$ all-0 rows which distinguish it from a uniformly random matrix even up to row permutations, and multiplication by a $\mathrm{GL}(n, \mathbb{Z}_2)$ matrix from the right. For the μ ensemble, we know that this probability is equal to 1 by construction. For the uniform ensemble, we can get a rough estimate on it as follows

- There are $\binom{k+m}{m}$ subsets of m rows of $\mathbf{M}_1$, and $m!$ permutations of those rows. Hence, there are this many matrices $\mathbf{Z}_i$ which bring the permutations of those rows into the form $\mathbf{E}$. Let us think of all of the $\mathbf{Z}_i$ as being different from each other.

- For each $\mathbf{Z}_i$, we can apply it to an arbitrary permutation of m rows of $\mathbf{M}_2$ and check the probability of getting $\mathbf{E}$ (of which there are $m!\binom{k+m}{m}$). This probability is given roughly by $2^{-m \cdot n + \Theta(n^2)}$ if we model each choice of m rows as iid. uniformly random, and the factor of $2^{\Theta(n^2)}$ comes from the fact that we only care about equality up to a $\mathrm{GL}(n, \mathbb{Z}_2)$ transformation on the right (of which there are roughly $2^{\Theta(n^2)}$). Modelling the submatrices as uniform iid. matrices is of course strictly speaking false. However, we expect it to be approximately true. To see why observe that if we choose a fixed subset of m rows of $\mathbf{M}$, denoted as $\mathbf{A} = \mathbf{ZE}$ and exchange only one of the rows of $\mathbf{A}$ for another row of the big matrix $\mathbf{M}$, yielding a matrix $\mathbf{A}^*$, then $\mathbf{Z}^{-1}\mathbf{A}^*$ will be far from $\mathbf{E}$ and look essentially random.

In total we roughly get a 'collision probability' of

$$\frac{(k+m)!^2}{k!^2} \cdot 2^{-m \cdot n + \Theta(n^2)} \sim 2^{2m \log k - n(m - \alpha n)} = 2^{2dcn \log n - (d - \alpha)n^2} \in \exp(-O(n^2)) \qquad (5.9)$$

where we chose $m = dn$ and $k = n^c$, and we fix $\Theta(n^2) = \alpha n^2$. This suggests that we can make the collision probability on the order of 1 (i.e., a constant TVD away from the hidden ensemble μ) if we choose $m - n = e \log n$ for some constant e, yielding

$$2^{m \log k - n(m-n)} = 2^{(2c-e)n \log n + 2ec \log n}. \tag{5.10}$$

Hence, we can choose c in $k = n^c$ large enough that $2c - e \geq 0$. This suggests that the worst-to-average-case reduction using random additional Hamiltonian terms and randomization by a CNOT circuit fails for $k \in \omega(\log n)$, but might still work for $k \in O(\log n)$.

The worst-to-average-case reduction over the architectures A does not require any information about the angles. In Section 5.3 we will show that HPS can be broken if the architecture A is known precisely in the regime $m \leq n$ in which the above worst-to-average reduction works. It can, however, be desirable to make part of the key publicly available. An example where making part of the key public is helpful for designing cryptographic primitives is the Learning with Errors (LWE) problem [Reg09]. The LWE problem is to recover an unknown vector $\mathbf{s} \in \mathbb{Z}_q^n$ from "noisy" scalar products with public but random vectors $\mathbf{a} \in \mathbb{Z}_q^n$. Clearly, the LWE problem could be made harder by keeping the random vectors $\mathbf{a}$ secret, but keeping them public is crucial for many applications such as public-key encryption. In the following, we study a variant of the HPS problem, where the architectures are randomly drawn but publicly available. We discover that even with public architecture, we can reduce the problem of learning the angles of a worst case to learning those of a typical instance.

Reduction for the angles. To show a worst-to-average case reduction for task of learning the angles with a public architecture, we show that given an (worst-case) $\mathrm{HPS}_{n,n,q,\chi}$ instance $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ with $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ in the clear, the angles can be completely randomized by a simple quantum algorithm. The key idea is just to one-time-pad the angles by applying uniformly random rotations around the public Hamiltonian terms to the given HPS $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$.

Lemma 11 (Worst-to-average-case reduction for the angles). *Suppose there is a quantum algorithm $\mathcal{A}$ which for every choice of $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$, given a HPS $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ with $\boldsymbol{\theta}$ chosen uniformly at random from Θ_q^m, succeeds in returning $(\boldsymbol{\theta}, \mathbf{A})$ with probability at least $1/\mathrm{poly}(m, n, \log q)$ over the choice of $\boldsymbol{\theta}$. Then there is a quantum algorithm $\mathcal{A}'$ which succeeds at this task for* every *choice of $\boldsymbol{\theta}$ with probability $1/\mathrm{poly}(m, n, \log q)$.*

Proof. To show the lemma, we observe that given any fixed HPS $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$, we can randomize it by choosing $\boldsymbol{\varphi} \sim \Theta_q^m$ uniformly at random and apply $\mathcal{A}$ to the re-randomized state

$$\exp\left(i \sum_{j=1}^{m} \varphi_j \prod_{k=1}^{m} Z_k^{\mathbf{A}_{jk}} \right) |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle \equiv |\Phi_{\boldsymbol{\theta}+\boldsymbol{\varphi}}^{\mathbf{A}}\rangle. \tag{5.12}$$

Repeating this procedure at most $\mathrm{poly}(n, m, \log q)$ times will recover $\boldsymbol{\theta} + \boldsymbol{\varphi}$ and $\mathbf{A}$, and thus $(\boldsymbol{\theta}, \mathbf{A})$. This implies that an average-case algorithm can solve worst-case instances and thus proves the reduction. $\qquad \square$

Learning with public architecture The worst-to-average-case reduction works for any parameters of m, n, but of course, for the hardness of HPS it is crucial that the learning problem is also worst-case hard. In particular, one might wonder whether the HPS assumption can be relaxed to a publicly known architecture, and unknown angles. We delineate the parameter regime in which we can expect this in the following. To this end, we show that for $m \leq n$ the angles can in fact be efficiently learned given a public, full-rank Hamiltonian architecture. Nonetheless, we believe such a stronger assumption to still be true whenever $m \gg n$.

Lemma 13 (Learning with public Hamiltonian architecture). *Let $n \in \mathbb{N}$ be the security parameter, and $m = \mathrm{poly}(n)$ and $q = \mathrm{poly}(n)$. Consider the family of $\mathsf{HPS}_{n,m,q,\chi}$ instances $\mathcal{E}_{\mathbf{A}} = \{|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle\}_{\boldsymbol{\theta} \in \Theta_q^m}$ with publicly known $\mathbf{A} \in \mathcal{R}(m, n)$. Then, there exists an efficient quantum algorithm which inverts any state in the ensemble $\mathcal{E}_{\mathbf{A}}$ whenever $m \leq n$.*

Proof. Let $\mathbf{R} \in \mathrm{GL}(n, \mathbb{Z}_2)$ be an invertible matrix such that if we write $\mathbf{A} = \sum_{i=1}^{m} |e_i\rangle \langle a_i|$, i.e., the rows of $\mathbf{A}$ are denoted by $|a_i\rangle \in \mathbb{Z}_2^n$, and the standard vectors by $|e_i\rangle \in \mathbb{Z}_2^m$, then $\langle a_i| \mathbf{R} = \langle e_i|$. Then $\mathbf{AR} = (\mathbb{I}_m | 0_{(n-m) \times (n-m)})$.

This means that we can apply a CNOT circuit $U_{\mathbf{R}}$ to $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$, yielding

$$U_{\mathbf{R}} |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(i \sum_{j=1}^{m} \theta_j \prod_{k=1}^{n} Z_k^{(\mathbf{AR})_{jk}}\right) |+\rangle^{\otimes n} \tag{5.14}$$

$$= \exp\left(i \sum_{j=1}^{m} \theta_j \prod_{k=1}^{n} Z_j\right) |+\rangle^{\otimes n} = \bigotimes_{j=1}^{m} \left(e^{i\theta_j Z_j} |+\rangle\right) \otimes |+\rangle^{\otimes(n-m)}, \tag{5.15}$$

which is a product state. The angles θ_j can now be learned to precision $1/\epsilon$ using $O(1/\epsilon^2)$ many copies of $U_{\mathbf{R}} |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ via a measurement in the X-basis. $\qquad \square$

The consequence of Lemma 13 is that at least in the parameter regime in which we can show a worst-to-average case reduction over the architecture, it is necessary for the HPS assumption to be true that the Hamiltonian architecture is hidden. We note that—like for the worst-to-average-case reduction over the architectures—we expect that the learning algorithm can be extended to $m = n + O(\log n)$.

5.2 Hamiltonian Phase States Form Approximate State Designs

In this subsection we show that the states in the HPS ensemble form approximate state designs if $m \geq Cn$ for a constant $C > 0$. It will be convenient to view HPS as a random walk of depth m on the diagonal group. We will therefore slightly adjust the notation.

Consider the following probability distribution ν on the diagonal subgroup of $SU(2^n)$: Draw a uniformly random bitstring $A_1 \in \{0,1\}^n$ and a uniformly random angle $\theta \in [0, 2\pi)$ and apply $e^{i\theta \otimes_{j=1}^n Z^{A_{1j}}}$. We can draw m such diagonal unitaries independently and multiply them. The resulting probability measure is denoted by ν^{*m}.

We will first show that $e^{i \sum_{i=1}^m \theta_i \otimes_{j=1}^n Z^{A_{ij}}}$ is an approximate t-design on the diagonal group. More precisely, we prove the following theorem:

Theorem 16. *For $m \geq 2t(2nt + \log(1/\varepsilon))$ the random unitary $e^{i\theta_i \sum_{i=1}^m \otimes_j Z^{A_{ij}}}$ with random A_{ij} and θ_i is a ε-approximate diagonal t-design in the sense of Definition 7. Moreover, the same bound holds if θ_i is drawn uniformly from $\{2\pi k/q\}_{k=1}^q$, where q is an integer satisfying $q > 2t$.*

We provide a proof of Theorem 16 in Appendix A. The proof of Theorem 16 is remarkably simple in comparison to the derivations of similar results for random quantum circuits [CHH+24, BHH16, Haf22]. Additionally, the constants in Theorem 16 are unusually small: In stark contrast the constants in these results are north of 10^{13}. A similar result was obtained in Ref. [HLT24] for the related random Pauli rotations $e^{i\theta P}$ for a random $\theta \in (0, 2\pi]$ and a random Pauli string P.

Theorem 16 almost directly implies the following corollary:

Corollary 17. *For $m \geq 2t(2nt + \log(1/\varepsilon))$ the state ensemble defined by $|\Phi_\theta^A\rangle = U|+^n\rangle$ for U drawn from ν^{*m} (or ν_q^{*m} for $q > 2t$) is a $\varepsilon + O(t^2/2^n)$-approximate state t-design.*

As a consequence no algorithm with access to t copies can distinguish the states $|\phi_\theta^A\rangle$ from Haar random. In particular, this rules out a large class of natural attacks which make use of a small number of samples. Prominent examples in classical cryptanalysis are linear attacks (2-wise independence rules this out), and differential attacks (t-wise independence rules out $\log_2(t)$ differential attacks). Moreover, the fact that HPS with sufficiently many terms can generate arbitrary state t designs makes it seem unlikely even that there is a distinguishing algorithm using just a few more than t samples. This would mean that there is a sharp transition in the complexity of distinguishing HPS states from uniform. Thus, the t-design property gives evidence for the security of the HPS assumption.

Proof of Corollary 17. $U|+^n\rangle = \sum_{x \in \{0,1\}^n} e^{i\phi_x}|x\rangle$ for a uniformly random diagonal U form a $O(t^2/2^n)$-approximate state design [BS19]. The statement then follows directly from the triangle inequality:

$$\left\| \mathop{\mathbb{E}}_{U \sim \nu^{*m}} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} - \mathop{\mathbb{E}}_{U \sim \mu(d)} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} \right\|_1$$

$$\leq \left\| \mathop{\mathbb{E}}_{U \sim \nu^{*m}} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} - \mathop{\mathbb{E}}_{U \sim \mu_D(d)} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} \right\|_1$$

$$+ \left\| \mathop{\mathbb{E}}_{U \sim \mu(d)} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} - \mathop{\mathbb{E}}_{U \sim \mu_D(d)} (U|+^n\rangle\langle+^n|U^\dagger)^{\otimes t} \right\|_1$$

$$\leq \varepsilon + O(t^2/2^n).$$

□

As a consequence we can also show that HPS contains many almost orthogonal states, yielding additional evidence for the HPS assumption:

Corollary 18. *Let $m = 100nt$, $\delta = 1 - 2^{-n/8}$ and $t \leq 2^{n/2}$. For any fixed state $|\psi\rangle$, we have with probability $1 - 2^{-\Omega(nt)}$ over the matrix $\mathbf{A}$ that*

$$
\Pr_{\boldsymbol{\theta}} \left[|\langle \psi | \exp \left(\sum_{i=1}^{m} i\theta_i \bigotimes_{j=1}^{n} Z^{\boldsymbol{A}_{ij}} \right) |+^n\rangle|^2 \geq 1 - \delta \right] \leq 2^{-\Omega(nt)}. \tag{5.19}
$$

We defer the proofs of Theorem 16 and Corollary 18 to Appendix A.

5.3 Algorithms for Learning Hamiltonian Phase States

Recall that our (search) HPS assumption can be thought of as a state discrimination task. The goal is to recover the architecture $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ and the set of angles $\boldsymbol{\theta} \in \Theta_q^m$ given many copies of a random Hamiltonian phase state from the ensemble

$$
\mathcal{E} = \left\{ |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp \left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} \right) |+^n\rangle \right\}_{\mathbf{A} \in \mathbb{Z}_2^{m \times n}, \boldsymbol{\theta}=(\theta_1,\ldots,\theta_m) \in \Theta_q^m} .
$$

In this section, we consider various learning algorithms for the (search) HPS probem. We observe that the HPS problem does in fact have polynomial quantum sample complexity, and can thus be solved information-theoretically. However, as we also observe, all known learning algorithms have exponential time complexity, which suggests that the HSP problem cannot be solved efficiently.

We distinguish between the *private-key* and *public-key* setting: the former is essentially the learning task from Definition 1, whereas in the latter we further assume that the learner also has access to the architecture matrix $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$. We provide evidence that the learning tasks remains hard even if we reveal additional information about $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ and the goal is simply to guess the angles $\boldsymbol{\theta}$.

Sample complexity of HPS and hypothesis selection While we believe that HPS is a computationally hard problem, it can be solved information-theoretically with only polynomially many samples. In full generality, the problem of finding a fixed state ρ_j among many hypothesis states $\rho_1, \ldots, \rho_M$ is called quantum hypothesis testing. Currently, the best known general algorithm is threshold search as described in [BO21, Theorem 1.5] requires $n \log^2(M)$ copies improving over the bound from Ref. [Aar18]. For the HPS problem this implies an upper bound on the sample complexity of $O(n \log^2(q^m 2^{nm})) = O(n^3 m^2 \log(q))$. As the fidelities for pure states are PSD observables of rank 1, we can also use the shadow tomography protocol of Ref. [HKP20]. Given a secret state $|\Phi_{\boldsymbol{\theta}}^A\rangle$ allows

us to estimate the fidelities of all the $M = q^m 2^{nm}$ phase states up to an error of ε from $O(\log(M)/\varepsilon^2) = O(mn \log(q)/\varepsilon^2)$ samples. Then, a solver can simply list all estimated fidelities and pick the state with the largest overlap up to an error of ε.

We expect these bounds to be tight in the regime where $m \leq O(n^{\log(q)})$. For $m \to \infty$ better bounds are available at least for $q = 2^d$. In this case, the HPS instance generated b unitaries in the dth level of the Clifford hierarchy and it was proven in Ref. [ABDY23, Theorem 15] that for any state of the form

$$\exp\left(i \sum_{y \in \{0,1\}^n} a_y \bigotimes_{j=1}^{n} Z^{y_j} \right) |+^n\rangle \tag{5.20}$$

with $a_y \in \mathbb{Z}$ a circuit description can be learned with $O(n^d)$ copies using only measurements in the standard basis.

Learning algorithms for HPS with a public architecture.   In the special case when the architecture is public, our HPS assumption does in fact admit an optimal[3] but nevertheless exponential-time learning algorithm.

We consider the following state discrimination task, where the goal is to recover the set of angles $\boldsymbol{\theta}$ given many copies from the ensemble

$$\mathcal{E}_{\mathbf{A}} = \left\{ |\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} \right) |+^n\rangle \right\}_{\boldsymbol{\theta}=(\theta_1,...,\theta_m)\in\Theta_q^m}$$

where the matrix $\mathbf{A} \in \mathbb{Z}_2^{m \times n}$ is a random but fixed *architecture* which is known to the learner. This fits exactly into the framework of the pretty good measurement (PGM) [BK00, Mon19]. The ensemble $\mathcal{E}$ now turns out to be *geometrically uniform* because it can be written as $\mathcal{E}_{\mathbf{A}} = \left\{ U_{\boldsymbol{\theta}}^{\mathbf{A}} |+^n\rangle \right\}_{\boldsymbol{\theta}=(\theta_1,...,\theta_m)}$ where $\{U_{\boldsymbol{\theta}}^{\mathbf{A}}\}_{\boldsymbol{\theta}}$ is an Abelian group of matrices. Eldar and Forney [EJ00] showed that the PGM is optimal for all geometrically uniform ensembles, which implies that it is also optimal for our variant of the HPS problem. Nevertheless, despite the optimality, the best known algorithm for implementing pretty good measurements has exponential-time complexity in the size of the ensemble [GLM+22]. Consequently, we believe that the HPS problem remains computationally intractable, even if the architecture is public.

5.4 Discussion of Full Quantumness

We now provide some evidence that our HPS assumption is *plausibly* fully quantum. To begin, note that it is quite challenging (in general) to argue that any concrete assumption (one that names a specific problem as not being solvable in polynomial time) does not

---

[3]Here, we mean an algorithm that achieves the optimal success probability for a given number of copies.

imply one-way functions. Indeed, if one-way functions exist by other means, then every concrete assumption implies a one-way function independent of whether it is true.

One way that researchers get around this is to take some "idealized" version of the assumption, and show that relative to some oracle, the idealized version exists while one-way functions do not. The oracle separations shown in [Kre21, KQST23] already prove this for us. Namely, if we assume that our circuits actually output Haar random states, then with an additional PSPACE oracle, BQP $=$ QMA, but our idealized Haar random states are genuinely Haar random. However, using this as evidence that our assumption is fully-quantum is barely better than simply stating it to be so, since our argument essentially reduced to "nothing that claims to output Haar random-looking states can imply one-way functions". However, the fact that our ensembles are actually t-designs for a very high t should make one-way functions very hard to build from these states in practice. Indeed, examining the proof of [Kre21], we can see that if we instead take a family of t-designs instead of Haar random unitaries, any BQP algorithm that calls the oracle fewer than t times will have a similar concentration as the Haar measure, and thus will be simulatable by the PSPACE oracle (without access to the t-design). This implies that any construction of one-way functions that directly uses our assumption must make many (more than $t = \sqrt{m/n}$) calls to our oracle, implying that any such construction cannot be simple.

6 Applications

6.1 One-Way State Generators

The search variant of HPS allows for a practical and efficient implementation of one-way state generators (Definition 8). The following observation is immediate.

Theorem 1 (One-Way State Generators from HPS). *Let $n \in \mathbb{N}$ be the security parameter, and let $m = \mathsf{poly}(n)$ and $q = \mathsf{poly}(n)$. Then, under the (search) $\mathsf{HPS}_{n,m,q,\chi}$ assumption, where χ is the uniform distribution, there exist one-way state generators.*

6.2 Pseudorandom State Generators

In a similar vein, the decision variant of HPS gives rise to an efficient implementaion of pseudo-random states (Definition 9). The following observation is immediate.

Theorem 2 (Pseudorandom State Generators from HPS). *Let $n \in \mathbb{N}$ be the security parameter, and let $m = \mathsf{poly}(n)$ and $q = \mathsf{poly}(n)$. Then, under the (decision) $\mathsf{HPS}_{n,m,q,\chi}$ assumption, where χ is the uniform distribution, there exist pseudorandom state generators.*

6.3 Quantum Trapdoor Functions

As a stepping stone towards building public-key cryptography from the HPS assumption, we show how to construct quantum trapdoor functions [Col23]. These are one-way state generators that have a trapdoor that can be used to invert them easily.

Definition 3 (Quantum Trapdoor Function). *Let $n \in \mathbb{N}$ denote the security parameter. A quantum trapdoor function is a tuple of QPT algorithms* $(\mathsf{GenTrap}, \mathsf{GenEval}, \mathsf{Eval}, \mathsf{Invert})$:

- $\mathsf{GenTrap}(1^n) \to \mathsf{td}$: *on input 1^n, it outputs a classical trapdoor td.*

- $\mathsf{GenEval}(\mathsf{td}) \to |\xi_{\mathsf{td}}\rangle$: *on input td, it outputs a quantum evaluation state $|\xi_{\mathsf{td}}\rangle$.*

- $\mathsf{Eval}(|\xi_{\mathsf{td}}\rangle, x) \to |\psi_x\rangle$: *on input $|\xi_{\mathsf{td}}\rangle$ and $x \in \{0,1\}^n$, it outputs a state $|\psi_x\rangle$.*

- $\mathsf{Invert}(\mathsf{td}, |\psi_x\rangle) \to x'$: *on input td and $|\psi_x\rangle$, it outputs a string $x' \in \{0,1\}^n$.*

We require that the algorithms satisfy the following properties:

(a) ***Hardness of inversion:*** *For any $n \in \mathbb{N}$, for any QPT algorithm $\mathcal{A}$ and for any number of copies $m = \mathsf{poly}(n)$, it holds that*

$$\Pr\left[x \leftarrow \mathcal{A}(|\psi_x\rangle, |\xi_{\mathsf{td}}\rangle^{\otimes m}) : \begin{array}{c} x \sim \{0,1\}^n \\ \mathsf{td} \leftarrow \mathsf{GenTrap}(1^n) \\ |\xi_{\mathsf{td}}\rangle \leftarrow \mathsf{GenEval}(\mathsf{td}) \\ |\psi_x\rangle \leftarrow \mathsf{Eval}(|\xi_{\mathsf{td}}\rangle, x) \end{array}\right] \leq \mathsf{negl}(n).$$

(b) ***Trapdoor inversion:*** *For any $n \in \mathbb{N}$ and for any $x \in \{0,1\}^n$, it holds that*

$$\Pr\left[x \leftarrow \mathsf{Invert}(\mathsf{td}, |\psi_x\rangle) : \begin{array}{c} \mathsf{td} \leftarrow \mathsf{GenTrap}(1^n) \\ |\xi_{\mathsf{td}}\rangle \leftarrow \mathsf{GenEval}(\mathsf{td}) \\ |\psi_x\rangle \leftarrow \mathsf{Eval}(|\xi_{\mathsf{td}}\rangle, x) \end{array}\right] = 1.$$

Quantum trapdoor functions from HPS. Let us now construct QTFs from HPS using the following construction.

Construction 4. *Let $n \in \mathbb{N}$ be the security parameter and let $m = \mathsf{poly}(n)$ and $q = \mathsf{poly}(n)$. Then, we define the QTF construction* $(\mathsf{GenTrap}, \mathsf{GenEval}, \mathsf{Eval}, \mathsf{Invert})$ *as follows:*

- $\mathsf{GenTrap}(1^n) \to \mathsf{td}$: *on input 1^n, it outputs a trapdoor $\mathsf{td} = (\boldsymbol{\theta}, \mathbf{A})$, where*

$$\boldsymbol{\theta} \sim \Theta_q^m \qquad \text{and} \qquad \mathbf{A} \sim \mathbb{Z}_2^{m \times n}.$$

- $\mathsf{GenEval}(\mathsf{td}) \to |\xi_{\mathsf{td}}\rangle$: *on input td, it outputs the quantum evaluation state*

$$|\xi_{\mathsf{td}}\rangle = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}}\right) |+^n\rangle.$$

- Eval($|\tilde{\zeta}_{\text{td}}\rangle$, x) → $|\psi_x\rangle$: *on input $|\tilde{\zeta}_{\text{td}}\rangle$ and $x \in \{0,1\}^n$, it outputs the state*

$$|\psi_x\rangle = (Z^{x_1} \otimes \cdots \otimes Z^{x_n}) |\tilde{\zeta}_{\text{td}}\rangle.$$

- Invert(td, $|\psi_x\rangle$) → x': *on input td and $|\psi_x\rangle$, it applies $H^{\otimes n} U_{\boldsymbol{\theta}}^{\mathbf{A},\dagger}$ to $|\psi_x\rangle$, where*

$$U_{\boldsymbol{\theta}}^{\mathbf{A}} = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} \right).$$

Then, it measures all qubits in the computational basis and outputs the result x'.

Theorem 5. *Let $n \in \mathbb{N}$ be the security parameter, and let $m = \text{poly}(n)$ and $q = \text{poly}(n)$. Then, under the (decisional) $\text{HPS}_{n,m,q,\chi}$ assumption, where χ is the uniform distribution, Construction 4 yields a quantum trapdoor function.*

Proof. We need to verify two properties:

(a) *Hardness of inversion:* First, we invoke the decisional $\text{HPS}_{n,m,q,\chi}$ assumption,

$$\Pr\left[x \leftarrow \mathcal{A}(|\psi_x\rangle, |\tilde{\zeta}_{\text{td}}\rangle^{\otimes m}) \; : \; \begin{matrix} x \sim \{0,1\}^n \\ \text{td} \leftarrow \text{GenTrap}(1^n) \\ |\tilde{\zeta}_{\text{td}}\rangle \leftarrow \text{GenEval}(\text{td}) \\ |\psi_x\rangle \leftarrow \text{Eval}(|\tilde{\zeta}_{\text{td}}\rangle, x) \end{matrix} \right]$$
$$\leq \Pr\left[x \leftarrow \mathcal{A}(Z^x |\psi\rangle, |\psi\rangle^{\otimes m}) \; : \; \begin{matrix} x \sim \{0,1\}^n \\ |\psi\rangle \leftarrow \text{Haar}(2^n) \end{matrix} \right] + \text{negl}(n).$$

The claim then follows from [Col23, Lemma 5].

(b) *Trapdoor inversion:* This follows from the fact that Z^x commutes with the diagonal part of the Hamiltonian Phase state. In particular, using that

$$Z^x H^{\otimes n} = H^{\otimes n} X^x, \quad \forall x \in \{0,1\}^n,$$

we find that

$$\begin{aligned} Z^x U_{\boldsymbol{\theta}}^{\mathbf{A}} |+^n\rangle &= U_{\boldsymbol{\theta}}^{\mathbf{A}} Z^x |+^n\rangle \\ &= U_{\boldsymbol{\theta}}^{\mathbf{A}} Z^x H^{\otimes n} |0^n\rangle \\ &= U_{\boldsymbol{\theta}}^{\mathbf{A}} H^{\otimes n} X^x |0^n\rangle \\ &= U_{\boldsymbol{\theta}}^{\mathbf{A}} H^{\otimes n} |x\rangle. \end{aligned}$$

Therefore, inversion takes place with probability 1. $\qquad\square$

6.4 Public-Key Encryption with Quantum Keys

Due to the work of [Col23], quantum trapdoor functions imply a public-key encrpytion scheme, with quantum public keys.

Lemma 6 (Theorem 5 from [Col23]). *If quantum trapdoor functions exist, then there is a public-key encryption scheme with quantum public keys.*

Note, that the work of Coladangelo [Col23] required the use of (post-quantum) one-way functions. Using our construction of quantum trapdoor functions from HPS (which is a potentially even weaker assumption than the existence of one-way functions), we immediately get that public-key encryption scheme with quantum public keys exist.

Corollary 7. *Let $n \in \mathbb{N}$ be the security parameter, and $m = \text{poly}(n)$ and $q = \text{poly}(n)$. Then under the (decisional) $\text{HPS}_{n,m,q,\chi}$ assumption, where χ is the uniform distribution, there is a construction of public-key encryption with quantum public keys.*

Proof. The construction results from inserting Construction 4 from HPS into the construction of public-key encryption with quantum public keys from [Col23]. □

6.5 Quantum Pseudoentanglement

In this section, we show that we use ensembles of Hamiltonian Phase states to get a highly efficient construction of quantum pseudoentanglement, a notion that was introduced in the work of Aaronson et al. [ABF+24].

Definition 8 (Quantum pseudoentanglement). *Let $n \in \mathbb{N}$ be the security parameter. A $(f(n), g(n))$-pseudoentangled state ensemble (PES) consists of a pair of n-qubit state ensembles $\{|\psi_k\rangle\}_{k \in \mathcal{K}_n}$ and $\{|\phi\rangle_k\}_{k \in \mathcal{K}_n}$ with key space $\mathcal{K}_n$ such that:*

- *(Efficiency:) There exists a pair of uniform $\text{poly}(n)$-sized quantum circuits C_ψ, C_ϕ which, on input $k \in \mathcal{K}_n$, prepare the states $|\psi_k\rangle$ and $|\phi_k\rangle$, respectively.*

- *(Entanglement gap:) With high probability (i.e., at least $1 - 1/\text{poly}(n)$) over the choice of $k \sim \mathcal{K}_n$, the entanglement entropy between the first $\frac{n}{2}$ and last $\frac{n}{2}$ qubits of $|\psi_k\rangle$ (respectively, $|\phi_k\rangle$) is in the order of $\Theta(f(n))$ (respectively, $\Theta(g(n))$).*

- *(Computational indistinguishability:) For any integer $n \in \mathbb{N}$ and any number of copies $t = \text{poly}(n)$, the following mixed states are computationally indistinguishable,*

$$\rho = \mathop{\mathbb{E}}_{k \sim \mathcal{K}_n} \left[|\psi_k\rangle\langle\psi_k|^{\otimes t(n)} \right] \quad \approx_c \quad \sigma = \mathop{\mathbb{E}}_{k \sim \mathcal{K}_n} \left[|\phi_k\rangle\langle\phi_k|^{\otimes t(n)} \right].$$

In other words, for any efficient quantum distinguisher $\mathcal{A}$ which outputs a single-bit,

$$\left| \mathop{\text{Pr}}_{k \sim \mathcal{K}_n(1^n)} [\mathcal{A}(|\psi_k\rangle^{\otimes t(n)}) = 1] - \mathop{\text{Pr}}_{k \sim \mathcal{K}_n} [\mathcal{A}(|\phi_k\rangle^{\otimes t(n)}) = 1] \right| \leq \text{negl}(n).$$

Quantum Pseudoentanglement from the HPS assumption. Let $n \in \mathbb{N}$ be the security parameter. We can construct a (k, n)-pseudoentangled state ensemble (PES) from the HPS assumption as follows:

- (high-entanglement state) choose a parameter $m \geq n$, and let $\{|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle\}_{\mathbf{A} \in \mathbb{Z}_2^{m \times n}, \boldsymbol{\theta} \in \Theta_q^m}$ be the Hamiltonian phase state family with

$$
|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} \right) H^{\otimes n} |0^n\rangle .
$$

- (low-entanglement state) choose a parameter $k \leq n$, and let $\{|\Phi_{\boldsymbol{\omega}}^{\mathbf{B}}\rangle\}_{\mathbf{B} \in \mathbb{Z}_2^{k \times n}, \boldsymbol{\omega} \in \Theta_q^k}$ be the Hamiltonian phase state family with

$$
|\Phi_{\boldsymbol{\omega}}^{\mathbf{B}}\rangle = \exp\left(i \sum_{i=1}^{k} \omega_i \bigotimes_{j=1}^{n} Z^{\mathbf{B}_{ij}} \right) H^{\otimes n} |0^n\rangle .
$$

The goal of this section is to prove the following theorem.

Theorem 9 (Quantum pseudoentanglement from HPS). *Let $n \in \mathbb{N}$. Let $k \leq n \leq m$ with $k = \omega(\log(n))$ and $q = \mathsf{poly}(n)$ with $q > 4$. Assuming the hardness of the (decisional) HPS assumption $\mathsf{HPS}_{n,\mu,q,\chi}$, for $\mu \in \{k, m\}$ and where χ is the uniform distribution, the pair $\{|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle\}_{\mathbf{A} \in \mathbb{Z}_2^{m \times n}, \boldsymbol{\theta} \in \Theta_q^m}$ and $\{|\Phi_{\boldsymbol{\omega}}^{\mathbf{B}}\rangle\}_{\mathbf{B} \in \mathbb{Z}_2^{k \times n}, \boldsymbol{\omega} \in \Theta_q^k}$ form a (n, k)-pseudoentangled state ensemble.*

Before we give a proof of the statement, let us first introduce some relevant notation. For brevity, we define the diagonal part of of a Hamiltonian Phase state $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$ as

$$
D_{\boldsymbol{\theta}}^{\mathbf{A}} = \exp\left(i \sum_{i=1}^{m} \theta_i \bigotimes_{j=1}^{n} Z^{\mathbf{A}_{ij}} \right)
$$

Next, we define the $\frac{n}{2} \times \frac{n}{2}$ partition matrix $\Lambda_{\boldsymbol{\theta}}^{\mathbf{A}}$ such that, for $k, l \in \{0, 1\}^{\frac{n}{2}}$,

$$
\Lambda_{\boldsymbol{\theta}, (k,l)}^{\mathbf{A}} := \langle k, l| D_{\boldsymbol{\theta}}^{\mathbf{A}} |k, l\rangle .
$$

To get a handle on the entanglement entropy of the state, we have to analyze its reduced

state across a particular cut $(A : B)$, say the first $n/2$ and the last $n/2$ many qubits.

$$\rho_A = 2^{-n} \left(\sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \sum_{k \in \{0,1\}^{\frac{n}{2}}} \sum_{l \in \{0,1\}^{\frac{n}{2}}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,j)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(k,j)} \mathrm{Tr}_B[|i,j\rangle\langle k,l|] \right)$$

$$= 2^{-n} \left(\sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \sum_{k \in \{0,1\}^{\frac{n}{2}}} \sum_{l \in \{0,1\}^{\frac{n}{2}}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,j)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(k,j)} |i\rangle\langle k| \right)$$

$$= 2^{-n} \left(\sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \sum_{k \in \{0,1\}^{\frac{n}{2}}} \sum_{l \in \{0,1\}^{\frac{n}{2}}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,j)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(k,j)} |i\rangle\langle j|j\rangle\langle k| \right)$$

$$= 2^{-n} \left(\sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,j)} |i\rangle\langle j| \right) \left(\sum_{j \in \{0,1\}^{\frac{n}{2}}} \sum_{k \in \{0,1\}^{\frac{n}{2}}} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,k)} |j\rangle\langle k| \right)$$

$$= 2^{-n} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}} (\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^{\dagger}.$$

From the work of Aaronson et al. [ABF+24], we know that the entanglement entropy $S(\rho)$ is bounded below and above by the following quantities

$$-\log\left(\|2^{-n}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^{\dagger}\|_F \right) \leq S(\rho) \leq \log\left(\mathrm{rank}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^{\dagger}) \right)$$

We now give a proof of the aforementioned theorem.

Proof of Theorem 9. The efficiency of the state ensembles is immediate. For pseudorandomness, it suffices to invoke the (decisional) HPS assumption $\mathsf{HPS}_{n,\mu,q,\chi}$, for $\mu \in \{k, m\}$ and the uniform distribution χ the uniform distribution, to argue that the Hamiltonian Phase state pair $\{|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle\}_{\mathbf{A} \in \mathbb{Z}_2^{m \times n}, \boldsymbol{\theta} \in \Theta_q^m}$ and $\{|\Phi^{\mathbf{B}}_{\boldsymbol{\omega}}\rangle\}_{\mathbf{B} \in \mathbb{Z}_2^{k \times n}, \boldsymbol{\omega} \in \Theta_q^k}$ is computationally indistinguishable. This is essentially a consequence of the triangle inequality, since we each of the states is itself indistinguishable from a Haar random state.

Next, we attempt to get a handle on the average entanglement entropy for a general Hamiltonian Phase state. For convenience, we analyze the case for a general $m \in \mathbb{N}$, which we can later choose to be either $m \geq n$ or $m \leq n$ with $m = \omega(\log n)$. To get a lower bound on the entanglement entropy, we need to bound the expectation

$$\mathop{\mathbb{E}}_{\substack{\mathbf{A} \sim \mathbb{Z}_2^{m \times n} \\ \boldsymbol{\theta} \sim \Theta_q^m}} \left[\|2^{-n}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^{\dagger}\|_F^2 \right]$$

$$= 2^{-2n} \sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \mathop{\mathbb{E}}_{\substack{\mathbf{A} \sim \mathbb{Z}_2^{m \times n} \\ \boldsymbol{\theta} \sim \Theta_q^m}} \left[\left(\sum_{k \in \{0,1\}^{\frac{n}{2}}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,k)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,k)} \right)^2 \right] \tag{6.10}$$

$$= 2^{-2n} \sum_{i \in \{0,1\}^{\frac{n}{2}}} \sum_{j \in \{0,1\}^{\frac{n}{2}}} \sum_{k,l \in \{0,1\}^{\frac{n}{2}}} \mathop{\mathbb{E}}_{\substack{\mathbf{A} \sim \mathbb{Z}_2^{m \times n} \\ \boldsymbol{\theta} \sim \Theta_q^m}} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,k)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,k)} \boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,l)} \overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,l)}.$$

Consider the expression

$$\underset{\substack{\mathbf{A}\sim\mathbb{Z}_2^{m\times n}\\ \boldsymbol{\theta}\sim\Theta_q^m}}{\mathbb{E}}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,k)}\overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,k)}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta},(i,l)}\overline{\boldsymbol{\Lambda}}^{\mathbf{A}}_{\boldsymbol{\theta},(j,l)}$$

$$=\underset{\substack{\mathbf{A}\sim\mathbb{Z}_2^{m\times n}\\ \boldsymbol{\theta}\sim\Theta_q^m}}{\mathbb{E}}\exp\left[\mathrm{i}\sum_{r=1}^m\theta_r((-1)^{\langle i,k|A^r\rangle}+(-1)^{\langle j,k|A^r\rangle}-(-1)^{\langle i,l|A^r\rangle}-(-1)^{\langle j,l|A^r\rangle})\right]$$

$$=\left(\underbrace{\underset{\substack{\mathbf{A}\sim\mathbb{Z}_2^{1\times n}\\ \boldsymbol{\theta}\sim\Theta_q^1}}{\mathbb{E}}\exp\left[\mathrm{i}\theta((-1)^{\langle i,k|A\rangle}+(-1)^{\langle j,k|A\rangle}-(-1)^{\langle i,l|A\rangle}-(-1)^{\langle j,l|A\rangle})\right]}_{=:e_{i,j,k,l}}\right)^m.$$

(6.11)

Clearly, if $\{(i,k),(j,k)\}=\{(i,l),(j,l)\}$ then $e_{i,j,k,l}=1$. We can upper bound the number of all tuples with this property by $2^{3n/2}\times 2!$, which leads to a contribution of $2\times 2^{-n/2}$ in Eq. (6.10). Assume that $q>4$. Then, for all other tuples (i,j,k,l) we notice that $e_{i,j,k,l}$ is an eigenvalue <1 of the moment operator $\mathbb{E}U^{\otimes 2}\otimes\overline{U}^{\otimes 2}$ for $U=e^{\mathrm{i}\otimes_{j=1}^n Z^{A_{1j}}}$. Then, we can invoke Lemma 4 with $t=2$ to obtain the bound $e_{i,j,k,l}\leq 1-\frac{1}{2t}=\frac{3}{4}$. Plugging this back into Eq. (6.10) yields the simple upper bound

$$\underset{\substack{\mathbf{A}\sim\mathbb{Z}_2^{m\times n}\\ \boldsymbol{\theta}\sim\Theta_q^n}}{\mathbb{E}}\left[\|2^{-n}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^\dagger\|_F^2\right]\leq 2\times 2^{-\frac{n}{2}}+\left(\frac{3}{4}\right)^m.$$

(6.12)

In particular, for $m\geq n$ we find with Markov's inequality that

$$S(\rho)\geq-\log\left(\|2^{-n}\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}}(\boldsymbol{\Lambda}^{\mathbf{A}}_{\boldsymbol{\theta}})^\dagger\|_F\right)=\Omega(n)$$

(6.13)

with probability $1-2^{-\Omega(n)}$.

To obtain a matching upper bound on the entanglement entropy observe that each unitary $e^{\mathrm{i}\phi\otimes_{j=1}^n Z^{y_j}}$ can be decomposed into a single qubit Z rotation and at most $2n$ CNOT and SWAP gates such that at most 2 CNOT or SWAP gates act across the cut $(A:B)$. In particular, CNOT has an operator Schmidt rank of 2 and SWAP has an operator Schmidt rank of 4. Therefore, any robation of the form $e^{\mathrm{i}\phi\otimes_{j=1}^n Z^{y_j}}$ can increase the Schmidt rank of a state across the cut $(A:B)$ by at most $4\times 4=16$. Overall, we find for the entanglement entropy

$$-\frac{1}{2}\log(2\times 2^{-n}+(3/4)^m)\leq S(\rho_A)\leq\log(\mathrm{rank}(\rho_A))\leq\log(16)m.$$

(6.14)

The theorem then follows from choosing the values $m=n$ and $m=\omega(\log(n))$. $\qquad\square$

6.6 Pseudorandom Unitaries

In this section, we show that a variant of our HPS assumption allow us construct *pseudorandom unitaries*. We first show that a unitary that repeatedly applies (uniformly random) phase oracles and Hadamards is indistinguishable from a truly Haar random unitary.

Let us first give a formal definiton of pseudorandom unitaries, as proposed by Ji, Liu, and Song [JLS18]. Broadly speaking, these are efficient ensembles of unitaries which emulate Haar random unitaries to computationally bounded adversaries.

Definition 15 (Pseudorandom unitary). *Let $n \in \mathbb{N}$ be the security parameter. An infinite sequence $\mathcal{U} = \{\mathcal{U}_n\}_{n \in \mathbb{N}}$ of n-qubit unitary ensembles $\mathcal{U}_n = \{U_k\}_{k \in \mathcal{K}}$ is a called a family of pseudorandom unitaries if it satisfies the following conditions.*

- *(**Efficient computation**) There exists an efficient quantum algorithm $\mathcal{Q}$ such that for all keys $k \in \mathcal{K}$, where $\mathcal{K}$ denotes the key space, and any $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$, it holds that*

$$\mathcal{Q}(k, |\psi\rangle) = U_k |\psi\rangle .$$

- *(**Pseudorandomness**) The unitary U_k, for a randomly chosen key $k \sim \mathcal{K}$, is computationally indistinguishable from a Haar random unitary $U \sim \mu(2^n)$. In other words, for any efficient quantum algorithm $\mathcal{A}$, it holds that*

$$\left| \Pr_{k \sim \mathcal{K}} [\mathcal{A}^{U_k}(1^n) = 1] - \Pr_{U \sim \mu(2^n)} [\mathcal{A}^U(1^n) = 1] \right| \leq \mathsf{negl}(n) .$$

Note that, whenever we write $\mathcal{U}_n = \{U_k\}_{k \in \mathcal{K}}$, it is implicit that the key space $\mathcal{K}$ depends on the security parameter $n \in \mathbb{N}$, and that the length of each key $k \in \mathcal{K}$ is polynomial in n.

6.6.1 Security of the FHFHFC pseudo-random unitary

A recent line of work has shown that a "PFC" ensemble [MPSY24], where P and F are a pseudo-random permutation and pseudo-random function respectively (and C is a uniformly random Clifford) is a pseudo-random unitary [MH24]. However, as these require one-way functions, there are no suitable post-quantum instantiations of pseudo-random unitaries. An older folklore construction of pseudo-random unitaries, first proposed by [JLS18], is to apply alternating pseudo-random phase unitaries and Hadamards. In this section, we prove that a version of this construction based on the HPS assumption is secure, and propose an instantiation of these pseudo-random unitaries using HPS oracles.

Let F a uniformly random diagonal matrix. We can think of F being an "ideal HPS" phase oracle, which acts as follows for a tuple of angles $\theta = (\theta_z)_z$:

$$\mathsf{F}_\theta = \sum_x \exp\left(i \sum_z \theta_z (-1)^{\langle x, z \rangle} \right) |x\rangle\langle x| .$$

Here, each angle θ_z will be randomly drawn from a discrete set $\{2\pi k/q\}_{k=0}^{q-1}$. In the following, we will choose $q = 2^n$. We note that any exponentially large q suffices. We make two

observations before continuing with the proof. First, we call this an ideal HPS because it involves sampling an exponential number of real phases, instead of only a polynomial number of them. Second, we note that these matrices are identically distributed to uniformly random diagonal matrices with complex phases, but we phrase it as an ideal HPS instance to motivate our conjectured pseudo-random unitary construction from HPS.

Remark 16. *As a matter of notation, we will write θ, or $\theta^{(1)}$, to mean the tuple of angles $(\theta_1, \ldots, \theta_{2^n})$ (or with the appropriate superscript). The association between the list of angles associated with a θ will be clear from context.*

We claim that the following unitary is pseudo-random, for three uniformly random diagonal unitaries with complex phases, F_1, F_2 and F_3:

$$F_3 \cdot H^{\otimes n} \cdot F_2 \cdot H^{\otimes n} \cdot F_1 \cdot C.$$

We call this the FHFHFC-ensemble, and by the previous observation, this is identical to the proposed construction from [JLS18]. The following is the main theorem of this section.

Theorem 17. *Let $n \in \mathbb{N}$ be the number of qubits as input to a quantum adversary. Then for all t-query adversaries $\mathcal{A}_t^{(\cdot)}$ making queries to a n qubit oracle, the following holds*

$$\mathsf{td}\left(\mathbb{E}_{U \sim \mu(2^n)}\left[\mathcal{A}_t^U(|0\rangle)\right], \mathbb{E}_{\mathcal{O} \sim \mathsf{FHFHFC}}\left[\mathcal{A}_t^{\mathcal{O}}(|0\rangle)\right]\right) \leq O\left(\sqrt{\frac{t^3}{2^n}}\right).$$

In order to prove that these are pseudo-random, we rely on the results of [MH24], which showed that a certain isometry, called the path-recording isometry, is indistinguishable from a Haar random unitary for bounded-query adversaries. We state the definition of the path-recording isometry here.

Definition 18 (Path recording isometry). *Define the path recording isometry* PR, *which acts on two registers, an input* A, *and database* XY *containing t input-output pairs, as follows:*

$$\mathsf{PR}\,|x\rangle_\mathsf{A}\,|R\rangle_\mathsf{XY} \mapsto \frac{1}{\sqrt{2^n - t}} \sum_{y \notin \mathrm{Im}(R)} |y\rangle_\mathsf{A}\,|R \cup (x, y)\rangle_\mathsf{XY}.$$

Before we begin with the proof, we comment on the construction itself: Why take three rounds of F? Intuitively, applying a single random function F records the input to the function, as noted in [Zha19]. Therefore, we can see that each of the three random phase oracles has a distinct "job": The first one records the input, and the last one records the output. The middle one records the relationship between the input and output, as we will see later. Where our proof differs from the proof of [MH24] is that the middle unitary is not able to record all *bijective* relations, but instead can only record so called "collision free" relations. We prove that, with high probability, a bijection with random output is collision free, which allows us to perform path recording with slightly more error than incurred in [MH24].

As the first step of the proof, we will imagine a purified version of the FHFHFC ensemble, where the phases are stored in a register that is not visible to the adversary. We will show that the FHFHF oracle exhibits behavior similar to the path recording isometry, up to an isometry that acts on the purifying register containing the phases.

We first argue that we can purify the action of FHFHF into another oracle we denote by FHFHFO.

Definition 19. FHFHFO *is a unitary oracle that acts on two registers as follows:*

$$\text{FHFHFO} \, |x\rangle \, |\theta^{(1)}, \theta^{(2)}, \theta^{(3)}\rangle = \mathsf{F}_{\theta^{(3)}} \cdot H^{\otimes n} \cdot \mathsf{F}_{\theta^{(2)}} \cdot H^{\otimes n} \cdot \mathsf{F}_{\theta^{(1)}} \, |x\rangle \, |\theta^{(1)}, \theta^{(2)}, \theta^{(3)}\rangle \, .$$

Lemma 20 (Purification technique for FHFHF). *The view of any adversary querying* FHF *is equivalent to the view of the adversary querying* FHFO *when the purifying register is initialized in the following state:*

$$\frac{1}{2^{3n}} \sum_{\substack{\theta_1^{(1)}, \dots, \theta_{2^n}^{(1)} \\ \theta_1^{(2)}, \dots, \theta_{2^n}^{(2)} \\ \theta_1^{(3)}, \dots, \theta_{2^n}^{(3)}}} |\theta^{(1)}, \theta^{(2)}, \theta^{(3)}\rangle \, .$$

Proof. Since the view of the adversary does not contain the purifying register, when we trace them out we will get a uniform mixture of states corresponding to FHFHF for uniformly random angles $\theta^{(1)}$, $\theta^{(2)}$, and $\theta^{(3)}$. This corresponds exactly to choosing $\theta^{(1)}$, $\theta^{(2)}$, and $\theta^{(3)}$ uniformly at random. $\qquad\square$

We now define a family of orthogonal states that correspond to a database R. We some new notation here to make the remaining equations less cumbersome. For a set S of n-bit strings and n-bit string z, define S_z to be:

$$S_z = \sum_{x \in S} (-1)^{\langle x, z \rangle} \, .$$

For a relation R, define the following sets derived from R:

$$\text{Dom}(R) = \{x : (x, y) \in R\}$$
$$\text{Im}(R) = \{y : (x, y) \in R\} \, .$$

As a first step, we define so-called FHFHF-set states, as follows:

Definition 21 (FHFHF-set states). *Given a set S, define the set $|\mathsf{fhfhf}_S\rangle$ as follows*

$$\frac{1}{2^n} \sum_{\theta_1, \dots, \theta_{2^n} \in [2^n]} \exp\left(i \sum_z \theta_z S_z \right) |\theta\rangle \, .$$

The following lemma will show that these set states are orthogonal whenever the sets are different.

We show that these states form an orthogonal basis. In order to prove this, we first prove the following property, which comes from Fourier analysis of Boolean functions.

Lemma 22. *Let X and Y be two sets of $t < 2^{n-1}$ elements that are not equal, then there exists a string z such that*

$$\sum_{x \in X}(-1)^{\langle z,x \rangle} \neq \sum_{y \in Y}(-1)^{\langle z,y \rangle}.$$

Proof. The value $\sum_x(-1)^{\langle z,x \rangle}$ is exactly the z'th Fourier coefficient of a Boolean function that takes value 1 in x. Since two functions have the same Fourier coefficients if and only if they are the same function, if no z exists, then we get a contradiction with the fact that x and y are different. Thus, there must exist a z for which these values differ. $\square$

We note that this proof stops working when t is allowed to be 2^{n-1}, as two sets whose union is $\{0,1\}^n$ will have identical Fourier coefficients.

Lemma 23 (Orthogonality of set states). *Set S, S' be two sets of n-bit strings, then the set states $|\mathrm{fhfhf}_S\rangle$ and $|\mathrm{fhfhf}_{S'}\rangle$ are orthogonal.*

Proof. Let z^* be the bit string on which S_z and S'_z differ (given by Lemma 22). Then we have the following

$$\langle \mathrm{fhfhf}_S | \mathrm{fhfhf}_{S'} \rangle = \frac{1}{2^{2n}} \sum_{\theta_1,\ldots,\theta_{2^n}} \left(i \sum_z \theta_z (S_z - S'_z) \right)$$

$$= \frac{1}{2^{2n}} \sum_{\theta_{z^*}} \exp\left(i\theta_{z^*} S_{z^*} - S'_{z^*} \right) \cdot \sum_{\theta_i | i \neq z^*} \exp\left(i \sum_{z \neq z^*} \theta_z (S_z - S'_z) \right)$$

$$= 0.$$

Here we use the fact that the following equality:

$$\mathbb{E}_\theta e^{i\theta x} = \begin{cases} 1 & x = 0 \\ 0 & \text{otherwise} \end{cases}.$$

Moreover, it is easy to check tht $\mathbb{E}_k e^{i2\pi k/q} = \mathbb{E}_\theta e^{i\theta x}$ for an integer x with $|x| < q$. Since $S_{z^*} - S'_{z^*}$ is not 0, this term evalutes to 0 and therefore the entire product is 0. $\square$

Now we define the FHFHF-relation state, which will be how the FHFHF oracle simulates the path recording oracle.

Definition 24 (FHFHF-relation states). *For $0 \leq t \leq 2^n$, and a relation R, define the* FHFHF-

42

relation states to be as follows

$$|\text{fhfhf}_R\rangle = \frac{1}{2^{nt/2}} \sum_{w_1,\dots,w_t \in \{0,1\}^n} \left(\prod_i (-1)^{\langle w_i, y_i \oplus x_i \rangle} \right)$$

$$\frac{1}{2^{3n}} \sum_{\substack{\theta_1^{(1)},\dots,\theta_{2^n}^{(1)} \\ \theta_1^{(2)},\dots,\theta_{2^n}^{(2)} \\ \theta_1^{(3)},\dots,\theta_{2^n}^{(3)}}} \exp\left(i \sum_{z,i} \theta_z^{(1)} (-1)^{\langle z, x_i \rangle} + \theta_z^{(2)} (-1)^{\langle z, w_i \rangle} + (-1)^{\langle z, y_i \rangle} \theta_z^{(3)} \right)$$

$$|\theta^{(1)}\rangle \otimes |\theta^{(2)}\rangle \otimes |\theta^{(3)}\rangle \, .$$

Remark 25. *We note that in terms of the set states, we can re-write the state $|\text{fhfhf}_R\rangle$ as follows:*

$$|\text{fhfhf}_R\rangle = |\text{fhfhf}_{\text{Dom}(R)}\rangle \left(\frac{1}{2^{nt/2}} \sum_{w_1,\dots,w_t} \prod_i \left((-1)^{\langle w_i, x_i \oplus y_i \rangle} \right) |\text{fhfhf}_{\{w_1,\dots,w_t\}}\rangle \right) |\text{fhfhf}_{\text{Im}(R)}\rangle \, .$$

Unlike in the analysis of the PFC-ensemble, the FHFHFC-relation states will not work on bijections, but on a slightly more restrictive set of relations, defined below.

Definition 26 (Collision-free relation). *Let $x_1, \dots x_t$ and $y_1, \dots, y_t$ be a collection of n-bit string. Then we say the pair is called "collision-free" if the following holds. For all permutations $\pi \in S_t$,*

$$\{x_i \oplus y_i\} \neq \{x_i \oplus y_{\pi(i)}\} \, .$$

Given a set of $x_1, \dots, x_t$, let $\text{CF}_\mathbf{x}$ be the set of collision free $y_1, \dots, y_t$.

Lemma 27 (Collision probability bound). *Fix distinct $x_1, \dots, x_t$, then the probability over uniformly chosen, distinct $y_1, \dots, y_t$ that the pair of sets is not collision free is at most $t^3/2^n$.*

Proof. We claim that being a collision-free relation implies that there exists a pair of y_i and y_j such that $y_i \oplus y_j$ is equal to $x_1 \oplus x_j$ (here note that the second indices are the same). We will then upper bound the probability that any of the t^2 possible $y_i \oplus y_j$ (for randomly chosen y_i and y_j) is equal to one of the at most t many $x_1 \oplus x_j$.

To see this, assume for the sake of contradiction that there is a permutation that breaks the collision-free property. Then it must be the case that $x_1 \oplus y_{\pi(1)} = x_a \oplus y_a$ for some index a, we see that the following holds:

$$x_1 \oplus y_1 \oplus x_1 \oplus y_{\pi(1)} = x_1 \oplus y_1 \oplus x_a \oplus y_a$$
$$\implies y_{\pi(1)} \oplus y_a = x_1 \oplus x_a \, .$$

Since we assume that $x_1, \dots, x_t$ are fixed, the probability that the XOR of a single pair of y_i, y_j fall in the set $\{x_1 \oplus x_j : j \in [t]\}$ is at most $t/2^n$. By a union bound, the probability that any of them fall into this set is at most $t^3/2^n$. By this argument, if none of the XORs of y_i and y_j fall into the XOR of all pairs of x's, the set must be collision free. $\square$

Remark 28. *In general, one might hope that by not just considering collision free bijections, but any set of x's and y's for which the ordering can be determined uniquely by the list of XORs of the two tuples, one can remove the need for the Clifford unitary at the end. However, we leave this for future work.*

Lemma 29 (Orthogonality of relation states). *Let R and R' be two different collision-free relations of size at most 2^{n-1} elements. Then $|\text{fhfhf}_R\rangle$ and $|\text{fhfhf}_{R'}\rangle$ are orthogonal.*

Proof. From Lemma 23, together with Remark 25, if the image or domain of the two relations is different, then the states will be orthogonal since the set states storing the image and domain will differ. Therefore, we only have to consider pairs of collision free relations that have the same inputs and outputs. However, by assumption these are collision free. Say that these two relations are related by a permutation π on the $y_1, \ldots, y_t$, and say that y_{i*} is the index on which $x_i \oplus y_i \neq x_i \oplus y_{\pi(i)}$, which exists because of the collision free property.

$$\langle \text{fhfhfh}_R | \text{fhfhfh}_{R'} \rangle = \frac{1}{2^{nt}} \sum_{\substack{w_1,\ldots,w_t \\ w'_1,\ldots,w'_t}} \prod_i (-1)^{\langle w_i,(x_i \oplus y_i)\rangle + \langle w'_i,(x_i \oplus y_{\pi(i)})\rangle} \langle \text{fhfhf}_{\{w_1,\ldots,w_t\}} | \text{fhfhf}_{\{w'_1,\ldots,w'_t\}} \rangle$$

$$= \frac{1}{2^{nt}} \sum_{w_1,\ldots,w_t} \prod_i (-1)^{\langle w_i,(x_i \oplus y_i) + (x_i \oplus y_{\pi(i)})\rangle}$$

$$= \frac{1}{2^{nt}} \left(\sum_{w_{i*}} (-1)^{\langle w_{i*},(x_{i*} \oplus (y_{i*} + y_{\pi(i*)}))\rangle} \right) \cdot \left(\sum_{\substack{w_1,\ldots,w_{i*-1}, \\ w_{i*+1},\ldots,w_t}} \prod_{i \neq i*} (-1)^{\langle w_i,(x_i \oplus y_i) - (x_i \oplus y_{\pi(i)})\rangle} \right)$$

$$= 0.$$

To get to the last line, we use the fact that in the first term, $x_{i*} \oplus (y_{i*} - y_{\pi_{i*}}) \neq 0$, so summing over the values of w_{i*}, half will have a sign of $+1$ and half will have a sign of -1, yielding 0. $\qquad\square$

Now, we have shown that there is an orthogonal set of states that correspond to relations. The next step is to show that the FHF oracle, when run on a state $|x\rangle |\text{FHF}_R\rangle$, actually acts to append (x, y) to the relation state.

Lemma 30. *For $t = \text{poly}(n)$, relation $R = \{(x_1, y_1), \ldots, (x_t, y_t)\}$, and $x \in \{0,1\}^n$, the following holds*

$$\text{FHFHFO} |x\rangle |\text{fhf}_R\rangle = \frac{1}{2^{n/2}} \sum_y |y\rangle |\text{fhfhf}_{R \cup \{(x,y)\}}\rangle .$$

Proof. The proof proceeds by evaluating the action of FHF on the input state. Let $\alpha_R^{\Theta_1,\Theta_2,\Theta_3}$ be the coefficient that appears in the integral in the definition of $|\text{FHF}_R\rangle$. We first apply

44

the first function, yielding the following

$$|x\rangle \frac{1}{2^{3n}} \frac{1}{2^{3n}} \sum_{\substack{\theta_1^{(1)},\dots,\theta_{2n}^{(1)} \\ \theta_1^{(2)},\dots,\theta_{2n}^{(2)} \\ \theta_1^{(3)},\dots,\theta_{2n}^{(3)}}} \alpha_R^{\theta^{(1)},\theta^{(2)},\theta^{(3)}} \exp\left(i\sum_z \theta_z^{(1)}(-1)^{\langle x,z\rangle}\right) |\theta^{(1)},\theta^{(2)},\theta^{(3)}\rangle \ .$$

Then we apply the Hadamard matrix, which yields the following state (without normalization)

$$\frac{1}{2^{n/2}} \sum_{y\in\{0,1\}^n} (-1)^{\langle x,w\rangle} |w\rangle \frac{1}{2^{3n}} \sum_{\substack{\theta_1^{(1)},\dots,\theta_{2n}^{(1)} \\ \theta_1^{(2)},\dots,\theta_{2n}^{(2)} \\ \theta_1^{(3)},\dots,\theta_{2n}^{(3)}}} \alpha_R^{\theta^{(1)},\theta^{(2)},\theta^{(3)}} \exp\left(i\sum_z \theta_z^{(1)}(-1)^{\langle x,z\rangle}\right) |\theta^{(1)},\theta^{(2)},\theta^{(3)}\rangle \ .$$

Finally, we apply the second ideal HPS to the state, yielding the following

$$\frac{1}{2^{n/2}} \sum_{w\in\{0,1\}^n} (-1)^{\langle x,w\rangle} |w\rangle \sum_{\substack{\theta_1^{(1)},\dots,\theta_{2n}^{(1)} \\ \theta_1^{(2)},\dots,\theta_{2n}^{(2)} \\ \theta_1^{(3)},\dots,\theta_{2n}^{(3)}}} \alpha_R^{\theta^{(1)},\theta^{(2)},\theta^{(3)}} \exp\left(i\sum_z \theta_z^{(1)}(-1)^{\langle x,z\rangle} + \theta_z^{(2)}(-1)^{\langle w,z\rangle}\right) |\theta^{(1)},\theta^{(2)},\theta^{(3)}\rangle \ .$$

Finally, we apply another Hadamard and a final ideal HPS unitary to get the following state.

$$\frac{1}{2^{n/2}} \sum_{w,y\in\{0,1\}^n} (-1)^{\langle y,x\oplus w\rangle} |y\rangle \frac{1}{2^{3n}} \sum_{\substack{\theta_1^{(1)},\dots,\theta_{2n}^{(1)} \\ \theta_1^{(2)},\dots,\theta_{2n}^{(2)} \\ \theta_1^{(3)},\dots,\theta_{2n}^{(3)}}} \alpha_R^{\theta^{(1)},\theta^{(2)},\theta^{(3)}}.$$

$$\exp\left(i\sum_z \theta_z^{(1)}(-1)^{\langle x,z\rangle} + \theta_z^{(2)}(-1)^{\langle w,z\rangle} + \theta_z^{(3)}(-1)^{\langle y,z\rangle}\right) |\theta^{(1)},\theta^{(2)},\theta^{(3)}\rangle \ .$$

Expanding out the definition of $\alpha_R^{\theta^{(1)},\theta^{(2)},\theta^{(3)}}$, pushing the sum over w into the purifying register, and combining the product of exponentials yields the desired state. $\qquad\square$

A subtle difference between this version and Claim 5 of [MH24] is the difference in normalization. However, since the size of the relation is polynomial in n, this will contribute a trace distance error of at most $t^2/2^n$, which we account for later.

Now, we have shown that the relation states for the FHFHF oracle are orthogonal whenever the relation is collision free, and that the action of the FHF oracle is to output a random y and store (x,y) in the relation state. Now, we make the connection to

the path recording oracle formal by defining two isometries that will act on the purifying register, Compress and Collide. The compress isometry acts as follows:

$$\text{Compress} = \sum_{R \text{ collision free}} |R\rangle\langle \text{fhfhf}_R| \, .$$

The biject isometry will act as follows.

$$\text{Collide} = \sum_{R \text{ collision free}} |\text{fhfhf}_R\rangle\langle \text{fhfhf}_R| \, .$$

Let $\Pi^{(t)}$ be the projection onto bijection *relation states*, as defined in [MH24], and $\Lambda^{(t)}$ be the projection onto *collision free* relation states. Then we have the following lemma.

Lemma 31. *For all adversaries $\mathcal{A}_t^{\mathcal{O}}$ that make t queries, let $|\mathcal{A}_t^{\mathcal{O}}\rangle$ be the state of the adversary after making t queries to $\mathcal{O}$. Then for all n-qubit unitaries G, the following holds.*

$$\text{td}\left(\text{Compress} \cdot \text{Collide} \cdot |\mathcal{A}_t^{\text{FHFHFO} \cdot G}\rangle\langle\mathcal{A}_t^{\text{FHFHFO} \cdot G}| \cdot \text{Collide} \cdot \text{Compress}, \Lambda^{(t)} \cdot |\mathcal{A}_t^{\text{PR} \cdot G}\rangle\langle\mathcal{A}_t^{\text{PR} \cdot G}| \cdot \Lambda^{(t)}\right)$$
$$\leq O\left(\frac{t^2}{2^n}\right) \, .$$

Proof. First, note that collision-free relations are a subset of bijective relations, and the projectors $\Pi^{(t)}$ and $\Lambda^{(t)}$ commute, which $\Lambda^{(t)} = \Lambda^{(t)}\Pi^{(t)}$. Similar to the proof from [MH24], we can explicitly write out the states $|\mathcal{A}_t^{\mathcal{O}}\rangle$ in either case. For a set S, let S^t denote the t-fold cartesian product of the set, and S_{dist}^t be the distinct subset of S^t (i.e. all sets of t elements that are distinct from each other). First, we note the following

$$|\mathcal{A}_t^{\text{PR} \cdot G}\rangle = \sqrt{\frac{(2^n - t)!}{2^n!}} \sum_{\substack{x_1, \ldots, x_t \in (\{0,1\}^n)^t \\ y_1, \ldots, y_t \in (\{0,1\}^n)_{\text{dist}}^t}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i) |0\rangle\right) \otimes |\{(x_i, y_i)\}_{i=1}^{t}\rangle \, .$$

Performing the projection onto bijective relations yields the following state:

$$\Pi^{(t)} \cdot |\mathcal{A}_t^{\text{PR} \cdot G}\rangle = \sqrt{\frac{(2^n - t)!}{2^n!}} \sum_{\substack{x_1, \ldots, x_t \in (\{0,1\}^n)_{\text{dist}}^t \\ y_1, \ldots, y_t \in (\{0,1\}^n)_{\text{dist}}^t}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i) |0\rangle\right) \otimes |\{(x_i, y_i)\}_{i=1}^{t}\rangle \, .$$

Further performing the projection onto collision free relations yields the following:

$$\Pi^{(t)} \cdot |\mathcal{A}_t^{\text{PR} \cdot G}\rangle = \sqrt{\frac{(2^n - t)!}{2^n!}} \sum_{\substack{x_1, \ldots, x_t \in (\{0,1\}^n)_{\text{dist}}^t \\ y_1, \ldots, y_t \in \text{CF}_{x_1, \ldots, x_t}}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i) |0\rangle\right) \otimes |\{(x_i, y_i)\}_{i=1}^{t}\rangle \, .$$

Now we turn to the FHF oracle. From the previous lemma, we have the following.

$$
|\mathcal{A}_t^{\mathsf{FHFHFO}\cdot G}\rangle = \sqrt{\frac{1}{2^{nt}}} \sum_{\substack{x_1,\dots,x_t\in(\{0,1\}^n)^t \\ y_1,\dots,y_t\in(\{0,1\}^n)^t}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i)\,|0\rangle \right) \otimes |\mathsf{fhfhf}_{\{(x_i,y_i)\}_{i=1}^t}\rangle\ .
$$

Applying Biject to this state, we get exactly the following state.

$$
\mathsf{Collide}\cdot|\mathcal{A}_t^{\mathsf{FHFHFO}\cdot G}\rangle = \sqrt{\frac{1}{2^{nt}}} \sum_{\substack{x_1,\dots,x_t\in(\{0,1\}^n)^t_{\mathrm{dist}} \\ y_1,\dots,y_t\in\mathsf{CF}_{x_1,\dots,x_t}}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i)\,|0\rangle \right) \otimes |\mathsf{fhfhf}_{\{(x_i,y_i)\}_{i=1}^t}\rangle
$$
$$
+ \mathsf{Collide}\cdot \sqrt{\frac{1}{2^{nt}}} \sum_{\substack{x_1,\dots,x_t\in(\{0,1\}^n)^t_{\mathrm{dist}} \\ y_1,\dots,y_t\notin\mathsf{CF}_{x_1,\dots,x_t}}} \left(\prod_{i=1}^{t} (|y_i\rangle\langle x_i| \cdot G \cdot A_i)\,|0\rangle \right) \otimes |\mathsf{fhfhf}_{\{(x_i,y_i)\}_{i=1}^t}\rangle\ .
$$

Since we showed that collision free relation states are orthogonal to each other, compress will map the first term to the correct relation states, and shrink the trace of the second term. However, since the probability that a y is not collision free relative to x is at most $t^3/2^n$, the trace of the second term is at most $t^3/2^n$, and the normalization difference in the first term contributes at most a $t^2/2^n$ trace distance error. Applying the triangle inequality on these two steps, we get the desired error bound. $\square$

Next, we show that the states after projecting onto $\Pi^{(t)}$ (bijections) and $\Lambda^{(t)}$ (collision free relations) are close to each other in trace distance.

Lemma 32. *For all t-query adversaries $\mathcal{A}_t^{\mathcal{O}}$ and unitaries G, the following bound holds:*

$$
\mathrm{td}\left(\Lambda^{(t)}\cdot|\mathcal{A}^{\mathsf{PR}\cdot G}\rangle\langle\mathcal{A}^{\mathsf{PR}\cdot G}|\cdot\Lambda^{(t)}, \Pi^{(t)}\cdot|\mathcal{A}^{\mathsf{PR}\cdot G}\rangle\langle\mathcal{A}^{\mathsf{PR}\cdot G}|\cdot\Pi^{(t)} \right) \le O\left(\sqrt{\frac{t^3}{2^n}} \right)\ .
$$

Proof. Since the probability that a uniformly random $y_1,\dots,y_t$ is collision free is at least $1 - t^3/2^n$ when measuring a uniformly random distinct $y_1,\dots,y_t$, applying the gentle measurement measurement lemma yields the desired trace distance bound. $\square$

The final important step is to take G to be a 2-design, which foces the relation states to be close to the subspace of bijective relations. Formally, we have the following lemma, which follows closely the proof of [MH24].

Lemma 33 (Trace distance of projecting onto bijective relation states)**.**

$$
\mathrm{td}\left(\mathbb{E}_{\mathsf{C}}\left[|\mathcal{A}^{\mathsf{FHFHFO}\cdot\mathsf{C}}\rangle\langle\mathcal{A}^{\mathsf{FHFHFO}\cdot\mathsf{C}}| \right], \mathsf{Collide}\cdot\mathbb{E}_{\mathsf{C}}\left[|\mathcal{A}^{\mathsf{FHFHFO}\cdot\mathsf{C}}\rangle\langle\mathcal{A}^{\mathsf{FHFHFO}\cdot\mathsf{C}}| \right]\cdot\mathsf{Collide} \right)
$$
$$
\le O\left(\sqrt{\frac{t^3}{2^n}} \right)\ .
$$

Proof. Following the proof in [MH24], we observe that $\mathsf{Compress}^\dagger \cdot \mathsf{Compress} \cdot \mathsf{Collide} = \mathsf{Collide}$. Then applying Lemma 2, we have the following bound on the trace distance:

$$1 - \mathrm{Tr}\left(\mathsf{Collide} \cdot \mathbb{E}_C\left[\left|\mathcal{A}^{\mathsf{FHFO\cdot C}}\middle\rangle\middle\langle\mathcal{A}^{\mathsf{FHFO\cdot C}}\right|\right] \cdot \mathsf{Collide}\right)$$

$$= 1 - \mathrm{Tr}\left(\mathsf{Compress} \cdot \mathsf{Collide} \cdot \mathbb{E}_C\left[\left|\mathcal{A}^{\mathsf{FHFHFO\cdot C}}\middle\rangle\middle\langle\mathcal{A}^{\mathsf{FHFHFO\cdot C}}\right|\right] \cdot \mathsf{Collide} \cdot \mathsf{Compress}^\dagger\right)$$

$$\leq 1 - \mathrm{Tr}\left(\Lambda^{(t)} \cdot \mathbb{E}_C\left[\left|\mathcal{A}^{\mathsf{PR\cdot C}}\middle\rangle\middle\langle\mathcal{A}^{\mathsf{PR\cdot C}}\right|\right] \cdot \Lambda^{(t)}\right) + O\left(\sqrt{\frac{t^3}{2^n}}\right)$$

$$\leq 1 - \mathrm{Tr}\left(\Pi^{(t)} \cdot \mathbb{E}_C\left[\left|\mathcal{A}^{\mathsf{PR\cdot C}}\middle\rangle\middle\langle\mathcal{A}^{\mathsf{PR\cdot C}}\right|\right] \cdot \Pi^{(t)}\right) + O\left(\sqrt{\frac{t^3}{2^n}}\right)$$

$$\leq \frac{t(t-1)}{2^n} + O\left(\sqrt{\frac{t^3}{2^n}}\right).$$

Here the second line uses Lemma 31, the second to last line uses the previous lemma, which states that the projection onto collision free relations is close to the projection onto bijections, and the final line uses Corollary 4.1 from [MH24]. This completes the proof of the lemma. $\square$

Now, we put everything together and prove Theorem 17, using the fact that [MH24] has already shown that the path recording oracle itself is both right-invariant and indistinguishable from a Haar random unitary.

Proof of Theorem 17. We consider the following sequence of hybrids

1. The state of the adversary querying the FHFHFC ensemble.

2. The adversary's view of the purification of the FHFHFC ensemble.

3. The adversary's view of purification of the FHFHFC ensemble, after projecting onto collision free relation states.

4. The adversary's view of the purification of the PR $\cdot$ C oracle, after projection onto collision free relations.

5. The adversary's view of the purification of the PR $\cdot$ C oracle.

6. The adversary's view of the PR oracle.

7. The state of the adversary querying a Haar random unitary.

1 is equivlant to 2 by the purification technique for the FHF oracle (Lemma 20). 2 is close to 3 by Lemma 33. 3 is close to 4 by Lemma 31. 4 is close to 5 by Corollary 4.1 from [MH24], and 5 is close to 6 by Lemma 4.3 from [MH24]. Applying the triangle inequality between all of these hybrids completes the proof. $\square$

Remark 34. *We note that the proof in this section also applies to the case when* F *is a binary phase state. The only property we used was the orthogonality of* set *states, which also applies to the set states formed using the purification technique to binary phase states.*

Conjecture 35. *We conjecture that it can be shown that* FHFHF *itself (without the random Clifford operation) for uniformly random phases is a strong (i.e. inverse secure) pseudorandom unitary as well. As noted before, proving this likely requires understanding when $((x_i, y_i))_i$ is uniquely determined by the set of x's, y's, and $x \oplus y's$.*

6.6.2   Removing the random Clifford unitary

In Ref. [JLS18] it was conjectured that a variant of the FHF ensemble is a pseudorandom unitaries. Here, we observe that the Clifford unitary in FHFHFC can be replaced by any approximate unitary 2-design. More precisely, the above proof works unchanged for C a relative error approximate 2-design ν in the sense that

$$(1-\varepsilon) \mathop{\mathbb{E}}_{U \sim \nu} U^{\otimes t}(.)(U^\dagger)^{\otimes t} \preceq \mathop{\mathbb{E}}_{U \sim \nu} U^{\otimes t}(.)(U^\dagger)^{\otimes t} \preceq (1+\varepsilon) \mathop{\mathbb{E}}_{U \sim \nu} U^{\otimes t}(.)(U^\dagger)^{\otimes t}, \qquad (6.36)$$

where $A \preceq B$ if $B - A$ is a completely positive map. Let ν_{FHF} denote the probability distribution defined by drawing two iid random diagonal untiaries $U_{1/2} = \sum_{x \in \{0,1\}^n} e^{i\varphi_{x,1/2}} |x\rangle\langle x|$ for uniformly random angles $\varphi_{x,1/2} \in \{2\pi k/2^n\}_{k=0}^{2^n-1}$ and implementing $U_1 H^{\otimes n} U_2$.

Theorem 37 ([NHMW17]). *We have*

$$||\mathbb{E}_{U \sim \nu_{\mathsf{FHF}}^{*l}} U^{\otimes 2}(.)(U^\dagger)^{\otimes 2} - \mathbb{E}_{U \sim \mu_H} U^{\otimes 2}(.)(U^\dagger)^{\otimes 2}||_\diamond \leq 4 \times 2^{-nl}, \qquad (6.38)$$

where $||(.)||_\diamond$ denotes the diamond norm, or channel indistinguishability defined as

$$||A||_\infty := \max_{\rho, ||\rho||_1 = 1} ||(A \otimes \mathbb{1}_d)\rho||_1. \qquad (6.39)$$

This is readily combined with the following standard lemma from Ref. [BHH16]:

Lemma 40. *If ν is an additive ε-approximate unitary t-design in the sense of Definition 5 then ν is a relative $\varepsilon \times 2^{2nt}$-approximate t-design.*

It thus suffices to choose $l = 3$ to obtain a relative error design from iterations of F and H. In particular, we can replace C in FHFHFC by a unitary drawn from ν_{FHF}^{*3}. Notice that the F taken in the proof is identically distributed to D (i.e. it is a uniformly random diagonal matrix). We arrive at the following statement

Theorem 41. *The ensemble* FHFHFHFHFHF *is indistinguishable from a Haar random unitary by any agent with access to* poly(n) *adaptive queries.*

In particular, the diagonal unitaries in the repetated FH can be replaced by cryptographic assumptions such as quantum secure one-way functions or HPS as discussed in the next section.

6.6.3 Instantiation under HPS

Note that the version of HPS learning proposed in Section 3 involves distinguishing an ensemble of states from Haar random states. As of this paper, we do not know how to build pseudo-random unitaries from pseudo-random states, which seems closely related to building pseudo-random unitaries from the distinguishing variant of HPS. We instead propose that building pseudo-random unitaries by layering HPS circuits with Hadamard gates yields a pseudo-random unitary.

Conjecture 42 (HPS pseudorandom unitaries). *Let* $\mathbf{A} = \{\mathbf{A}_\ell \sim \mathbb{Z}_2^{m \times n}\}_{\ell=1}^6$, *and let* $\mathbf{k} = \{\mathbf{k}_\ell = (\theta_1, \dots, \theta_m)\}_{\ell=1}^6$ *be chosen uniformly at random (according to some discretization), with* $m \geq n + \log(n)$. *We conjecture that the following gives rise to a pseudorandom unitary:*

$$
\mathsf{U}^{\mathbf{A},\mathbf{k}} = \prod_{i=1}^6 \left(H^{\otimes n} \exp\left(i \sum_{i=1}^m (\theta_\ell)_i \bigotimes_{j=1}^n Z^{(\mathbf{A}_\ell)_{ij}} \right) \right).
$$

Of course, we note that this conjecture is an assumption about the distinguishing advantage of any computationally bounded adversary. Our evidence that this conjecture is true comes from the security of the repeated DH pseudo-random unitary that we proved in the previous section. In that section, we took "ideal" HPS instances (i.e. completely random diagonal unitary) for all the diagonal gates, which can be viewed as running an exponentially deep HPS circuit. Thus, our conjecture essentially amounts to the fact that stopping this evolution by the HPS unitary after some polynomial-amount of time yields a pseudorandom unitary instead of a truly Haar random one.

It is an interesting question to interrogate the security of this construction, as beyond the security of the idealized scheme, we have no evidence for or against the security of this pseudorandom unitary.

7 Discussion and Future Work

Here we discuss connections between the HPS assumption and other assumptions in post-quantum cryptography and fully quantum cryptography. We also provide directions for potential future work.

Applications of HPS in the real world. In the work of [BFV19], the authors consider pseudo-random states formed by starting from a maximally entangled state, and applying time evolution by a fixed operator (called H_{CFT}) that represents the black hole. In order to justify their assumption, they model the black hole as a Haar random unitary, and construct pseudo-random states by alternating random Pauli gates with applications of the Haar random unitary. The HPS assumption, while having no formal connection to the physics of black holes, has a similar flavor to the random-circuit toy models of black hole evolution. In particular, if we imagine that the CFT Hamiltonian looks like the

X-program Hamiltonian, then the HPS assumption becomes akin to distinguishing the states after time-evolution of a black hole from a known state. Furthermore, if we take the view that black hole evolution appends a uniformly random θ_i, the HPS Hamiltonian allows us to both perform black hole time evolution, and interweave Pauli X gates, yielding a similar construction to the black-box PRS from [BFV19]. We believe that the HPS assumption could serve as a toy model for black hole scrambling dynamics, and would be interested in seeing a small scale implementation on a real quantum computer.

Building more cryptography. In this work, we built a number of cryptographic primitives from the HPS assumption. These included all primitives contained in "MicroCrypt" (which are implied by one-way functions but do not imply one-way functions in a black box way), and even some primitives (like public key cryptography with quantum public keys) which are not generally considered MicroCrypt primitives. We believe that the concrete assumption we propose should allow for much more interesting cryptography, in the same way that cryptographers have extended the *Learning with Errors* [Reg09] problem to build increasingly powerful and complex tools like functional encryption, indistinguishability and obfuscation, and zero-knowledge proofs. We hope that similar extensions of our results will yield efficient implementations of more useful cryptographic primitives on near-term quantum computers as well. We point to recent work describing possible constructions of indistinguishability obfuscation (iO) from classical reversible circuits as a possible future application of the HPS assumption: Can HPS be used in a similar way to build a quantum analog of random output iO?

Establishing the security of HPS. This paper proposes a concrete problem that we consider hard to solve. While there have been attempts to solve related problems, and these problems have evaded polynomial time solutions so far, it would be foolish to take the word of a single paper on the computational hardness of any problem. In order for both the authors and the community at large to build confidence in the security of the HPS assumption, effort must go into attempting to break the assumption, or providing more evidence of its hardness. This paper (the authors believe) provides ample evidence towards the security of HPS. We provide an worst-to-average case reduction, an attack in certain parameter regimes, and an analysis of run-time of state of the art algorithms for attacking HPS. We believe that our work makes the HPS assumption worthy of future study and an interesting algorithmic problem in its own right. We hope our work prompts researchers to design attacks and attempt to break the security of our assumptions.

References

[Aar18]   Scott Aaronson. "Shadow tomography of quantum states". In: *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*. 2018, pp. 325–338 (cit. on pp. 21, 30).

[ABDY23]   Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. *Optimal algorithms for learning quantum phase states*. 2023. arXiv: 2208.07851 [quant-ph]. URL: https://arxiv.org/abs/2208.07851 (cit. on pp. 6–8, 31).

[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)*. Ed. by Venkatesan Guruswami. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 2:1–2:21. ISBN: 978-3-95977-309-6. DOI: 10.4230/LIPIcs.ITCS.2024.2. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2024.2 (cit. on pp. 3, 4, 10, 35, 37).

[AC17]   Scott Aaronson and Lijie Chen. "Complexity-Theoretic Foundations of Quantum Supremacy Experiments". In: *32nd Computational Complexity Conference (CCC 2017)*. Ed. by Ryan O'Donnell. Vol. 79. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, 22:1–22:67. ISBN: 978-3-95977-040-8. DOI: 10.4230/LIPIcs.CCC.2017.22. (Visited on 09/20/2018) (cit. on p. 12).

[ACD+22]   Gorjan Alagic, David Cooper, Quynh Dang, Thinh Dang, John M. Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl A. Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, and Daniel Apon. *Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process*. en. 2022-07-05 04:07:00 2022. DOI: https://doi.org/10.6028/NIST.IR.8413. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=934458 (cit. on p. 3).

[AGQY22]   Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. "Pseudorandom (function-like) quantum state generators: New definitions and applications". In: *Theory of Cryptography Conference*. Springer. 2022, pp. 237–265 (cit. on pp. 3, 10).

[Ale03]   Michael Alekhnovich. "More on Average Case vs Approximation Complexity". In: *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*. FOCS '03. USA: IEEE Computer Society, 2003, p. 298. ISBN: 0769520405 (cit. on p. 3).

[AQY22]   Prabhanjan Ananth, Luowen Qian, and Henry Yuen. "Cryptography from Pseudorandom Quantum States". In: *Advances in Cryptology – CRYPTO 2022*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Cham: Springer Nature Switzerland, 2022, pp. 208–236. ISBN: 978-3-031-15802-5 (cit. on p. 4).

[BBBV97]   Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. "Strengths and weaknesses of quantum computing". In: *SIAM journal on Computing* 26.5 (1997), pp. 1510–1523 (cit. on p. 12).

[BCH+21]   Fernando GSL Brandão, Wissam Chemissany, Nicholas Hunter-Jones, Richard Kueng, and John Preskill. "Models of quantum complexity growth". In: *PRX Quantum* 2.3 (2021), p. 030316 (cit. on p. 12).

[BCJ23]   Michael J. Bremner, Bin Cheng, and Zhengfeng Ji. "IQP Sampling and Verifiable Quantum Advantage: Stabilizer Scheme and Classical Security". In: arXiv:2308.07152 (Aug. 2023). arXiv: 2308.07152. (Visited on 08/15/2023) (cit. on pp. 6, 24).

[BCQ22]   Zvika Brakerski, Ran Canetti, and Luowen Qian. "On the computational hardness needed for quantum cryptography". In: *arXiv preprint arXiv:2209.04101* (2022) (cit. on pp. 3, 12).

[BEG+24]   Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. "Logical Quantum Processor Based on Reconfigurable Atom Arrays". In: *Nature* 626.7997 (Feb. 2024), pp. 58–65. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06927-3. (Visited on 02/01/2024) (cit. on p. 11).

[BEM+23]   John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. *Unitary Complexity and the Uhlmann Transformation Problem*. 2023. arXiv: 2306.13073 [quant-ph]. URL: https://arxiv.org/abs/2306.13073 (cit. on p. 3).

[Beu22]   Ward Beullens. *Breaking Rainbow Takes a Weekend on a Laptop*. Cryptology ePrint Archive, Paper 2022/214. 2022. URL: https://eprint.iacr.org/2022/214 (cit. on p. 4).

[BFG+24]   Adam Bouland, Bill Fefferman, Soumik Ghosh, Tony Metger, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. "Public-Key Pseudoentanglement and the Hardness of Learning Ground State Entanglement Structure". In: *39th Computational Complexity Conference (CCC 2024)*. Ed. by Rahul Santhanam. Vol. 300. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 21:1–21:23. ISBN: 978-3-95977-331-7. DOI: 10.4230/LIPIcs.CCC.2024.21. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CCC.2024.21 (cit. on pp. 3, 4, 10).

[BFKL94]   Avrim Blum, Merrick Furst, Michael Kearns, and Richard J. Lipton. "Cryptographic Primitives Based on Hard Learning Problems". In: *Advances in Cryptology — CRYPTO' 93*. Ed. by Douglas R. Stinson. Berlin, Heidelberg: Springer Berlin Heidelberg, 1994, pp. 278–291. ISBN: 978-3-540-48329-8 (cit. on p. 3).

[BFV19]   Adam Bouland, Bill Fefferman, and Umesh Vazirani. "Computational pseudorandomness, the wormhole growth paradox, and constraints on the AdS/CFT duality". In: *arXiv preprint arXiv:1910.14646* (2019) (cit. on pp. 12, 50, 51).

[BH08]   Alex Brodsky and Shlomo Hoory. "Simple permutations mix even better". In: *Random Structures & Algorithms* 32.3 (2008), pp. 274–289 (cit. on pp. 5, 13).

[BHH16]   Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. "Local random quantum circuits are approximate polynomial-designs". In: *Communications in Mathematical Physics* 346 (2016), pp. 397–434 (cit. on pp. 5, 13, 29, 49).

[BJ24]   Rishabh Batra and Rahul Jain. "Commitments are equivalent to one-way state generators". In: *arXiv preprint arXiv:2404.03220* (2024) (cit. on pp. 3, 12).

[BJS10]   M. J. Bremner, R. Jozsa, and D. J. Shepherd. "Classical Simulation of Commuting Quantum Computations Implies Collapse of the Polynomial Hierarchy". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 467.2126 (2010), pp. 459–472. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2010.0301. (Visited on 07/26/2016) (cit. on pp. 6, 7, 12).

[BK00]   H. Barnum and E. Knill. *Reversing quantum dynamics with near-optimal quantum and classical fidelity*. 2000. arXiv: quant-ph/0004088 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/0004088 (cit. on pp. 21, 31).

[BM24]   Zvika Brakerski and Nir Magrafta. "Real-Valued Somewhat-Pseudorandom Unitaries". In: *arXiv preprint arXiv:2403.16704* (2024) (cit. on p. 10).

[BMM+24]   Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. *Oracle Separation Between Quantum Cryptography with Classical Communication and One-wayness*. 2024 (cit. on p. 12).

[BMM24]   Amos Beimel, Tal Malkin, and Noam Mazor. *Structural Lower Bounds on Black-Box Constructions of Pseudorandom Functions*. Cryptology ePrint Archive, Paper 2024/1104. 2024. URL: https://eprint.iacr.org/2024/1104 (cit. on p. 4).

[BMS16]   Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. "Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations". In: *Physical Review Letters* 117.8 (Aug. 2016). ISSN: 1079-7114. DOI: 10.1103/physrevlett.117.080501. URL: http://dx.doi.org/10.1103/PhysRevLett.117.080501 (cit. on pp. 6, 7).

[BN24]   John Bostanci and Barak Nehoran. *Oracle Separation Between Quantum Commitments and One-wayness*. 2024 (cit. on p. 12).

[BO21]   Costin Bădescu and Ryan O'Donnell. "Improved quantum data analysis". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1398–1411 (cit. on p. 30).

[BS18]   Adam R Brown and Leonard Susskind. "Second law of quantum complexity". In: *Physical Review D* 97.8 (2018), p. 086015 (cit. on p. 12).

[BS19]   Zvika Brakerski and Omri Shmueli. *(Pseudo) Random Quantum States with Binary Phase*. 2019. arXiv: 1906.10611 [quant-ph] (cit. on pp. 3, 6, 8, 29).

[CBB⁺24]   Chi-Fang Chen, Adam Bouland, Fernando GSL Brandão, Jordan Docter, Patrick Hayden, and Michelle Xu. "Efficient unitary designs and pseudorandom unitaries from permutations". In: *arXiv preprint arXiv:2404.16751* (2024) (cit. on pp. 10, 13).

[CCMR24]   Ran Canetti, Claudio Chamon, Eduardo Mucciolo, and Andrei Ruckenstein. "Towards general-purpose program obfuscation via local mixing". In: *Cryptology ePrint Archive* (2024) (cit. on pp. 5, 13).

[CD22]   Wouter Castryck and Thomas Decru. *An efficient key recovery attack on SIDH*. Cryptology ePrint Archive, Paper 2022/975. 2022. URL: https://eprint.iacr.org/2022/975 (cit. on p. 4).

[CGGH24]   Bruno P. Cavalar, Eli Goldin, Matthew Gray, and Peter Hall. *A Meta-Complexity Characterization of Quantum Cryptography*. 2024. arXiv: 2410.04984 [cs.CR]. URL: https://arxiv.org/abs/2410.04984 (cit. on p. 14).

[CHH⁺24]   Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan. "Incompressibility and spectral gaps of random circuits". In: *arXiv preprint arXiv:2406.07478* (2024) (cit. on pp. 5, 13, 29, 61).

[Col23]   Andrea Coladangelo. *Quantum trapdoor functions from classical one-way functions*. 2023. arXiv: 2302.12821 [quant-ph]. URL: https://arxiv.org/abs/2302.12821 (cit. on pp. 9, 33–35).

[EFL⁺24]   Netta Engelhardt, Åsmund Folkestad, Adam Levine, Evita Verheijden, and Lisa Yang. *Cryptographic Censorship*. 2024. arXiv: 2402.03425 [hep-th]. URL: https://arxiv.org/abs/2402.03425 (cit. on p. 3).

[EJ00]   Yonina C. Eldar and G. David Forney Jr. *On Quantum Detection and the Square-Root Measurement*. 2000. arXiv: quant-ph/0005132 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/0005132 (cit. on p. 31).

[FGO⁺13]   Jean-Charles Faugere, Valérie Gauthier-Umana, Ayoub Otmani, Ludovic Perret, and Jean-Pierre Tillich. "A distinguisher for high-rate McEliece cryptosystems". In: *IEEE Transactions on Information Theory* 59.10 (2013), pp. 6830–6844 (cit. on p. 25).

[FGSY24]   Bill Fefferman, Soumik Ghosh, Makrand Sinha, and Henry Yuen. *The Hardness of Quantum Learning and its Cryptographic Applications*. 2024 (cit. on p. 14).

[FI24]   Xiaozhou Feng and Matteo Ippoliti. *Dynamics of Pseudoentanglement*. 2024. arXiv: 2403.09619 [quant-ph]. URL: https://arxiv.org/abs/2403.09619 (cit. on pp. 3, 4, 10).

[GGM86]   Oded Goldreich, Shafi Goldwasser, and Silvio Micali. "How to construct random functions". In: *J. ACM* 33.4 (Aug. 1986), pp. 792–807. ISSN: 0004-5411. DOI: 10.1145/6490.6503. URL: https://doi.org/10.1145/6490.6503 (cit. on p. 3).

[GH23]   David Gross and Dominik Hangleiter. "Secret Extraction Attacks against Obfuscated IQP Circuits". In: (Dec. 2023). arXiv: 2312.10156 [quant-ph]. (Visited on 01/07/2024) (cit. on pp. 6, 24).

[GHP24]   Lucas Gretta, William He, and Angelos Pelecanos. "More Efficient *k*-wise Independent Permutations from Random Reversible Circuits via log-Sobolev Inequalities". In: *arXiv preprint arXiv:2406.08499* (2024) (cit. on p. 13).

[GK15]   Shafi Goldwasser and Yael Tauman Kalai. *Cryptographic Assumptions: A Position Paper*. Cryptology ePrint Archive, Paper 2015/907. 2015. URL: https://eprint.iacr.org/2015/907 (cit. on p. 4).

[GLM+22]   András Gilyén, Seth Lloyd, Iman Marvian, Yihui Quek, and Mark M. Wilde. "Quantum Algorithm for Petz Recovery Channels and Pretty Good Measurements". In: *Physical Review Letters* 128.22 (June 2022). ISSN: 1079-7114. DOI: 10.1103/physrevlett.128.220502. URL: http://dx.doi.org/10.1103/PhysRevLett.128.220502 (cit. on p. 31).

[Gol06]   Oded Goldreich. *Foundations of Cryptography: Volume 1*. USA: Cambridge University Press, 2006. ISBN: 0521035368 (cit. on p. 3).

[Gow96]   W Timothy Gowers. "An almost m-wise independent random permutation of the cube". In: *Combinatorics, Probability and Computing* 5.2 (1996), pp. 119–130 (cit. on pp. 5, 13).

[Haf22]   Jonas Haferkamp. "Random quantum circuits are approximate unitary t-designs in depth $O(nt^{5+o(1)})$". In: *Quantum* 6 (2022), p. 795 (cit. on pp. 5, 13, 29).

[Haf23]   Jonas Haferkamp. "On the Moments of Random Quantum Circuits and Robust Quantum Complexity". In: (June 2023). arXiv: 2303.16944 [quant-ph]. (Visited on 04/10/2024) (cit. on p. 13).

[HE23]   Dominik Hangleiter and Jens Eisert. "Computational Advantage of Quantum Random Sampling". In: *Rev. Mod. Phys.* 95.3 (July 2023), p. 035001. DOI: 10.1103/RevModPhys.95.035001. arXiv: 2206.04079. (Visited on 07/20/2023) (cit. on pp. 6, 7, 12).

[HH24]   Taiga Hiroka and Min-Hsiu Hsieh. *Computational Complexity of Learning Efficiently Generatable Pure States*. 2024. arXiv: 2410.04373 [quant-ph]. URL: https://arxiv.org/abs/2410.04373 (cit. on p. 14).

[HILL99]   Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. "A Pseudorandom Generator from any One-way Function". In: *SIAM Journal on Computing* 28.4 (1999), pp. 1364–1396. DOI: 10.1137/S0097539793244708. eprint: https://doi.org/10.1137/S0097539793244708. URL: https://doi.org/10.1137/S0097539793244708 (cit. on p. 3).

[HKP20]   Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements". In: *Nature Physics* 16.10 (2020), pp. 1050–1057 (cit. on pp. 21, 30).

[HLT24]   Jeongwan Haah, Yunchao Liu, and Xinyu Tan. "Efficient approximate unitary designs from random Pauli rotations". In: *arXiv preprint arXiv:2402.05239* (2024) (cit. on pp. 13, 29).

[HM24]   Taiga Hiroka and Tomoyuki Morimae. *Quantum Cryptography from Meta-Complexity*. 2024. arXiv: 2410.01369 [quant-ph]. URL: https://arxiv.org/abs/2410.01369 (cit. on p. 14).

[HMMR05] Shlomo Hoory, Avner Magen, Steven Myers, and Charles Rackoff. "Simple permutations mix well". In: *Theoretical computer science* 348.2-3 (2005), pp. 251–261 (cit. on pp. 5, 13).

[HO24]   William He and Ryan O'Donnell. "Pseudorandom permutations from random reversible circuits". In: *arXiv preprint arXiv:2404.14648* (2024) (cit. on pp. 5, 13).

[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*. Springer. 2018, pp. 126–152 (cit. on pp. 3, 6, 10, 12, 19, 20, 39, 40, 49).

[Kah23]   Gregory D. Kahanamoku-Meyer. "Forging Quantum Data: Classically Defeating an IQP-based Quantum Test". In: *Quantum* 7 (Sept. 2023), p. 1107. ISSN: 2521-327X. DOI: 10.22331/q-2023-09-11-1107. arXiv: 1912.05547. (Visited on 11/16/2023) (cit. on pp. 6, 24).

[KKNY05] Akinori Kawachi, Takeshi Koshiba, Harumichi Nishimura, and Tomoyuki Yamakami. "Computational indistinguishability between quantum states and its cryptographic application". In: *Advances in Cryptology–EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005. Proceedings 24*. Springer. 2005, pp. 268–284 (cit. on p. 13).

[KP23]   Isaac H. Kim and John Preskill. "Complementarity and the unitarity of the black hole S-matrix". In: *Journal of High Energy Physics* 2023.2 (Feb. 2023). ISSN: 1029-8479. DOI: 10.1007/jhep02(2023)233. URL: http://dx.doi.org/10.1007/JHEP02(2023)233 (cit. on pp. 3, 4).

[KQST23]   William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. "Quantum cryptography in algorithmica". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1589–1602 (cit. on pp. 3, 12, 32).

[Kre21]   William Kretschmer. "Quantum pseudorandomness and classical complexity". In: *arXiv preprint arXiv:2103.09320* (2021) (cit. on pp. 3, 9, 12, 32).

[KT24a]   Dakshita Khurana and Kabir Tomer. "Commitments from quantum one-wayness". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024, pp. 968–978 (cit. on pp. 3, 12).

[KT24b]   Dakshita Khurana and Kabir Tomer. *Founding Quantum Cryptography on Quantum Advantage, or, Towards Cryptography from Sharp P-Hardness*. 2024. arXiv: 2409.15248 [quant-ph]. URL: https://arxiv.org/abs/2409.15248 (cit. on pp. 12, 14).

[LC97]   Hoi-Kwong Lo and Hoi Fung Chau. "Is quantum bit commitment really possible?" In: *Physical Review Letters* 78.17 (1997), p. 3410 (cit. on p. 12).

[LC99]   Hoi-Kwong Lo and Hoi Fung Chau. "Unconditional security of quantum key distribution over arbitrarily long distances". In: *science* 283.5410 (1999), pp. 2050–2056 (cit. on p. 12).

[Low10]   Richard A Low. "Pseudo-randomness and learning in quantum computation". In: *arXiv preprint arXiv:1006.5227* (2010) (cit. on p. 61).

[LPTV23]   Tianren Liu, Angelos Pelecanos, Stefano Tessaro, and Vinod Vaikuntanathan. "Layout Graphs, Random Walks and the t-Wise Independence of SPN Block Ciphers". In: *Annual International Cryptology Conference*. Springer. 2023, pp. 694–726 (cit. on p. 13).

[LR88]   Michael Luby and Charles Rackoff. "How to construct pseudorandom permutations from pseudorandom functions". In: *SIAM Journal on Computing* 17.2 (1988), pp. 373–386 (cit. on p. 5).

[LTV21]   Tianren Liu, Stefano Tessaro, and Vinod Vaikuntanathan. "The t-wise independence of substitution-permutation networks". In: *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part IV 41*. Springer. 2021, pp. 454–483 (cit. on p. 13).

[May01]   Dominic Mayers. "Unconditional security in quantum cryptography". In: *Journal of the ACM (JACM)* 48.3 (2001), pp. 351–406 (cit. on p. 12).

[McE78]   Robert J McEliece. "A public-key cryptosystem based on algebraic coding theory". In: *Coding Thv* 4244 (1978), pp. 114–116 (cit. on p. 24).

[Mer78]   Ralph C. Merkle. "Secure communications over insecure channels". In: *Commun. ACM* 21.4 (Apr. 1978), pp. 294–299. ISSN: 0001-0782. DOI: 10.1145/359460.359473. URL: https://doi.org/10.1145/359460.359473 (cit. on p. 3).

[MH24]   Fermi Ma and Robert Huang. *How to Construct Random Unitaries*. In preparation. Preliminary version available at https://fermima.com/pru.pdf. 2024. URL: https://fermima.com/pru.pdf (cit. on pp. 10, 12, 16, 39, 40, 45–48).

[Mon19]   Ashley Montanaro. *Pretty simple bounds on quantum state discrimination*. 2019. arXiv: 1908.08312 [quant-ph] (cit. on pp. 21, 31).

[MPSY24]   Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. *Simple constructions of linear-depth t-designs and pseudorandom unitaries*. 2024. arXiv: 2404.12647 [quant-ph]. URL: https://arxiv.org/abs/2404.12647 (cit. on pp. 3, 10, 12, 13, 39).

[MX24]   Tomoyuki Morimae and Keita Xagawa. *Quantum Group Actions*. 2024. arXiv: 2410.04777 [quant-ph]. URL: https://arxiv.org/abs/2410.04777 (cit. on p. 15).

[MY22]   Tomoyuki Morimae and Takashi Yamakawa. "One-wayness in quantum cryptography". In: *arXiv preprint arXiv:2210.03394* (2022) (cit. on pp. 3, 4, 12, 18–20).

[NHMW17]   Yoshifumi Nakata, Christoph Hirche, Ciara Morgan, and Andreas Winter. "Unitary 2-designs from random X-and Z-diagonal unitaries". In: *Journal of Mathematical Physics* 58.5 (2017) (cit. on p. 49).

[NKM14]   Yoshifumi Nakata, Masato Koashi, and Mio Murao. "Generating a statet-design by diagonal quantum circuits". In: *New Journal of Physics* 16.5 (May 2014), p. 053043. ISSN: 1367-2630. DOI: 10.1088/1367-2630/16/5/053043. URL: http://dx.doi.org/10.1088/1367-2630/16/5/053043 (cit. on p. 8).

[RAC⁺24]   Ben W. Reichardt, David Aasen, Rui Chao, Alex Chernoguzov, Wim van Dam, John P. Gaebler, Dan Gresh, Dominic Lucchetti, Michael Mills, Steven A. Moses, Brian Neyenhuis, Adam Paetznick, Andres Paz, Peter E. Siegfried, Marcus P. da Silva, Krysta M. Svore, Zhenghan Wang, and Matt Zanner. "Demonstration of Quantum Computation and Error Correction with a Tesseract Code". In: arXiv:2409.04628 (Sept. 2024). arXiv: 2409.04628. (Visited on 09/10/2024) (cit. on p. 11).

[RBB⁺24]   C. Ryan-Anderson, N. C. Brown, C. H. Baldwin, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, N. Hewitt, C. Holliman, C. V. Horst, J. Johansen, D. Lucchetti, T. Mengle, M. Matheny, Y. Matsuoka, K. Mayer, M. Mills, S. A. Moses, B. Neyenhuis, J. Pino, P. Siegfried, R. P. Stutz, J. Walker, and D. Hayes. "High-Fidelity and Fault-tolerant Teleportation of a Logical Qubit Using Transversal Gates and Lattice Surgery on a Trapped-ion Quantum Computer". In: (Apr. 2024). arXiv: 2404.16728. (Visited on 04/26/2024) (cit. on p. 11).

[Reg09]   Oded Regev. "On lattices, learning with errors, random linear codes, and cryptography". In: *Journal of the ACM (JACM)* 56.6 (2009), pp. 1–40 (cit. on pp. 3, 21, 27, 51).

[RSA78]   R. L. Rivest, A. Shamir, and L. Adleman. "A method for obtaining digital signatures and public-key cryptosystems". In: *Commun. ACM* 21.2 (Feb. 1978), pp. 120–126. ISSN: 0001-0782. DOI: 10.1145/359340.359342. URL: https://doi.org/10.1145/359340.359342 (cit. on p. 3).

[SB09]   Dan Shepherd and Michael J. Bremner. "Temporally Unstructured Quantum Computation". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 465.2105 (May 2009), pp. 1413–1439. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2008.0443. (Visited on 05/08/2017) (cit. on pp. 6, 7, 24).

[SF]   N. J. A. Sloane and Hieronymus Fischer. *Decimal expansion of* $\prod_{k>=1}(1 - 1/2^k)$. OEIS Entry A048651. URL: https://oeis.org/A048651 (cit. on p. 22).

[SHH24]   Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. "Random Unitaries in Extremely Low Depth". In: arXiv:2407.07754 (July 2024). arXiv: 2407.07754. (Visited on 07/11/2024) (cit. on pp. 5, 13).

[Sho94]   P.W. Shor. "Algorithms for Quantum Computation: Discrete Logarithms and Factoring". In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. Nov. 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700 (cit. on p. 4).

[Sho97]   Peter W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1484–1509. ISSN: 1095-7111. DOI: 10.1137/s0097539795293172. URL: http://dx.doi.org/10.1137/S0097539795293172 (cit. on p. 3).

[Vad12]   Salil P. Vadhan. "Pseudorandomness". In: *Foundations and Trends in Theoretical Computer Science* 7 (2012). ISSN: 1551-305X. DOI: 10.1561/0400000010. URL: http://dx.doi.org/10.1561/0400000010 (cit. on p. 3).

[Win99]   Andreas Winter. "Coding theorem and strong converse for quantum channels". In: *IEEE Transactions on Information Theory* 45.7 (1999), pp. 2481–2485. DOI: 10.1109/18.796385 (cit. on p. 16).

[ZCCC07]   Bei Zeng, Hyeyoun Chung, Andrew W. Cross, and Isaac L. Chuang. "Local Unitary versus Local Clifford Equivalence of Stabilizer and Graph States". In: *Phys. Rev. A* 75.3 (Mar. 2007), p. 032325. ISSN: 1050-2947, 1094-1622. DOI: 10.1103/PhysRevA.75.032325. arXiv: quant-ph/0611214. (Visited on 10/01/2024) (cit. on p. 11).

[Zha19]   Mark Zhandry. "How to record quantum queries, and applications to quantum indifferentiability". In: *Advances in Cryptology–CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II 39*. Springer. 2019, pp. 239–268 (cit. on p. 40).

[ZLK+23]   Haimeng Zhao, Laura Lewis, Ishaan Kannan, Yihui Quek, Hsin-Yuan Huang, and Matthias C. Caro. *Learning quantum states and unitaries of bounded gate complexity*. 2023. arXiv: 2310.19882 [quant-ph]. URL: https://arxiv.org/abs/2310.19882 (cit. on p. 3).

A Proofs for the Design Property of Hamiltonian Phase States

We will prove Theorem 16 via a bound on the spectral gap of the random walk ν on the diagonal group $D(2^n)$. Instead of bounding the diamond norm directly, we consider the essential norm (see the discussion in Ref. [CHH+24]) in the diagonal group:

$$g_D(\nu, t) := \left|\left| \mathbb{E}_{U\sim\nu} U^{\otimes t} \otimes \overline{U}^{\otimes t} - \mathbb{E}_{U\sim\mu_D} U^{\otimes t} \otimes \overline{U}^{\otimes t} \right|\right|_\infty . \tag{A.1}$$

This equals the operator norm of $\mathbb{E}_{U\sim\nu} U^{\otimes t} \otimes \overline{U}^{\otimes t}$ restricted to orthocomplement of the vectors invariant under $U^{\otimes t} \otimes \overline{U}^{\otimes t}$ for $U \in D(2^n)$. Therefore, g_D inherits the submultiplicativity of the operator norm:

$$g_D(\nu^{*m}, t) \leq g_D(\nu, t)^m. \tag{A.2}$$

Our strategy is therefore to establish a bound on $g_D(\nu, t)$ and then amplify to apply the following standard lemma:

Lemma 3. *A probability measure ν on $D(2^n)$ is a $g_D(\nu, t) \times 2^{nt}$-approximate diagonal t-design.*

Proof. $g_D(\nu, t)$ equals the induced 2-norm of $\mathcal{M}_\nu^{(t)} - \mathcal{M}_{\mu_H}^{(t)}$. The claim now follows from the inequality $||(.)||_\diamond \leq 2^d ||(.)||_\infty$ (see e.g. [Low10]), where $(.)$ acts on the space of linear operators on $\mathbb{C}^d$. In our case $\mathcal{M}_\nu^{(t)} - \mathcal{M}_{\mu_D}^{(t)}$ acts on $\mathbb{C}^{2^{nt} \times 2^{nt}} \cong \mathbb{C}^{2^{2nt}}$. □

In the following, we denote by ν the probability measure over $D(2^n)$ defined by $\exp(i\theta \bigotimes_{j=1}^n Z^{A_{1j}})$, with y drawn uniformly from $\{0,1\}$ and θ drawn uniformly from $(0, 2\pi]$. In the remainder of this section we will use the simplified notation a for the vector $a = A_{1i}$. Moreover, we denote by ν_q the measure that instead draws θ uniformly from the discrete set $\{2\pi l/q\}_{l=0}^{q-1}$. Theorem 16 will follow from the following lemma:

Lemma 4. *For all $t \geq 1$ and all $n \geq 1$ we have $g_D(\nu, t) \leq 1 - \frac{1}{2t}$. Moreover, $g_D(\nu_q, t) \geq 1 - \frac{1}{2t}$ if $q > 2t$.*

Proof. We begin by expanding the moment operators in the computational basis:

$$\mathop{\mathbb{E}}_{U\sim\nu} U^{\otimes t}\otimes\overline{U}^{\otimes t} = \mathop{\mathbb{E}}_{\theta\in(0,2\pi]}\mathop{\mathbb{E}}_{a\in\{0,1\}^n}\sum_{\mathbf{x},\mathbf{x}'\in\{0,1\}^{nt}}e^{i\theta\sum_{j=1}^t(-1)^{\langle a,x_j\rangle}-(-1)^{\langle a,x_j'\rangle}}|\mathbf{x},\mathbf{x}'\rangle\langle\mathbf{x},\mathbf{x}'|.$$

$$\mathop{\mathbb{E}}_{U\sim\mu_D} U^{\otimes t}\otimes\overline{U}^{\otimes t} = \sum_{\mathbf{x},\mathbf{x}'\in\{0,1\}^{nt}}\mathop{\mathbb{E}}_{\phi_x\in(0,2\pi]}e^{i\sum_{j=1}^t\phi_{x_j}-\phi_{x_j'}}|\mathbf{x},\mathbf{x}'\rangle\langle\mathbf{x},\mathbf{x}'|.$$

(A.5)

$\mathbb{E}_{\mathbf{x},\mathbf{x}'\in\{0,1\}^n}\mathbb{E}_{\phi_x}e^{i\sum_{j=1}^t\phi_{x_j}-\phi_{x_j'}} = 0$ if and only if $\mathbf{x}$ and $\mathbf{x}'$ are related by a permutation of the t bitstrings. If two bitstrings are related by such a permutation we denote it by $\mathbf{x}\sim\mathbf{x}'$. Therefore, we find

$$g_D(\nu,t) = \max_{\mathbf{x},\mathbf{x}';\mathbf{x}\not\sim\mathbf{x}'}\mathop{\mathbb{E}}_{a\in\{0,1\}^n}\mathop{\mathbb{E}}_{\theta\in(0,2\pi]}e^{i\theta\sum_{j=1}^t(-1)^{\langle a|x_j\rangle}-(-1)^{\langle a|x_j'\rangle}}.$$

(A.6)

Notice that $\mathbb{E}_{\theta\in(0,2\pi]}e^{i\theta\sum_{j=1}^t(-1)^{\langle a,x_j\rangle}-(-1)^{\langle a,x_j'\rangle}} = 0$ if and only if $\sum_{j=1}^t(-1)^{\langle a,x_j\rangle}-(-1)^{\langle a,x_j'\rangle}\neq 0$. Therefore

$$g_D(\nu,t) = \max_{\mathbf{x},\mathbf{x}';\mathbf{x}\not\sim\mathbf{x}'}\mathrm{Pr}_a\left[\sum_{j=1}^t(-1)^{\langle a,x_j\rangle}-(-1)^{\langle a,x_j'\rangle}=0\right].$$

(A.7)

To bound this probability we use Fourier analysis of Boolean functions. More precisely, we define the function $f^{\mathbf{x},\mathbf{x}'}:\{0,1\}\to\mathbb{R}$ as

$$f^{\mathbf{x},\mathbf{x}'}(a):=\sum_{j=1}^t(-1)^{\langle a,x_j\rangle}-(-1)^{\langle a,x_j'\rangle}.$$

(A.8)

Recall that the Fourier transform of a Boolean function f is defined as

$$\hat{f}(a) = \mathop{\mathbb{E}}_{x\in\{0,1\}}f(x)(-1)^{\langle x|a\rangle}.$$

(A.9)

We can easily verify that

$$\hat{f}^{\mathbf{x},\mathbf{x}'} = \sum_{j=1}^t\delta_{x_j}-\delta_{x_j'},$$

(A.10)

where $\delta_x(z) = \begin{cases}0 & \text{if } x\neq z \\ 1 & \text{if } x=z\end{cases}$.

If a string appears in both $\mathbf{x}$ and $\mathbf{x}'$ we can remove it without changing $f^{\mathbf{x},\mathbf{x}'}$. Therefore, we can remove strings until we obtain two tuples of length $0 < r \leq t$ that have no

common strings. Now, via Paseval's identity, we find the following bound:

$$(2r)^2 \left|\left\{ a, \sum_{j=1}^r (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} \neq 0 \right\}\right| \geq \sum_a \left(\sum_{j=1}^r (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} \right)^2$$

$$= 2^n \mathbb{E}_a \left(f^{\mathbf{x}, \mathbf{x}'}(a) \right)^2 \tag{A.11}$$

$$= 2^n \sum_z \left(\hat{f}^{\mathbf{x}, \mathbf{x}'}(z) \right)^2$$

$$\geq 2^n 2r.$$

Therefore, we find

$$g_D(\nu, t) = \max_{\mathbf{x} \not\sim \mathbf{x}'} 1 - 2^{-n} \left|\left\{ a \in \{0, 1\}, \sum_{j=1}^t (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} \neq 0 \right\}\right| \leq 1 - \frac{1}{2r} \leq 1 - \frac{1}{2t},$$

$$\tag{A.12}$$

which completes the proof.

For the bound on $g_D(\nu_q, t)$ we observe just as above that we have

$$g_D(\nu, t) = \max_{\mathbf{x}, \mathbf{x}'; \mathbf{x} \not\sim \mathbf{x}'} \Pr_a \left[\sum_{j=1}^t (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} = 0 \mod q \right]. \tag{A.13}$$

But as $\left| \sum_{j=1}^t (-1)^{\langle a|x_j \rangle} - (-1)^{\langle a|x_j' \rangle} \right| \leq 2t$, and $q > 2t$ we have $\sum_{j=1}^t (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} = 0 \mod q$ if and only if $\sum_{j=1}^t (-1)^{\langle a, x_j \rangle} - (-1)^{\langle a, x_j' \rangle} = 0$. The rest of the argument thus goes through without change. $\square$

Proof of Corollary 18. We use the notation $|\psi\rangle^{\otimes t, t} := |\psi\rangle^{\otimes t} \otimes \overline{|\psi\rangle}^{\otimes t}$. For any state $|\psi\rangle$:

$$\mathbb{E}_{U \sim \mu_D} |\langle \psi | U | +^n \rangle|^{2t} = \langle \psi |^{\otimes t, t} \mathbb{E}_{U \sim \mu_D} U^{\otimes t} \otimes \overline{U}^{\otimes t} |+^n\rangle^{\otimes 2t}$$

$$\leq \sqrt{\langle +^n |^{\otimes t, t} (\mathbb{E}_{U \sim \mu_D} U^{\otimes t} \otimes \overline{U}^{\otimes t})^\dagger \mathbb{E}_{U \sim \mu_D} U^{\otimes t} \otimes \overline{U}^{\otimes t} |+\rangle^{\otimes t, t}} \tag{A.14}$$

$$\leq \sqrt{\mathbb{E}_{U \sim \mu_D} |\langle +^n | U | +^n \rangle|^2}.$$

We compute

$$\mathbb{E}_{U \sim \mu_D} |\langle +^n | U | +^n \rangle|^2 = \frac{1}{2^{2nt}} \sum_{x_1, \dots, x_t, x_1', \dots, x_t'} \mathbb{E}_{\phi_x, x \in \{0,1\}} e^{i \sum_{j=1}^t \phi_{x_j} - \phi_{x_j'}} \tag{A.15}$$

As long as $\sum_{j=1}^t \phi_{x_j} - \phi_{x_j'} = 0$, for all ϕ, we have

$$\mathbb{E}_{\phi_j} e^{i \sum_{j=1}^t \phi_{x_j} - \phi_{x_j'}} = 0.$$

But $\sum_{j=1}^{t} \phi_{x_j} - \phi_{x_j'} = 0$ if and only if $(x_1, \ldots, x_t)$ and $(x_1', \ldots x_t')$ are related by a permutation in S_t. Clearly, for every fixed tuple $(x_1, \ldots, x_t)$ there are at most $t!$ permutations. Consequently, we find

$$\mathbb{E}_{U \sim \mu_D} |\langle +^n | U | +^n \rangle|^2 \leq \frac{t!}{2^{nt}} \tag{A.16}$$

and therefore, plugging into Eq. (A.14),

$$\mathbb{E}_{U \sim \mu_D} |\langle \psi | U | +^n \rangle|^{2t} \leq \sqrt{\frac{t!}{2^{nt}}}. \tag{A.17}$$

By Theorem 16 we have that ν^{*m} is a ε-approximate diagonal t-design with $\varepsilon \leq \sqrt{t!/2^{nt}}$ and thus

$$\mathbb{E}_{U \sim \nu^{*d}} |\langle \psi | U | +^n \rangle|^{2t} \leq \mathbb{E}_{U \sim \mu_D} |\langle \psi | U | +^n \rangle|^{2t} + \varepsilon \leq 2\sqrt{\frac{t!}{2^{nt}}}. \tag{A.18}$$

We can split the expectation value in Eq. (A.14) into the discrete part and the angles. Applying Markov's inequality over the distribution of A we find

$$\Pr_A \left[\mathbb{E}_\theta |\langle \psi | \prod_i e^{\sum_{i=1}^m i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^{2t} \geq (2t!/2^{nt})^{\frac{1}{4}}) \right] \leq (2(t!/2^{nt})^{\frac{1}{4}}). \tag{A.19}$$

Therefore, with high probability over the Z strings we have

$$\mathbb{E}_\theta |\langle \psi | e^{\sum_{i=1}^m i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^{2t} \geq (2t!/2^{nt})^{\frac{1}{4}}). \tag{A.20}$$

We can now use Markov's inequality again to prove that most states for any state $|\psi\rangle$ almost all states $U|0^n\rangle$ have small overlap. Now consider any A such that Eq. (A.20) holds. For any $\delta > 0$, we have

$$\Pr_\theta [|\langle \psi | \prod_i e^{i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^2 \geq (1 - \delta)] = \Pr_\theta [|\langle \psi | e^{\sum_{i=1}^m i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^{2t} \geq (1 - \delta)^t]$$

$$\leq \frac{\mathbb{E}_\theta |\langle \psi | e^{\sum_{i=1}^m i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^{2t}}{(1 - \delta)^t}$$

$$\leq 2^{-\frac{1}{4}[n - \log_2(t) + \log_2(1/(1-\delta)]t}. \tag{A.21}$$

In particular, even if we choose $1 - \delta = 2^{-n/8}$ and $t = \mathsf{poly}(n)$, we find that

$$\langle \psi | e^{\sum_{i=1}^m i\theta_i \otimes_{j=1}^n Z^{A_{ij}}} | +^n \rangle|^2 \leq 2^{-n/8} \tag{A.22}$$

with probability $1 - 2^{-\frac{1}{10} nt}$. From a union bound, we then find that there are at least 2^{cnt} distinct instances that are (almost) mutually orthogonal to each other. $\qquad \square$