Efficient Quantum Pseudorandomness from Hamiltonian Phase States

John Bostanci
Columbia University, New York, NY, USA

Jonas Haferkamp
Harvard University, Cambridge, MA, USA

Dominik Hangleiter
QuICS, University of Maryland & NIST, College Park, MD, USA
Simons Institute for the Theory of Computing, University of California at Berkeley, CA, USA

Alexander Poremba
Massachusetts Institute of Technology, Cambridge, MA, USA

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.

2012 ACM Subject Classification Theory of computation → Cryptographic primitives

Keywords and phrases Quantum pseudorandomness, quantum phase states, quantum cryptography

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.9

Related Version *Full Version*: https://arxiv.org/abs/2410.08073

Funding *John Bostanci*: Supported under AFORS (award FA9550-21-1-036) and NSF CAREER (award CCF2144219).
Jonas Haferkamp: Supported by the Harvard Quantum Initiative postdoctoral fellowship.
Dominik Hangleiter: Supported by the QuICS Hartree fellowship.
Alexander Poremba: Supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under contract number DE-SC0012704.

Acknowledgements This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by NSF QLCI Grant No. 2016245.

1 Introduction

Pseudorandomness [34, 68] 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 [44] 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* [35], 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 [61] and discrete logarithms [54], decoding error correcting codes [13, 4], or finding short vectors in high-dimensional lattices [59]. 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 [65], which has led to the design of *post-quantum assumptions* [3].

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 [46, 50, 2]. In particular, Ji, Liu and Song [46] proposed the notion of pseudorandom states [46] and pseudorandom unitaries as the natural quantum analogues of pseudorandom generators [44] and pseudorandom functions [35], respectively. The work of Kretschmer [50, 51] 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 [2, 16, 32], quantum learning theory [70], to models of scrambling phenomena in chaotic quantum systems [49, 31], and, more generally, even in the foundations of quantum cryptography [46, 50, 51, 57, 5, 17, 15, 48, 10].

Limitations of existing constructions. Despite strong evidence that MicroCrypt primitives such as pseudorandom quantum states and pseudorandom unitaries lie "below" one-way functions [50, 51], all known constructions implicitly make use of one-way functions (or other assumptions which are themselves synonymous with the existence of one-way functions) [46, 19, 55]. This naturally 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 [6, 57].

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 [11], 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 [2, 16], or when studying the emergence of thermal equilibria in isolated many-body systems [32], or when modeling scrambling phenomena in chaotic quantum systems [49]. 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 [66, 26, 12]. What makes a new cryptographic assumption reasonable? While no widely agreed upon standards exist [36], 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. [63] implies that a simple ensemble of random quantum circuits in a 1D architecture of depth polylog(n) is also pseudorandom.

Conjecture 1 can be seen as a direct quantum analogue of a claim that was first proposed by Gowers [37] 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 [42] that the Luby-Rackoff [52] 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 [25].

Gowers originally conjectured the emergence of pseudorandomness when attempting to prove that random quantum circuits converge quickly to ensembles of t-wise independent permutations [37] (this bound was further improved later on towards an optimal scaling [43, 24, 28]. In fact, this property can itself be viewed as evidence for pseudorandomness. It turns out that random quantum circuits satisfy an analogous property by converging nearly optimally towards approximate t-designs [28, 20, 40, 63].

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. 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 [64]. 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 [47, 22, 38]. At the same time, these circuits are also believed to be sufficiently mixing and hard to simulate classically [64, 21, 23, 41]. Moreover, since Hamiltonian phase states can be generated by a commuting Hamiltonian, they admit an 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 [7].

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)

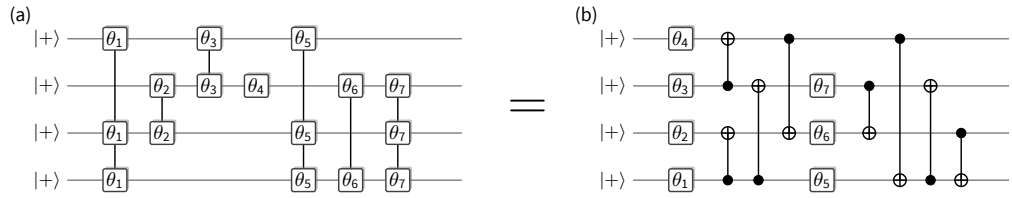

Figure 1 Hamiltonian Phase States (HPS) are generated by sequentially applying Ising-type rotations around angles θ_i to the state $|+^n\rangle = H^{\otimes n}|0^n\rangle$. (a) Example of a HPS on 4 qubits. Connected boxes at sites i, j, k with angle θ represent the unitary $\exp(i\theta Z_i Z_j Z_k)$. (b) HPS can be implemented using only single-qubit Z rotations interlaced with CNOT circuits.

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 [46]. These are generated using a single-bit output quantum-secure pseudorandom function family $\{f_k\}_k$ with

$$|\phi_k\rangle \propto \sum_x \omega_q^{f_k(x)} |x\rangle\,, \tag{1}$$

where ω_q is a q-th root of unity, for example $q = 2$ [19]. 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. [46, 7] 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.

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 $q = \mathsf{poly}(n)$ parts. We consider phase states of the form

$$|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle = \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i \bigotimes_{j=1}^{n} \mathsf{Z}^{\mathbf{A}_{ij}}\right) H^{\otimes n}|0^n\rangle\,. \tag{2}$$

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

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

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 [64]. 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 [64, 21, 23, 41].

Our cryptographic assumption rests on the apparent hardness of *learning* Hamiltonian Phase States (or generalized X programs), which was highlighted in recent work [7]. 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).
- 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).

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 [19] 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 [58] 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 [7], and has been found to have exponential time complexity in the worst case (despite only having polynomial sample complexity).
- (Notion of a random instance) A random Hamiltonian phase state $|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\rangle$, e.g., as in Equation (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. We separately show how to re-randomize the architecture and the set of angles. Thus, the hardness of our problem reduces to a worst-case version.
- (Relation to other problems) 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.

Finally, we also provide evidence that $\mathsf{HPS}_{n,m,q,\chi}$ is plausibly fully-quantum and does not allow one to construct one-way functions. In particular, we note that the result of [50] 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 and pseudorandom state generators, 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 [29] 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 [29]. Using a construction based on binary-phase states, Coladangelo [29] 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 the full version, 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 [2, 16] has found many applications in quantum physics, for example to study the emergence of thermal equilibria in isolated many-body systems [32]. 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 the full version, 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 believe 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 [5, 18, 55, 27], eventually resulting in the most recent breakthrough result by Ma and Huang [53].

The result of [53] show that the ensemble of unitaries, colloquially known as the PFC-ensemble [55], 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. In the full version, we provide a plausible construction of efficient pseudorandom unitaries from an natural assumption which is directly related to our HPS assumption: alternating applications of HPS unitaries and Hadamards.

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,

$$\mathsf{CNOT}_{k,l} e^{i\theta Z_l} = e^{i\theta Z_k Z_l} \mathsf{CNOT}_{k,l}, \tag{3}$$

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

$$\mathsf{CNOT}_{k,l} \left|+^n\right\rangle = \left|+^n\right\rangle. \tag{4}$$

More generally, if C is a circuit comprised of CNOT gates, then $C\left|x\right\rangle = \left|\mathbf{C}\cdot x\right\rangle$, 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 $\left|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\right\rangle$ can then be prepared as

$$\left|\Phi_{\boldsymbol{\theta}}^{\mathbf{A}}\right\rangle = \left(C_{\lceil m/n \rceil} \prod_{i=(\lceil m/n \rceil - 1)n}^{m} e^{i\theta_i Z_i}\right) \cdots \left(C_1 \prod_{i=1}^{n} e^{i\theta_i Z_i}\right) \left|+^n\right\rangle \tag{5}$$

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 Figure 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

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

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 [69]. 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 [14], or trapped ion processors [62, 60] in which arbitrary inter-block connectivities can be achieved.

3 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.

3.1 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 [57].

We now give a formal definition.

▶ **Definition 2** (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 $\mathcal{A}$,*

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

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 $|\Phi^{\mathbf{A}}_{\boldsymbol{\theta}}\rangle$ *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 $|\Phi^{\mathbf{A}'}_{\boldsymbol{\theta}'}\rangle$ which has non-vanishing fidelity with the input state, on average over the choice of architecture and set of angles.

3.2 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 state. Therefore, our (decision) assumption says that an ensemble of Hamiltonian Phase states forms a pseudorandom state generator [46, 57].

▶ **Definition 3** (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* $\mathcal{D}$*,*

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

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* $1/\mathsf{poly}(n, m, \log q)$*.*

4 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 4.1, we show two worst-to-average-case reductions for the $\mathsf{HPS}_{n,m,q,\chi}$ 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 $\mathbf{A}$.

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 4.2. 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.

In the full version, 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 [9, 56] and a simple algorithm that uses classical shadows [1, 45]. Moreover, we also discuss why the HPS assumption is fully quantum. To this end, we give evidence against the possibility of building one-way functions from HPS.

4.1 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 [59], with different levels of public knowledge.

Reduction for the architecture for $m \leq n$

First, we observe that for any fixed choice of angles $\boldsymbol{\theta}$, the Hamiltonian architecture can be re-randomized if $m \leq n$ and χ is the uniform distribution over full-rank matrices $\mathcal{R}(m,n) \coloneqq \{\mathbf{A} \in \mathbb{Z}_2^{m\times n} \,|\, \mathrm{rank}(\mathbf{A}) = \min(m,n)\}$. 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 probability[2] $\prod_{k=1}^{\min(m,n)}(1-2^{-k}) \geq 0.288$ [67]. 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 4** (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 $T \cdot \mathsf{poly}(n)$ 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 $t = \mathsf{poly}(n)$. 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 $T \cdot \mathsf{poly}(n)$.

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

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

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

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

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

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

$$\begin{aligned}
U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger} &= \sum_{x\in\{0,1\}^n}\langle x|\,U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger}|x\rangle\cdot|x\rangle\langle x| \\
&= \sum_{x\in\{0,1\}^n}\langle \mathbf{R}x|\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)|\mathbf{R}x\rangle\cdot|x\rangle\langle x| \\
&= \sum_{x\in\{0,1\}^n}(-1)^{\sum_{j=1}^{n}\mathbf{C}_{ij}(\mathbf{R}x)_j}|x\rangle\langle x| \\
&= \sum_{x\in\{0,1\}^n}(-1)^{\sum_{j=1}^{n}(\mathbf{C}\cdot\mathbf{R})_{ij}x_j}|x\rangle\langle x| \\
&= \sum_{x\in\{0,1\}^n}\langle x|\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}\right)|x\rangle\cdot|x\rangle\langle x| \\
&= \bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}.
\end{aligned}$$

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

$$\begin{aligned}
U_{\mathbf{R}}\exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger} &= \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\,U_{\mathbf{R}}\left(\bigotimes_{j=1}^{n}\mathsf{Z}^{\mathbf{C}_{ij}}\right)U_{\mathbf{R}}^{\dagger}\right) \\
&= \exp\left(\mathrm{i}\sum_{i=1}^{m}\theta_i\bigotimes_{j=1}^{n}\mathsf{Z}^{(\mathbf{C}\cdot\mathbf{R})_{ij}}\right).
\end{aligned}$$

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 $U_{\mathbf{R}}|\psi_{\boldsymbol{\theta}}^{\mathbf{C}}\rangle)^{\otimes t}$ is distributed precisely as in the $\mathsf{HPS}_{n,n,q,\chi}$ 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. ◀

4.2 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 $\boldsymbol{A}_1 \in \{0,1\}^n$ and a uniformly random angle $\theta \in [0, 2\pi)$ and apply $e^{\mathrm{i}\theta \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{1j}}}$. 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 $e^{\mathrm{i}\sum_{i=1}^m \theta_i \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{ij}}}$ is an approximate t-design on the diagonal group. More precisely, we prove the following theorem:

▶ **Theorem 5.** *For $m \geq 2t(2nt + \log(1/\varepsilon))$ the random unitary $e^{\mathrm{i}\theta_i \sum_{i=1}^m \bigotimes_j Z^{\boldsymbol{A}_{ij}}}$ with random A_{ij} and θ_i is a ε-approximate diagonal t-design. 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 5 in the full version. The proof of Theorem 5 is remarkably simple in comparison to the derivations of similar results for random quantum circuits [28, 20, 40]. Additionally, the constants in Theorem 5 are unusually small: In stark contrast the constants in these results are north of 10^{13}. A similar result was obtained in Ref. [39] for the related random Pauli rotations $e^{\mathrm{i}\theta P}$ for a random $\theta \in (0, 2\pi]$ and a random Pauli string P.

Theorem 5 almost directly implies the following corollary:

▶ **Corollary 6.** *For $m \geq 2t(2nt + \log(1/\varepsilon))$ the state ensemble defined by $|\Phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle = U|+^n\rangle$ 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 $|\phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle$ 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.

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

▶ **Corollary 7.** *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 $\boldsymbol{A}$ that*

$$\Pr_{\boldsymbol{\theta}}\left[|\langle\psi|\exp\left(\sum_{i=1}^m \mathrm{i}\theta_i \bigotimes_{j=1}^n Z^{\boldsymbol{A}_{ij}}\right)|+^n\rangle|^2 \geq 1-\delta\right] \leq 2^{-\Omega(nt)}. \tag{6}$$

We defer the proofs of Theorem 5 and Corollary 7 to the full version of the paper.

4.3 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

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

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 2, 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 [8, Theorem 1.5] requires $n\log^2(M)$ copies improving over the bound from Ref. [1]. 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. [45]. Given a secret state $|\Phi_{\boldsymbol{\theta}}^{\boldsymbol{A}}\rangle$ 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. [7, Theorem 15] that for any state of the form

$$\exp\left(\mathrm{i} \sum_{y\in\{0,1\}^n} a_y \bigotimes_{j=1}^{n} Z^{y_j} \right) |+^n\rangle \tag{7}$$

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.

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

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

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

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) [9, 56]. The ensemble $\mathcal{E}$ now turns out to be *geometrically uniform* because it can be written as $\mathcal{E}_{\mathbf{A}} = \left\{U_{\boldsymbol{\theta}}^{\mathbf{A}} |+^n\rangle\right\}_{\boldsymbol{\theta}=(\theta_1,\ldots,\theta_m)}$ where $\{U_{\boldsymbol{\theta}}^{\mathbf{A}}\}_{\boldsymbol{\theta}}$ is an Abelian group of matrices. Eldar and Forney [30] 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 [33]. Consequently, we believe that the HPS problem remains computationally intractable, even if the architecture is public.

References

1 Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pages 325–338, 2018.

2 Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum Pseudoentanglement. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:21, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2024.2.

3 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, 2022-07-05 04:07:00 2022. doi:10.6028/NIST.IR.8413.

4 Michael Alekhnovich. More on average case vs approximation complexity. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, page 298, USA, 2003. IEEE Computer Society.

5 Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In *Theory of Cryptography Conference*, pages 237–265. Springer, 2022.

6 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 208–236, Cham, 2022. Springer Nature Switzerland.

7 Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. Optimal algorithms for learning quantum phase states, 2023. arXiv:2208.07851.

8 Costin Bădescu and Ryan O'Donnell. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1398–1411, 2021.

9 H. Barnum and E. Knill. Reversing quantum dynamics with near-optimal quantum and classical fidelity, 2000. arXiv:quant-ph/0004088.

10 Rishabh Batra and Rahul Jain. Commitments are equivalent to one-way state generators. *arXiv preprint arXiv:2404.03220*, 2024.

11 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.

12 Ward Beullens. Breaking rainbow takes a weekend on a laptop. Cryptology ePrint Archive, Paper 2022/214, 2022. URL: https://eprint.iacr.org/2022/214.

13 Avrim Blum, Merrick Furst, Michael Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology — CRYPTO' 93*, pages 278–291, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.

14 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. *Nature*, 626(7997):58–65, February 2024. doi:10.1038/s41586-023-06927-3.

15 John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the uhlmann transformation problem, 2023. arXiv:2306.13073.

16 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 Rahul Santhanam, editor, *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:23, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2024.21.

17 Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. *arXiv preprint arXiv:2209.04101*, 2022.

18 Zvika Brakerski and Nir Magrafta. Real-valued somewhat-pseudorandom unitaries. *arXiv preprint arXiv:2403.16704*, 2024.

19 Zvika Brakerski and Omri Shmueli. (pseudo) random quantum states with binary phase, 2019. arXiv:1906.10611.

20 Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.

21 M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2010. doi:10.1098/rspa.2010.0301.

22 Michael J. Bremner, Bin Cheng, and Zhengfeng Ji. IQP Sampling and Verifiable Quantum Advantage: Stabilizer Scheme and Classical Security, August 2023. arXiv:2308.07152.

23 Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical Review Letters*, 117(8), August 2016. doi:10.1103/physrevlett.117.080501.

24 Alex Brodsky and Shlomo Hoory. Simple permutations mix even better. *Random Structures & Algorithms*, 32(3):274–289, 2008.

25 Ran Canetti, Claudio Chamon, Eduardo Mucciolo, and Andrei Ruckenstein. Towards general-purpose program obfuscation via local mixing. *Cryptology ePrint Archive*, 2024.

26 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.

27 Chi-Fang Chen, Adam Bouland, Fernando GSL Brandão, Jordan Docter, Patrick Hayden, and Michelle Xu. Efficient unitary designs and pseudorandom unitaries from permutations. *arXiv preprint arXiv:2404.16751*, 2024.

28 Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan. Incompressibility and spectral gaps of random circuits. *arXiv preprint arXiv:2406.07478*, 2024.

29 Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions, 2023. arXiv:2302.12821.

30 Yonina C. Eldar and G. David Forney Jr. On quantum detection and the square-root measurement, 2000. arXiv:quant-ph/0005132.

31 Netta Engelhardt, Åsmund Folkestad, Adam Levine, Evita Verheijden, and Lisa Yang. Cryptographic censorship, 2024. arXiv:2402.03425.

32 Xiaozhou Feng and Matteo Ippoliti. Dynamics of pseudoentanglement, 2024. arXiv:2403.09619.

33 András Gilyén, Seth Lloyd, Iman Marvian, Yihui Quek, and Mark M. Wilde. Quantum algorithm for petz recovery channels and pretty good measurements. *Physical Review Letters*, 128(22), June 2022. doi:10.1103/physrevlett.128.220502.

34 Oded Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, USA, 2006.

35 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, August 1986. doi:10.1145/6490.6503.

36 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.

37 W Timothy Gowers. An almost m-wise independent random permutation of the cube. *Combinatorics, Probability and Computing*, 5(2):119–130, 1996.

38 David Gross and Dominik Hangleiter. Secret extraction attacks against obfuscated IQP circuits, December 2023. arXiv:2312.10156.

39 Jeongwan Haah, Yunchao Liu, and Xinyu Tan. Efficient approximate unitary designs from random pauli rotations. *arXiv preprint arXiv:2402.05239*, 2024.

40 Jonas Haferkamp. Random quantum circuits are approximate unitary t-designs in depth $o(nt^{5+o(1)})$. *Quantum*, 6:795, 2022.

41 Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95(3):035001, July 2023. doi:10.1103/RevModPhys.95.035001.

42 William He and Ryan O'Donnell. Pseudorandom permutations from random reversible circuits. *arXiv preprint arXiv:2404.14648*, 2024.

43 Shlomo Hoory, Avner Magen, Steven Myers, and Charles Rackoff. Simple permutations mix well. *Theoretical computer science*, 348(2-3):251–261, 2005.

44 Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.

45 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

46 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018.

47 Gregory D. Kahanamoku-Meyer. Forging quantum data: Classically defeating an IQP-based quantum test. *Quantum*, 7:1107, September 2023. doi:10.22331/q-2023-09-11-1107.

48 Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 968–978, 2024.

49 Isaac H. Kim and John Preskill. Complementarity and the unitarity of the black hole s-matrix. *Journal of High Energy Physics*, 2023(2), February 2023. doi:10.1007/jhep02(2023)233.

50 William Kretschmer. Quantum pseudorandomness and classical complexity. *arXiv preprint arXiv:2103.09320*, 2021.

51 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*, pages 1589–1602, 2023.

52 Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

53 Fermi Ma and Robert Huang. How to construct random unitaries, 2024. In preparation. Preliminary version available at https://fermima.com/pru.pdf. URL: https://fermima.com/pru.pdf.

54 Ralph C. Merkle. Secure communications over insecure channels. *Commun. ACM*, 21(4):294–299, April 1978. doi:10.1145/359460.359473.

55 Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries, 2024. arXiv:2404.12647.

56 Ashley Montanaro. Pretty simple bounds on quantum state discrimination, 2019. arXiv:1908.08312.

57 Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint arXiv:2210.03394*, 2022.

58 Yoshifumi Nakata, Masato Koashi, and Mio Murao. Generating a statet-design by diagonal quantum circuits. *New Journal of Physics*, 16(5):053043, May 2014. doi:10.1088/1367-2630/16/5/053043.

59 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.

60 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, September 2024. arXiv:2409.04628.

61 R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, February 1978. doi:10.1145/359340.359342.

62 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, April 2024. arXiv:2404.16728.

63 Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth, July 2024. arXiv:2407.07754.

64 Dan Shepherd and Michael J. Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 465(2105):1413–1439, May 2009. doi:10.1098/rspa.2008.0443.

65 Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. doi:10.1137/s0097539795293172.

66 P.W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, November 1994. doi:10.1109/SFCS.1994.365700.

67 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.

68 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7, 2012. doi:10.1561/0400000010.

69 Bei Zeng, Hyeyoun Chung, Andrew W. Cross, and Isaac L. Chuang. Local unitary versus local Clifford equivalence of stabilizer and graph states. *Phys. Rev. A*, 75(3):032325, March 2007. doi:10.1103/PhysRevA.75.032325.

70 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.