Efficient simulation of random states and random unitaries

Gorjan Alagic[1], Christian Majenz[2], and Alexander Russell[3]

[1] QuICS, University of Maryland, and NIST, Gaithersburg, Maryland
[2] QuSoft, Amsterdam and Centrum Wiskunde & Informatica, Amsterdam
[3] Department of Computer Science and Engineering, University of Connecticut

Abstract. We consider the problem of efficiently simulating random quantum states and random unitary operators, in a manner which is convincing to unbounded adversaries with black-box oracle access.

This problem has previously only been considered for restricted adversaries. Against adversaries with an a priori bound on the number of queries, it is well-known that t-designs suffice. Against polynomial-time adversaries, one can use pseudorandom states (PRS) and pseudorandom unitaries (PRU), as defined in a recent work of Ji, Liu, and Song; unfortunately, no provably secure construction is known for PRUs.

In our setting, we are concerned with unbounded adversaries. Nonetheless, we are able to give stateful quantum algorithms which simulate the ideal object in both settings of interest. In the case of Haar-random states, our simulator is polynomial-time, has negligible error, and can also simulate verification and reflection through the simulated state. This yields an immediate application to quantum money: a money scheme which is information-theoretically unforgeable and untraceable. In the case of Haar-random unitaries, our simulator takes polynomial space, but simulates both forward and inverse access with zero error.

These results can be seen as the first significant steps in developing a theory of lazy sampling for random quantum objects.

1 Introduction

1.1 Motivation

Efficient simulation of randomness is a task with countless applications, ranging from cryptography to derandomization. In the setting of classical probabilistic computation, such simulation is straightforward in many settings. For example, a random function which will only be queried an a priori bounded number of times t can be perfectly simulated using a t-wise independent function [30]. In the case of unbounded queries, one can use pseudorandom functions (PRFs), provided the queries are made by a polynomial-time algorithm [16]. These are examples of *stateless* simulation methods, in the sense that the internal memory of the simulator is initialized once (e.g., with the PRF key) and then remains fixed regardless of how the simulator is queried. Against arbitrary adversaries, one must typically pass to *stateful* simulation. For example, the straightforward and

well-known technique of *lazy sampling* suffices to perfectly simulate a random function against arbitrary adversaries; however, the simulator must maintain a list of responses to all previous queries.

Each of these techniques for simulating random classical primitives has a plethora of applications in theoretical cryptography, both as a proof tool and for cryptographic constructions. These range from constructing secure cryptosystems for encryption and authentication, to proving security reductions in a wide range of settings, to establishing security in idealized models such as the Random Oracle Model [7].

Quantum randomness. As is well-known, quantum sources of randomness exhibit dramatically different properties from their classical counterparts [23,8]. Compare, for example, uniformly random n-bit classical states (i.e., n-bit strings) and uniformly random n-qubit (pure) quantum states. A random string x is obviously trivial to sample perfectly given probabilistic classical (or quantum) computation, and can be copied and distributed arbitrarily. However, it is also (just as obviously) deterministic to all parties who have examined it before. By contrast, a random state $|\varphi\rangle$ would take an unbounded amount of information to describe perfectly. Even if one manages to procure such a state, it is then impossible to copy due to the no-cloning theorem. On the other hand, parties who have examined $|\varphi\rangle$ many times before, can *still* extract almost exactly n bits of randomness from any fresh copy of $|\varphi\rangle$ they receive – even if they use the *exact same measurement procedure* each time.

The differences between random classical and random quantum maps are even more stark. The outputs of a classical random function are of course classical random strings, with all of the aforementioned properties. Outputs which have already been examined become effectively deterministic, while the rest remain uniformly random and independent. This is precisely what makes efficient simulation possible via lazy sampling. A Haar-random unitary U queried on two inputs $|\psi\rangle$ and $|\phi\rangle$ also produces (almost) independent and uniformly random states when queried, but only if the queries are *orthogonal*, i.e., $\langle\psi \mid \phi\rangle = 0$. Unitarity implies that overlapping queries must be answered consistently, i.e., if $\langle\psi \mid \phi\rangle = \delta$ then $\langle(U\psi) \mid (U\phi)\rangle = \delta$. This possibility of querying with a distinct pure state which is not linearly independent from previous queries simply doesn't exist for classical functions.

We emphasize that the above differences should not be interpreted as quantum random objects simply being "stronger" than their classical counterparts. In the case of classical states, i.e. strings, the ability to copy is quite useful, e.g., in setting down basic security definitions [9,3,2] or when rewinding an algorithm [28,29,14]. In the case of maps, determinism is also quite useful, e.g., for verification in message authentication.

1.2 The problem: efficient simulation

Given the dramatic differences between classical and quantum randomness, and the usefulness of both, it is reasonable to ask if there exist quantum analogues

of the aforementioned efficient simulators of classical random functions. In fact, given the discussion above, it is clear that we should begin by asking if there even exist efficient simulators of random quantum states.

Simulating random states. The first problem of interest is thus to efficiently simulate the following ideal object: an oracle $\mathfrak{IS}(n)$ which contains a description of a perfectly Haar-random n-qubit pure state $|\varphi\rangle$, and which outputs a copy of $|\varphi\rangle$ whenever it is invoked. We first make an obvious observation: the classical analogue, which is simply to generate a random bitstring $x \leftarrow \{0,1\}^n$ and then produce a copy whenever asked, is completely trivial. In the quantum case, efficient simulation is only known against limited query algorithms (henceforth, adversaries.)

If the adversary has an a priori bound on the number of queries, then *state t-designs* suffice. These are indexed families $\{|\varphi_{k,t}\rangle : k \in K_t\}$ of pure states which perfectly emulate the standard uniform "Haar" measure on pure states, up to the first t moments. State t-designs can be sampled efficiently, and thus yield a stateless simulator for this case [5]. A recent work of Ji, Liu and Song considered the case of polynomial-time adversaries [18]. They defined a notion of *pseudorandom states* (PRS), which appear Haar-random to polynomial-time adversaries who are allowed as many copies of the state as they wish. They also showed how to construct PRS efficiently, thus yielding a stateless simulator for this class of constrained adversaries [18]; see also [10].

The case of arbitrary adversaries is, to our knowledge, completely unexplored. In particular, before this work it was not known whether simulating $\mathfrak{IS}(n)$ against adversaries with no a priori bound on query or time complexity is possible, even if given polynomial space (in n and the number of queries) and unlimited time. Note that, while the state family constructions from [18,10] could be lifted to the unconditional security setting by instantiating them with random instead of pseudorandom functions, this would require space exponential in n regardless of the number of queries.

Simulating random unitaries. In the case of simulating random unitaries, the ideal object is an oracle $\mathfrak{IU}$ (n) which contains a description of a perfectly Haar-random n-qubit unitary operator U, and applies U to its input whenever it is invoked. The classical analogue is the well-known Random Oracle, and can be simulated perfectly using the aforementioned technique of lazy sampling. In the quantum case, the situation is even less well-understood than in the case of states.

For the case of query-limited adversaries, we can again rely on design techniques: (approximate) *unitary t-designs* can be sampled efficiently, and suffice for the task [11,21]. Against polynomial-time adversaries, Ji, Liu and Song defined the natural notion of a *pseudorandom unitary* (or PRU) and described candidate constructions [18]. Unfortunately, at this time there are no provably secure constructions of PRUs. As in the case of states, the case of arbitrary adversaries is completely unexplored. Moreover, one could a priori plausibly conjecture that

simulating $\mathfrak{IU}$ might even be impossible. The no-cloning property seems to rule out examining input states, which in turn seems to make it quite difficult for a simulator to correctly identify the overlap between multiple queries, and then answer correspondingly.

Extensions. While the above problems already appear quite challenging, we mention several natural extensions that one might consider. First, for the case of repeatedly sampling a random state $|\varphi\rangle$, one would ideally want some additional features, such as the ability to apply the two-outcome measurement $\{|\varphi\rangle\langle\varphi|, \mathbb{1} - |\varphi\rangle\langle\varphi|\}$ (*verification*) or the reflection $\mathbb{1} - 2|\varphi\rangle\langle\varphi|$. In the case of pseudorandom simulation, these additional features can be used to create a (computationally secure) quantum money scheme [18]. For the case of simulating random unitaries, we might naturally ask that the simulator for a unitary U also has the ability to respond to queries to $U^{-1} = U^\dagger$.

1.3 This work

In this work, we make significant progress on the above problems, by giving the first simulators for both random states and random unitaries, which are convincing to arbitrary adversaries. We also give an application of our sampling ideas: the construction of a new quantum money scheme, which provides information-theoretic security guarantees against both forging and tracing.

We begin by remarking that our desired simulators must necessarily be stateful, for both states and unitaries. Indeed, since approximate t-designs have $\Omega((2^{2n}/t)^{2t})$ elements (see, e.g., [25] which provides a more fine-grained lower bound), a stateless approach would require superpolynomial space simply to store an index from a set of size $\Omega((2^{2n}/t(n))^{2t(n)})$ for all polynomials $t(n)$.

In the following, we give a high-level overview of our approach for each of the two simulation problems of interest.

Simulating random states. As discussed above, we wish to construct an efficient simulator $\mathfrak{ES}(n)$ for the ideal oracle $\mathfrak{IS}(n)$. For now we focus on simulating the procedure which generates copies of the fixed Haar-random state; we call this $\mathfrak{IS}(n).\mathsf{Gen}$. We first note that the mixed state observed by the adversary after t queries to $\mathfrak{IS}(n).\mathsf{Gen}$ is the expectation of the projector onto t copies of $|\psi\rangle$. Equivalently, it is the (normalized) projector onto the *symmetric subspace* $\mathbf{Sym}_{n,t}$ of $(\mathbb{C}^{2^n})^{\otimes t}$:

$$\tau_t = \mathbb{E}_{\psi\sim\mathrm{Haar}}|\psi\rangle\langle\psi|^{\otimes t} \propto \Pi_{\mathrm{Sym}^t\mathbb{C}^{2^n}} . \tag{1}$$

Recall that $\mathbf{Sym}_{n,t}$ is the subspace of $(\mathbb{C}^{2^n})^{\otimes t}$ of vectors which are invariant under permutations of the t tensor factors. Our goal will be to maintain an entangled state between the adversary $\mathcal{A}$ and our oracle simulator $\mathfrak{ES}$ such that the reduced state on the side of $\mathcal{A}$ is τ_t after t queries. Specifically, the joint state will be the maximally entangled state between the $\mathbf{Sym}_{n,t}$ subspace of the

t query output registers received by $\mathcal{A}$, and the $\mathbf{Sym}_{n,t}$ subspace of t registers held by $\mathfrak{ES}$. If we can maintain this for the first t queries, then it's not hard to see that there exists an isometry $V^{t\to t+1}$ which, by acting only on the state of $\mathfrak{ES}$, implements the extension from the t-fold to the $(t+1)$-fold joint state.

The main technical obstacle, which we resolve, is showing that $V^{t\to t+1}$ can be performed efficiently. To achieve this, we develop some new algorithmic tools for working with symmetric subspaces, including an algorithm for coherent preparation of its basis states. We let A denote an n-qubit register, A_j its indexed copies, and $A^t = A_1 \cdots A_t$ t-many indexed copies (and likewise for B.) We also let $\{|\mathrm{Sym}(\alpha)\rangle : \alpha \in S^{\uparrow}_{n,t}\}$ denote a particular orthonormal basis set for $\mathbf{Sym}_{n,t}$, indexed by some set $S^{\uparrow}_{n,t}$ (see Section 3 for definitions of these objects.)

Theorem 1. *For each n and t, there exists a polynomial-time quantum algorithm which implements an isometry $V = V^{t\to t+1}$ from B^t to $A_{t+1}B^{t+1}$ such that, up to negligible trace distance,*

$$(\mathbb{1}_{A^t} \otimes V) \sum_{\alpha\in S^{\uparrow}_{n,t}} |\mathrm{Sym}(\alpha)\rangle_{A^t} |\mathrm{Sym}(\alpha)\rangle_{B^t} = \sum_{\beta\in S^{\uparrow}_{n,t+1}} |\mathrm{Sym}(\beta)\rangle_{A^{t+1}} |\mathrm{Sym}(\beta)\rangle_{B^{t+1}} .$$

Above, V is an operator defined to apply to a specific subset of registers of a state. When no confusion can arise, in such settings we will abbreviate $\mathbb{1} \otimes V$—the application of this operator on the entire state—as simply V.

It will be helpful to view $V^{t\to t+1}$ as first preparing $|0^n\rangle_{A_{t+1}}|0^n\rangle_{B_{t+1}}$ and then applying a unitary $U^{t\to t+1}$ on $A_{t+1}B^{t+1}$. Theorem 1 then gives us a way to answer Gen queries efficiently, as follows. For the first query, we prepare a maximally entangled state $|\phi^+\rangle_{A_1B_1}$ across two n-qubit registers A_1 and B_1, and reply with register A_1. Note that $\mathbf{Sym}_{n,1} = \mathbb{C}^{2^n}$. For the second query, we prepare two fresh registers A_2 and B_2, both in the $|0^n\rangle$ state, apply $U^{1\to 2}$ on $A_2B_1B_2$, return A_2, and keep B_1B_2. For the t-th query, we proceed similarly, preparing fresh blank registers $A_{t+1}B_{t+1}$, applying $U^{t\to t+1}$, and then outputting the register A_{t+1}.

With this approach, as it turns out, there is also a natural way to respond to verification queries Ver and reflection queries $\mathsf{Reflect}$. The ideal functionality $\mathfrak{IS}.\mathsf{Ver}$ is to apply the two-outcome measurement $\{|\varphi\rangle\langle\varphi|, \mathbb{1} - |\varphi\rangle\langle\varphi|\}$ corresponding to the Haar-random state $|\varphi\rangle$. To simulate this after producing t samples, we apply the inverse of $U^{t-1\to t}$, apply the measurement $\{|0^{2n}\rangle\langle 0^{2n}|, \mathbb{1} - |0^{2n}\rangle\langle 0^{2n}|\}$ to A_tB_t, reapply $U^{t-1\to t}$, and then return A_t together with the measurement outcome (i.e., yes/no). For $\mathfrak{IS}.\mathsf{Reflect}$, the ideal functionality is to apply the reflection $\mathbb{1} - 2|\varphi\rangle\langle\varphi|$ through the state. To simulate this, we perform a sequence of operations analogous to Ver, but apply a phase of -1 on the $|0^{2n}\rangle$ state of A_tB_t instead of measuring.

Our main result on simulating random states is to establish that this collection of algorithms correctly simulates the ideal object $\mathfrak{IS}$, in the following sense.

Theorem 2. *There exists a stateful quantum algorithm $\mathfrak{ES}(n,\epsilon)$ which runs in time polynomial in n, $\log(1/\epsilon)$, and the number of queries q submitted to it, and*

satisfies the following. For all oracle algorithms $\mathcal{A}$,

$$\left|\Pr\left[\mathcal{A}^{\mathfrak{IS}(n)}=1\right]-\Pr\left[\mathcal{A}^{\mathfrak{ES}(n,\epsilon)}=1\right]\right|\leq\epsilon\,.$$

A complete description of our construction, together with the proofs of Theorem 1 and Theorem 2, are given in Section 3.

We remark that, if one can give a certain mild a-priori bound on the number of queries that will be made to the state sampler, an alternative construction[4] based on the compressed oracle technique of Zhandry [31] and the aforementioned work by Ji, Liu and Song [18] becomes possible. We describe this construction in Section 3.3.

Application: untraceable quantum money. To see that the efficient state sampler leads to a powerful quantum money scheme, consider building a scheme where the bank holds the ideal object $\mathfrak{IS}$. The bank can mint bills by $\mathfrak{IS}.\mathsf{Gen}$, and verify them using $\mathfrak{IS}.\mathsf{Ver}$. As each bill is guaranteed to be an identical and Haar-random state, it is clear that this scheme should satisfy perfect unforgeability and untraceability, under quite strong notions of security.

By Theorem 7, the same properties should carry over for a money scheme built on $\mathfrak{ES}$, provided ϵ is sufficiently small. We call the resulting scheme *Haar money*. Haar money is an information-theoretically secure analogue of the scheme of [18], which is based on pseudorandom states. We remark that our scheme requires the bank to have quantum memory and to perform quantum communication with the customers. However, given that quantum money already requires customers to have large-scale, high-fidelity quantum storage, these additional requirements seem reasonable.

The notions of correctness and unforgeability (often called completeness and soundness) for quantum money are well-known (see, e.g., [1].) Correctness asks that honestly generated money schemes should verify, i.e., $\mathsf{Ver}(\mathsf{Mint})$ should always accept. Unforgeability states that an adversary with k bills and oracle access to Ver should not be able to produce a state on which $\mathsf{Ver}^{\otimes k+1}$ accepts. In this work, we consider untraceable quantum money (also called "quantum coins" [24].) We give a formal security definition for untraceability, which states that an adversary $\mathcal{A}$ with oracle access to Ver and Mint cannot do better than random guessing in the following experiment:

1. $\mathcal{A}$ outputs some candidate bill registers $\{M_j\}$ and a permutation π;
2. $b\leftarrow\{0,1\}$ is sampled, and if $b=1$ the registers $\{M_j\}$ are permuted by π; each candidate bill is verified and the failed ones are discarded;
3. $\mathcal{A}$ receives the rest of the bills and the entire internal state of the bank, and outputs a guess b' for b.

Theorem 3. *The Haar money scheme $\mathfrak{HM}$, defined by setting*

1. $\mathfrak{HM}.\mathsf{Mint}=\mathfrak{ES}(n,\mathsf{negl}(n)).\mathsf{Gen}$

[4] We thank Zvika Brakerski for pointing out this alternative approach.

2. $\mathfrak{HM}.\mathsf{Ver} = \mathfrak{ES}(n, \mathsf{negl}(n)).\mathsf{Ver}$

is a correct quantum money scheme which satisfies information-theoretic unforgeability and untraceability.

One might reasonably ask if there are even stronger definitions of security for quantum money. Given its relationship to the ideal state sampler, we believe that Haar money should satisfy almost any notion of unforgeability and untraceability, including composable notions. We also remark that, based on the structure of the state simulator, which maintains an overall pure state supported on two copies of the symmetric subspace of banknote registers, it is straightforward to see that the scheme is also secure against an "honest but curious" or "specious" [26,15] bank. We leave the formalization of these added security guarantees to future work.

Sampling Haar-random unitaries. Next, we turn to the problem of simulating Haar-random unitary operators. In this case, the ideal object $\mathfrak{IU}(n)$ initially samples a description of a perfectly Haar-random n-qubit unitary U, and then responds to two types of queries: $\mathfrak{IU}.\mathsf{Eval}$, which applies U, and $\mathfrak{IU}.\mathsf{Invert}$, which applies $U^\dagger$. In this case, we are able to construct a stateful simulator that runs in space polynomial in n and the number of queries q, and is *exactly indistinguishable* from $\mathfrak{IU}(n)$ to arbitrary adversaries. Our result can be viewed as a polynomial-space quantum analogue of the classical technique of lazy sampling for random oracles.

Our high-level approach is as follows. For now, suppose the adversary $\mathcal{A}$ only makes parallel queries to Eval. If the query count t of $\mathcal{A}$ is a priori bounded, we can simply sample an element of a unitary t-design. We can also do this coherently: prepare a quantum register I in uniform superposition over the index set of the t-design, and then apply the t-design controlled on I. Call this efficient simulator $\mathfrak{EU}_t$. Observe that the effect of t parallel queries is just the application of the *t-twirling channel* $\mathcal{T}^{(t)}$ to the t input registers [11], and that $\mathfrak{EU}_t$ simulates $\mathcal{T}^{(t)}$ faithfully. What is more, it applies a *Stinespring dilation*[5] [27] of $\mathcal{T}^{(t)}$ with dilating register I.

Now suppose $\mathcal{A}$ makes an "extra" query, i.e., query number $t+1$. Consider an alternative Stinespring dilation of $\mathcal{T}^{(t)}$, namely the one implemented by $\mathfrak{EU}_{t+1}$ when queried t times. Recall that all Stinespring dilations of a quantum channel are equivalent, up to a partial isometry on the dilating register. It follows that there is a partial isometry, acting on the private space of $\mathfrak{EU}_t$, that transforms the dilation of $\mathcal{T}^{(t)}$ implemented by $\mathfrak{EU}_t$ into the dilation of $\mathcal{T}^{(t)}$ implemented by $\mathfrak{EU}_{t+1}$. If we implement this transformation, and then respond to $\mathcal{A}$ as prescribed by $\mathfrak{EU}_{t+1}$, we have achieved perfect indistinguishability against the additional query. By iterating this process, we see that the a priori bound on the number

[5] The Stinespring dilation of a quantum channel is an isometry with the property that the quantum channel can be implemented by applying the isometry and subsequently discarding an auxiliary register.

of queries is no longer needed. We let $\mathfrak{EU}$ denote the resulting simulator. The complete construction is described in Construction 4 below.

Our high-level discussion above did not take approximation into account. All currently known efficient constructions of t-designs are approximate. Here, we take a different approach: we will implement our construction using *exact t-designs*. This addresses the issue of adaptive queries: if there exists an adaptive-query distinguisher with nonzero distinguishing probability, then by post-selection there also exists a parallel-query one via probabilistic teleportation. This yields that the ideal and efficient unitary samplers are perfectly indistinguishable to arbitrary adversaries.

Theorem 4. *For all oracle algorithms $\mathcal{A}$,* $\Pr\left[\mathcal{A}^{\mathfrak{IU}(n)} = 1\right] = \Pr\left[\mathcal{A}^{\mathfrak{EU}(n)} = 1\right]$.

The existence of exact unitary t-designs for all t is a fairly recent result. It follows as a special case of a result of Kane [19], who shows that designs exist for all finite-dimensional vector spaces of well-behaved functions on path-connected topological spaces. He also gives a simpler result for homogeneous spaces when the vector space of functions is invariant under the symmetry group action. Here, the number of elements of the smallest design is bounded just in terms of the dimension of the space of functions. The unitary group is an example of such a space, and the dimension of the space of homogeneous polynomials of degree t in both U and $U^\dagger$ can be explicitly derived, see e.g. [25]. This yields the following.

Corollary 1. *The space complexity of $\mathfrak{EU}(n)$ for q queries is bounded from above by $2q(2n + \log e) + O(\log q)$.*

An alternative approach. We now sketch another potential approach to lazy sampling of unitaries. Very briefly, this approach takes a representation-theoretic perspective and suggests that the Schur transform [6] could lead to a polynomial-time algorithm for lazy sampling Haar-random unitaries. The discussion below uses tools and language from quantum information theory and the representation theory of the unitary and symmetric groups to a much larger extent than the rest of the article, and is not required for understanding our main results.

We remark that the analogous problem of lazy sampling a quantum oracle for a random classical function was recently solved by Zhandry [31]. One of the advantages of Zhandry's technique is that it partly recovers the ability to inspect previously made queries, an important feature of classical lazy sampling. The key insight is that the simulator can implement the Stinespring dilation of the oracle channel, and thus record the output of the *complementary channel*.[6] As the classical function is computed via XOR, changing to the $\mathbb{Z}_2^n$-Fourier basis makes the recording property explicit. It also allows for an efficient implementation.

In the case of Haar-random unitary oracles, we can make an analogous observation. Consider an algorithm that makes t parallel queries to U. The relevant

[6] The complementary channel of a quantum channel maps the input to the auxiliary output of the Stinespring dilation isometry.

Fourier transform is now over the unitary group, and is given by the Schur transform [6]. By Schur-Weyl duality (see e.g. [13]), the decomposition of $\left(\mathbb{C}^{2^n}\right)^{\otimes t}$ into irreducible representations is given by

$$\left(\mathbb{C}^d\right)^{\otimes t} \cong \bigoplus_{\lambda \vdash_d t} [\lambda] \otimes V_{\lambda,d}. \tag{2}$$

Here $\lambda \vdash_d t$ means λ is any partition of t into at most d parts, $[\lambda]$ is the Specht module of S_t, and $V_{\lambda,d}$ is the Weyl module of $U(d)$, corresponding to the partition λ, respectively. By Schur's lemma, the t-twirling channel acts as

$$\mathcal{T}^{(t)} = \bigoplus_{\lambda \vdash_d t} \mathrm{id}_{[\lambda]} \otimes \Lambda_{V_{\lambda,d}}, \tag{3}$$

where id is the identity channel, and $\Lambda = \mathrm{Tr}(\cdot)\tau$ with the maximally mixed state τ is the depolarizing channel. We therefore obtain a Stinespring dilation of the t-twirling channel as follows. Let $\tilde{B}, \tilde{B}'$ be registers with Hilbert spaces

$$\mathcal{H}_{\tilde{B}} = \mathcal{H}_{\tilde{B}'} = \bigotimes_{\lambda \vdash_d t} V_{\lambda,d} \tag{4}$$

and denote the subregisters by $\tilde{B}_\lambda$ and $\tilde{B}'_\lambda$, respectively. Let further $|\phi^+\rangle_{\tilde{B}\tilde{B}'}$ be the standard maximally entangled state on these registers, and let C be a register whose dimension is the number of partitions of t (into at most 2^n parts). Define the isometry

$$\hat{V}_{A^t\tilde{B}\to A^t\tilde{B}C} = \bigoplus_{\lambda \vdash_d t} F_{V_{\lambda,d}\tilde{B}_\lambda} \otimes \mathbb{I}_{[\lambda]} \otimes |\lambda\rangle_C \tag{5}$$

In the above equation $V_{\lambda,d}$ and $[\lambda]$ are understood to be subspaces of A^t, the identity operators on $\tilde{B}_\mu$, $\mu \neq \lambda$ are omitted and F is the swap operator. By (3), a Stinespring dilation of the t-twirling channel is then given by

$$V_{A^t\to A^t\tilde{B}\tilde{B}'C} = \hat{V}_{A^t\tilde{B}\to A^t\tilde{B}C}|\phi^+\rangle_{\tilde{B}\tilde{B}'}. \tag{6}$$

By the equivalence of all Stinespring dilations, the exists an isometry $W_{\hat{B}_t\to\tilde{B}\tilde{B}'C}$ that transforms the state register of $\mathfrak{EU}(n)$ after t parallel queries so that the global state is the same as if the Stinespring dilation above had been applied to the t input registers. But now the quantum information that was contained in the subspace $V_{\lambda,d}$ of the algorithm's query registers can be found in register $\tilde{B}_\lambda$.

1.4 Organization

The remainder of the paper is organized as follows. In Section 2, we recall some basic notation and facts, and some lemmas concerning coherent preparation of certain generic families of quantum states. The proofs for these lemmas are given in the full version [4]. We also describe stateful machines, which will be our model

for thinking about the aforementioned ideal objects and their efficient simulators. In Section 3 we describe our efficient simulator for Haar-random states, and in Section 4 we describe our polynomial-space simulator for Haar-random unitaries. We end by describing the Haar money scheme and establishing its security in Section 5.

1.5 Acknowledgments

The authors thank Zvika Brakerski for suggesting the alternative construction based on compressed oracles. We thank Yi-Kai Liu, Carl Miller, and Fang Song for helpful comments on an earlier draft. CM thanks Michael Walter for discussions about t-designs. CM was funded by a NWO VIDI grant (Project No. 639.022.519) and a NWO VENI grant (Project No. VI.Veni.192.159). GA acknowledges support from NSF grant CCF-1763736. GA was supported by the Dutch Research Council (NWO) through a travel grant - 040.11.708.

2 Preliminaries

Given a fixed-size (e.g., n-qubit) register A, we will use $A_1, A_2, \ldots$ to denote indexed copies of A. We will use A^t to denote a register consisting of t indexed copies of A, i.e., $A^t = A_1 A_2 \cdots A_t$. Unless stated otherwise, distances of quantum states are measured in the trace distance, i.e., $d(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ where $\|X\|_1 = \mathrm{Tr}\left[\sqrt{X^\dagger X}\right]$. Distances of unitary operators are measured in the operator norm.

We will frequently apply operators to some subset of a larger collection of registers. In that context, we will use register indexing to indicate which registers are being acted upon, and suppress identities to simplify notation. The register indexing will also be suppressed when it is clear from context. For example, given an operator $X_{A\to B}$ and some state ρ on registers A and C, we will write $X(\rho)$ in place of $(X \otimes \mathbb{1}_C)(\rho)$ to denote the state on BC resulting from applying X to the A register of ρ.

We let $|\phi^+\rangle_{AA'}$ denote the maximally entangled state on registers A and A'. For a linear operator X and some basis choice, we denote its transpose by X^T.

Lemma 1 (Mirror lemma; see, e.g., [22]). *For $X_{A\to B}$ a linear operator,*

$$X_{A\to B}|\phi^+\rangle_{AA'} = \sqrt{\frac{\dim(B)}{\dim(A)}} X^T_{B'\to A'}|\phi^+\rangle_{BB'}\,.$$

2.1 Unitary designs

Let μ_n be the Haar measure on the unitary group $\mathrm{U}(2^n)$. We define the Haar t-twirling channel $\mathcal{T}^{(t)}_{\mathrm{Haar}}$ by

$$\mathcal{T}^{(t)}_{\mathrm{Haar}}(X) = \int_{\mathrm{U}(2^n)} U^{\otimes t} X \left(U^{\otimes t}\right)^\dagger \mathrm{d}\mu(U). \tag{7}$$

For a finite subset $D \subset \mathrm{U}(2^n)$, we define the t-twirling map with respect to D as

$$\mathcal{T}_D^{(t)}(X) = \frac{1}{|D|} \sum_{U \in D} U^{\otimes t} X \left(U^{\otimes t}\right)^\dagger . \tag{8}$$

An n-qubit unitary t-design is a finite set $D \subset \mathrm{U}(2^n)$ such that

$$\mathcal{T}_D^{(t)} = \mathcal{T}_{\mathrm{Haar}}^{(t)}(X) \tag{9}$$

Another twirling channel is the mixed twirling channels with ℓ applications of the unitary and $t-\ell$ applications of it's inverse,

$$\mathcal{T}_{\mathrm{Haar}}^{(\ell,t-\ell)}(\Gamma) = \int_{\mathrm{U}(2^n)} U^{\otimes \ell} \otimes \left(U^{\otimes (t-\ell)}\right)^\dagger \Gamma \left(U^{\otimes \ell}\right)^\dagger \otimes U^{\otimes (t-\ell)} \mathrm{d}\mu(U). \tag{10}$$

The mixed twirling channel $\mathcal{T}_D^{(\ell,t-\ell)}$ for a finite set $D \subset \mathrm{U}(2^n)$ is also defined analogous to Equation (8). As our definition of unitary t-designs is equivalent to one based on the expectation values of polynomials (see, e.g., [21]), we easily obtain the following.

Proposition 1. *Let D be an n-qubit unitary t-design and $0 \leq \ell \leq t$. Then*

$$\mathcal{T}_{\mathrm{Haar}}^{(\ell,t-\ell)} = \mathcal{T}_D^{(\ell,t-\ell)} \tag{11}$$

Finite exact unitary t-designs exist. In particular, one can apply the following theorem to obtain an upper bound on their minimal size. Here, a design for a function space W on a topological space X with measure μ is a finite set $D \subset X$ such that the expectation of a function $f \in W$ is the same whether it is taken over X according to μ or over the uniform distribution on D.

Theorem 5 ([19], Theorem 10). *Let X be a homogeneous space, μ an invariant measure on X and W a M-dimensional vector subspace of the space of real functions on X that is invariant under the symmetry group of X, where $M > 1$. Then for any $N > M(M-1)$, there exists a W-design for X of size N. Furthermore, there exists a design for X of size at most $M(M-1)$.*

The case of unitary t-designs is the one where $X = \mathrm{U}(2^n)$ is acting on itself (e.g., on the left), μ is the Haar measure, and W is the vector space of homogeneous polynomials of degree t in both U and $U^\dagger$[7]. The dimension of this space is

$$M_t = \binom{2^{2n}+t-1}{t}^2 \leq \left(\frac{e(2^{2n}+t-1)}{t}\right)^t , \tag{12}$$

see e.g. [25]. We therefore get

Corollary 2. *For all n, there exists an exact n-qubit unitary t-design with a number of elements which is at most*

$$\left(\frac{e(2^{2n}+t-1)}{t}\right)^{2t} .$$

[7] The output of the twirling channel (7) is a matrix of such polynomials.

2.2 Real and ideal stateful machines

We will frequently use stateful algorithms with multiple "interfaces" which allow a user to interact with the algorithm. We will refer to such objects as *stateful machines*. We will use stateful machines to describe functionalities (and implementations) of collections of oracles which relate to each other in some way. For example, one oracle might output a fixed state, while another oracle reflects about that state.

Definition 1 (Stateful machine). *A stateful machine $\mathcal{S}$ consists of:*

- *A finite set Λ, whose elements are called interfaces. Each interface $\mathcal{I} \in \Lambda$ has two fixed parameters $n_{\mathcal{I}} \in \mathbb{N}$ (input size) and $m_{\mathcal{I}} \in \mathbb{N}$ (output size), and a variable $t_{\mathcal{I}}$ initialized to 1 (query counter.)*
- *For each interface $\mathcal{I} \in \Lambda$, a sequence of quantum algorithms $\{\mathcal{S}.\mathcal{I}_j : j = 1, 2, \dots\}$. Each $\mathcal{S}.\mathcal{I}_j$ has an input register of $n_{\mathcal{I}}$ qubits, an output register of $m_{\mathcal{I}}$ qubits, and is allowed to act on an additional shared work register R (including the ability to add/remove qubits in R.) In addition, each $\mathcal{S}.\mathcal{I}_j$ increments the corresponding query counter $t_{\mathcal{I}}$ by one.*

The typical usage of a stateful machine $\mathcal{S}$ is as follows. First, the work register R is initialized to be empty, i.e., no qubits. After that, whenever a user invokes an interface $\mathcal{S}.\mathcal{I}$ and supplies $n_{\mathcal{I}}$ qubits in an input register M, the algorithm $\mathcal{S}.\mathcal{I}_{t_{\mathcal{I}}}$ is invoked on registers M and R. The contents of the output register are returned to the user, and the new, updated work register remains for the next invocation. We emphasize that the work register is shared between all interfaces.

We remark that we will also sometimes define *ideal machines*, which behave outwardly like a stateful machine but are not constrained to apply only maps which are implementable in finite space or time. For example, an ideal machine can have an interface that implements a perfectly Haar-random unitary U, and another interface which implements $U^\dagger$.

2.3 Some state preparation tools

We now describe some algorithms for efficient coherent preparation of certain quantum state families. The proofs for the following lemmas can be found in the full version [4]. We begin with state families with polynomial support.

Lemma 2. *Let $|\varphi\rangle = \sum_{x\in\{0,1\}^n} \varphi(x)|x\rangle$ be a family of quantum states whose amplitudes φ have an efficient classical description $\tilde{\varphi}$, and such that $|\{x : \varphi(x) \neq 0\}| \leq \mathsf{poly}(n)$. Then there exists a quantum algorithm $\mathcal{P}$ which runs in time polynomial in n and $\log(1/\epsilon)$ and satisfies $\|\mathcal{P}|\tilde{\varphi}\rangle|0^n\rangle - |\tilde{\varphi}\rangle|\varphi\rangle\|_2 \leq \epsilon$.*

Given a set $S \subset \{0,1\}^n$, we let

$$|S\rangle := \frac{1}{\sqrt{|S|}} \sum_{x\in S} |x\rangle \qquad \text{and} \qquad |\bar{S}\rangle := \frac{1}{\sqrt{2^n - |S|}} \sum_{x\in\{0,1\}\setminus S} |x\rangle$$

denote the states supported only on S and its set complement $\bar{S}$, respectively. Provided that S has polynomial size, we can perform coherent preparation of both state families efficiently: the former by Lemma 2 and the latter via the below.

Lemma 3. *Let $S \subset \{0,1\}^n$ be a family of sets of size $\mathsf{poly}(n)$ with efficient description $\tilde{S}$, and let $\epsilon > 0$. There exists a quantum algorithm $\mathcal{P}$ which runs in time polynomial in n and $\log(1/\epsilon)$ and satisfies*

$$\left\| \mathcal{P}|\tilde{S}\rangle_A|0^n\rangle_B - |\tilde{S}\rangle_A|\bar{S}\rangle_B \right\|_2 \leq \epsilon\,.$$

Finally, we show that if two orthogonal quantum states can be prepared, then so can an arbitrary superposition of the two.

Lemma 4. *Let $|\zeta_{0,j}\rangle, |\zeta_{1,j}\rangle$ be two familes of n-qubit quantum states such that $\langle\zeta_{0,j} \mid \zeta_{1,j}\rangle = 0$ for all j, and such that there exists a quantum algorithm $\mathcal{P}_b$ which runs in time polynomial in n and $\log(1/\epsilon)$ and satisfies $\|\mathcal{P}_b|j\rangle|0^n\rangle - |j\rangle|\zeta_{b,j}\rangle\|_2 \leq \epsilon$ for $b \in \{0,1\}$.*

For $z_0, z_1 \in \mathbb{C}$ such that $|z_0|^2 + |z_1|^2 = 1$, let $\tilde{z}$ denote a classical description of (z_0, z_1) to precision at least ϵ. There exists a quantum algorithm $\mathcal{Q}$ which runs in time polynomial in n and $\log(1/\epsilon)$ and satisfies

$$\left\| \mathcal{Q}|j\rangle|\tilde{z}\rangle|0^n\rangle - |j\rangle|\tilde{z}\rangle\big(z_0|\zeta_{0,j}\rangle + z_1|\zeta_{1,j}\rangle\big) \right\|_2 \leq \epsilon\,. \tag{13}$$

3 Simulating a Haar-random state oracle

3.1 The problem, and our approach

We begin by defining the ideal object we'd like to emulate. Here we deviate slightly from the discussion above, in that we ask for the reflection oracle to also accept a (quantum) control bit.

Construction 1 (Ideal state sampler) *The ideal n-qubit state sampler is an ideal machine $\mathfrak{IS}(n)$ with interfaces* $(\mathsf{Init}, \mathsf{Gen}, \mathsf{Ver}, \mathsf{CReflect})$*, defined as follows.*

1. $\mathfrak{IS}(n).\mathsf{Init}$: *takes no input; samples a description $\tilde{\varphi}$ of an n-qubit state $|\varphi\rangle$ from the Haar measure.*
2. $\mathfrak{IS}(n).\mathsf{Gen}$: *takes no input; uses $\tilde{\varphi}$ to prepare a copy of $|\varphi\rangle$ and outputs it.*
3. $\mathfrak{IS}(n).\mathsf{Ver}$: *receives n-qubit input; uses $\tilde{\varphi}$ to apply the measurement $\{|\varphi\rangle\langle\varphi|, \mathbb{1} - |\varphi\rangle\langle\varphi|\}$; return the post-measurement state and output acc in the first case and rej in the second.*
4. $\mathfrak{IS}(n).\mathsf{CReflect}$: *receives $(n+1)$-qubit input; uses $\tilde{\varphi}$ to implement the controlled reflection $R_\varphi := |0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes (\mathbb{1} - 2\,|\varphi\rangle\langle\varphi|)$ about $|\varphi\rangle$.*

We assume that Init is called first, and only once; the remaining oracles can then be called indefinitely many times, and in any order. If this is inconvenient for some application, one can easily adjust the remaining interfaces to invoke Init if that has not been done yet. We remark that Ver can be implemented with a single query to $\mathsf{CReflect}$.

Lemma 5. Ver *can be simulated with one application of* CReflect.

Proof. Prepare an ancilla qubit in the state $|+\rangle$ and apply reflection on the input controlled on the ancilla. Then apply H to the ancilla qubit and measure it. Output all the qubits, with the ancilla interpreted as $1 = \mathsf{acc}$ and $0 = \mathsf{rej}$. □

Our goal is to devise a stateful simulator for Construction 1 which is efficient. Efficient here means that, after t total queries to all interfaces (i.e., Init, Gen, Ver, and CReflect), the simulator has expended time polynomial in n, t, and $\log(1/\epsilon)$.

As described in Section 1.3, our approach will be to ensure that, for every t, the state shared between the adversary $\mathcal{A}$ and our stateful oracle simulator $\mathfrak{ES}$ will be maximally entangled between two copies of the t-fold symmetric subspace $\mathbf{Sym}_{n,t}$: one held by $\mathcal{A}$, and the other by $\mathfrak{ES}$. The extension from the t-fold to the $(t+1)$-fold joint state will be performed by an isometry $V^{t\to t+1}$ which acts only on the state of $\mathfrak{ES}$ and two fresh n-qubit registers A_{t+1} and B_{t+1} initialized by $\mathfrak{ES}$. After V is applied, A_{t+1} will be given to $\mathcal{A}$. As we will show, V can be performed efficiently using some algorithmic tools for working with symmetric subspaces, which we will develop in the next section. This will yield an efficient way of simulating Gen. Simulation of Ver and CReflect will follow without much difficulty, as outlined in Section 1.3.

3.2 Some tools for symmetric subspaces

A basis for the symmetric subspace. We recall an explicit orthonormal basis of the symmetric subspace (see, e.g., [18] or [17].) Let

$$S^{\uparrow}_{n,t} = \left\{ \alpha \in (\{0,1\}^n)^t \,\middle|\, \alpha_1 \leq \alpha_2 \leq ... \leq \alpha_t \right\} \tag{14}$$

be the set of lexicographically-ordered t-tuples of n bit strings. For each $\alpha \in S^{\uparrow}_{n,t}$, define the unit vector

$$|\mathrm{Sym}(\alpha)\rangle = \left(t! \prod_{x\in\{0,1\}^n} f_x(\alpha)! \right)^{-\frac{1}{2}} \sum_{\sigma\in S_t} |\alpha_{\sigma(1)}\rangle|\alpha_{\sigma(2)}\rangle ... |\alpha_{\sigma(t)}\rangle. \tag{15}$$

Here, $f_x(\alpha)$ is the number of times the string x appears in the tuple α. The set $\{|\mathrm{Sym}(\alpha)\rangle : \alpha \in S^{\uparrow}_{n,t}\}$ is an orthonormal basis for $\mathrm{Sym}^t\mathbb{C}^{2^n}$. We remark that the Schmidt decomposition of $|\mathrm{Sym}(\alpha)\rangle$ with respect to the bipartition formed by the t-th register vs. the rest is given by

$$|\mathrm{Sym}(\alpha)\rangle = \sum_{x\in\{0,1\}^n} \sqrt{\frac{f_x(\alpha)}{t}} |\mathrm{Sym}(\alpha^{-x})\rangle|x\rangle, \tag{16}$$

where $\alpha^{-x} \in S^{\uparrow}_{n,t-1}$ is the tuple α with one copy of x removed.

Some useful algorithms. We now describe some algorithms for working in the above basis. Let A and B denote n-qubit registers. Recall that A_j denotes indexed copies of A and that A^t denotes $A_1 A_2 \cdots A_t$, and likewise for B. In our setting, the various copies of A will be prepared by the oracle simulator and then handed to the query algorithm at query time. The copies of B will be prepared by, and always remain with, the oracle simulator.

Proposition 2. *For each n, t and $\epsilon = 2^{-\mathrm{poly}(n,t)}$, there exists an efficiently implementable unitary $U_{n,t}^{\mathrm{Sym}}$ on A^t such that for all $\alpha \in S_{n,t}^{\uparrow}$, $U_{n,t}^{\mathrm{Sym}}|\alpha\rangle = |\mathrm{Sym}(\alpha)\rangle$ up to trace distance ϵ.*

Proof. Clearly, the operation

$$|\mathrm{Sym}(\alpha)\rangle|\beta\rangle \mapsto |\mathrm{Sym}(\alpha)\rangle|\beta \oplus \alpha\rangle \tag{17}$$

is efficiently implementable exactly, by XORing the classical sort function of the first register into the second register.

Let us now show that the operation $|\alpha\rangle \mapsto |\alpha\rangle|\mathrm{Sym}(\alpha)\rangle$ is also efficiently implementable (up to the desirable error) by exhibiting an explicit algorithm. We define it recursively in t, as follows. For $t = 1$, $\mathrm{Sym}(x) = x$ for all $x \in \{0,1\}^n$, so this case is simply the map $|x\rangle \mapsto |x\rangle|x\rangle$. Suppose now the operation $|\alpha\rangle \mapsto |\alpha\rangle|\mathrm{Sym}(\alpha)\rangle$ can be implemented for any $\alpha \in S_{n,t-1}^{\uparrow}$. The t-th level algorithm will begin by applying

$$|\alpha\rangle \mapsto |\alpha\rangle \sum_{x \in \{0,1\}^n} \sqrt{\frac{f_x(\alpha)}{t}}|x\rangle\,.$$

Since $f_x(\alpha)$ is nonzero for only t-many $x \in \{0,1\}^n$, this can be implemented efficiently by Lemma 2. Next, we perform $|\alpha\rangle|x\rangle \mapsto |\alpha\rangle|x\rangle|\alpha^{-x}\rangle$. Using the algorithm for $t-1$, we then apply $|\alpha\rangle|x\rangle|\alpha^{-x}\rangle \mapsto |\alpha\rangle|x\rangle|\alpha^{-x}\rangle|\mathrm{Sym}(\alpha^{-x})\rangle$, and uncompute α^{-x}. By (16), we have in total applied $|\alpha\rangle \mapsto |\alpha\rangle|\mathrm{Sym}(\alpha)\rangle$ so far. To finish the t-th level algorithm for approximating $|\alpha\rangle \mapsto |\mathrm{Sym}(\alpha)\rangle$, we simply apply (17) to uncompute α from the first register. □

Theorem 6 (Restatement of Theorem 1). *For each n, t and $\epsilon = 2^{-\mathrm{poly}(n,t)}$, there exists an efficiently implementable isometry $V^{t \to t+1}$ from B^t to $A_{t+1}B^{t+1}$ such that, up to trace distance ϵ,*

$$V : \sum_{\alpha \in S_{n,t}^{\uparrow}} |\mathrm{Sym}(\alpha)\rangle_{A^t}|\mathrm{Sym}(\alpha)\rangle_{B^t} \longmapsto \sum_{\beta \in S_{n,t+1}^{\uparrow}} |\mathrm{Sym}(\beta)\rangle_{A^{t+1}}|\mathrm{Sym}(\beta)\rangle_{B^{t+1}}\,.$$

We expect the techniques used here to generalize to other irreducible representations of the unitary group.

Proof. We describe the algorithm assuming all steps can be implemented perfectly. It is straightforward to check that each step can be performed to a sufficient accuracy that the accuracy of the entire algorithm is at least ϵ.

We will need a couple of simple subroutines. First, given $\alpha \in S^{\uparrow}_{n,t}$ and $x \in \{0,1\}^n$, we define α^{+x} to be the element of $S^{\uparrow}_{n,t+1}$ produced by inserting x at the first position such that the result is still lexicographically ordered. One can perform this reversibly via $|\alpha\rangle|0^n\rangle|x\rangle \mapsto |\alpha\rangle|x\rangle|x\rangle \mapsto |\alpha^{+x}\rangle|x\rangle$. Second, we will need to do coherent preparation of the state

$$|\psi_\alpha\rangle = \sum_{x\in\{0,1\}^n} \sqrt{\frac{1+f_x(\alpha)}{2^n+t}}|x\rangle\,. \tag{18}$$

For any given $\alpha \in S^{\uparrow}_{n,t}$, the state $|\psi_\alpha\rangle$ can be prepared via the preparation circuit for the two orthogonal components of the state whose supports are $\{x : f_x(\alpha) > 0\}$ and $\{x : f_x(\alpha) = 0\}$. These two components can be prepared coherently using Lemma 2 and Lemma 3, respectively. Their superposition can be prepared with Lemma 4. All together, we get an algorithm for $|\alpha\rangle|0^n\rangle \mapsto |\alpha\rangle|\psi_\alpha\rangle$.

The complete algorithm is a composition of several efficient routines. We describe this below, explicitly calculating the result for the input states of interest. For readability, we omit overall normalization factors.

$$\begin{aligned}
&\sum_\alpha |\mathrm{Sym}(\alpha)\rangle_{A^t}|\mathrm{Sym}(\alpha)\rangle_{B^t} && \\
\longmapsto &\sum_\alpha |\mathrm{Sym}(\alpha)\rangle_{A^t}|0^n\rangle|\mathrm{Sym}(\alpha)\rangle_{B^t}|0^n\rangle && \text{add working registers} \\
\longmapsto &\sum_\alpha |\mathrm{Sym}(\alpha)\rangle_{A^t}|0^n\rangle|\alpha\rangle_{B^t}|0^n\rangle && \text{apply } \left(U^{\mathrm{Sym}}_{n,t}\right)^\dagger \text{ to } B^t \\
\longmapsto &\sum_{\alpha,x} \sqrt{\frac{1+f_x(\alpha)}{2^n+t}}|\mathrm{Sym}(\alpha)\rangle_{A^t}|x\rangle|\alpha\rangle_{B^t}|0^n\rangle && \text{prepare } |\psi_\alpha\rangle \\
\longmapsto &\sum_{\alpha,x} \sqrt{\frac{1+f_x(\alpha)}{2^n+t}}|\mathrm{Sym}(\alpha)\rangle_{A^t}|x\rangle|\alpha^{+x}\rangle_{B^{t+1}} && \text{insert } x \text{ into } \alpha \\
\longmapsto &\sum_{\alpha,x} \sqrt{\frac{1+f_x(\alpha)}{2^n+t}}|\mathrm{Sym}(\alpha)\rangle_{A^t}|x\rangle_{A_{t+1}}|\mathrm{Sym}(\alpha^{+x})\rangle_{B^{t+1}} && \text{apply } U^{\mathrm{Sym}}_{n,t+1} \text{ to } B^{t+1}
\end{aligned}$$

To see that the last line above is the desired result, we observe that we can index the sum in the last line above in a more symmetric fashion: the sum is just taken over all pairs (α, β) such that the latter can be obtained from the former by adding one entry (i.e., the string x). But that is the same as summing over all pairs (α, β), such that the former can be obtained from the latter by *removing*

one entry.

$$\begin{aligned}
&\sum_{\alpha,x} \sqrt{\frac{1+f_x(\alpha)}{2^n+t}} |\mathrm{Sym}(\alpha)\rangle_{A^t} |x\rangle_{A_{t+1}} |\mathrm{Sym}(\alpha^{+x})\rangle_{B^{t+1}} \\
&= \sum_{\beta,x} \sqrt{\frac{f_x(\beta)}{2^n+t}} |\mathrm{Sym}(\beta^{-x})\rangle_{A^t} |x\rangle_{A_{t+1}} |\mathrm{Sym}(\beta)\rangle_{B^{t+1}} \\
&= \sqrt{\frac{t}{2^n+t}} \sum_{\beta} \left(\sum_{x} \sqrt{\frac{f_x(\beta)}{t}} |\mathrm{Sym}(\beta^{-x})\rangle_{A^t} |x\rangle_{A_{t+1}} \right) |\mathrm{Sym}(\beta)\rangle_{B^{t+1}} \\
&= \sqrt{\frac{t}{2^n+t}} \sum_{\beta} |\mathrm{Sym}(\beta)\rangle_{A^{t+1}} |\mathrm{Sym}(\beta)\rangle_{B^{t+1}}.
\end{aligned}$$

Here, the last equality is (16), and the prefactor is the square root of the quotient of the dimensions of the t- and $(t+1)$-copy symmetric subspaces, as required for a correct normalization of the final maximally entangled state. □

3.3 State sampler construction and proof

Construction 2 (Efficient state sampler) *Let n be a positive integer and ϵ a negligible function of n. The efficient n-qubit state sampler with precision ϵ is a stateful machine $\mathfrak{ES}(\epsilon, n)$ with interfaces* $(\mathsf{Init}, \mathsf{Gen}, \mathsf{Reflect})$*, defined below. For convenience, we denote the query counters by $t = t_{\mathsf{Gen}}$ and $q = t_{\mathsf{Reflect}}$ in the following.*

1. $\mathfrak{ES}(\epsilon, n).\mathsf{Init}$: *prepares the standard maximally entangled state $|\phi^+\rangle_{A_1B_1}$ on n-qubit registers A_1 and B_1, and stores both A_1 and B_1.*
2. $\mathfrak{ES}(\epsilon, n).\mathsf{Gen}$: *On the first query, outputs register A_1. On query t, takes as input registers B^{t-1} and produces registers A_tB^t by applying the isometry $V^{t-1\to t}$ from Theorem 6 with accuracy $\epsilon 2^{-(t+2q)}$; then it outputs A_t and stores B^t.*
3. $\mathfrak{ES}(\epsilon, n).\mathsf{CReflect}$: *On query q with input registers CA^*, do the following controlled on the qubit register C: apply $\left(U^{t-1\to t}\right)^\dagger$, a unitary implementation of $V^{t-1\to t}$, with accuracy $\epsilon 2^{-(t+2(q-1))}$, in the sense that $V^{t-1\to t} = U^{t-1\to t}|0^{2n}\rangle_{A_tB_t}$, with A^* playing the role of A_t. Subsequently, apply a phase -1 on the all-zero state of the ancilla registers A_t and B_t, and reapply $U^{t-1\to t}$, this time with accuracy $\epsilon 2^{-(t+2(q-1)+1)}$.*

We omitted defining $\mathfrak{ES}.\mathsf{Ver}$ since it is trivial to build from $\mathsf{CReflect}$, as described in Lemma 5. By Theorem 6, the runtime of $\mathfrak{ES}(\epsilon, n)$ is polynomial in n, $\log(1/\epsilon)$ and the total number of queries q that are made to its various interfaces.

We want to show that the above sampler is indistinguishable from the ideal sampler to any oracle algorithm, in the following sense. Given a stateful machine $\mathcal{C} \in \{\mathfrak{IS}(n), \mathfrak{ES}(n, \epsilon)\}$ and a (not necessarily efficient) oracle algorithm $\mathcal{A}$, we define the process $b \leftarrow \mathcal{A}^{\mathcal{C}}$ as follows:

1. $\mathcal{C}.\mathsf{Init}$ is called;
2. $\mathcal{A}$ receives oracle access to $\mathcal{C}.\mathsf{Gen}$ and $\mathcal{C}.\mathsf{CReflect}$;
3. $\mathcal{A}$ outputs a bit b.

Theorem 7. *For all oracle algorithms $\mathcal{A}$ and all $\epsilon > 0$ that can depend on n in an arbitrary way,*

$$\left|\Pr\left[\mathcal{A}^{\mathfrak{IS}(n)} = 1\right] - \Pr\left[\mathcal{A}^{\mathfrak{ES}(n,\epsilon)} = 1\right]\right| \leq \epsilon\,. \tag{19}$$

Proof. During the execution of $\mathfrak{ES}(\epsilon, n)$, the i-th call of $V^{t-1\to t}$ (for any t) incurs a trace distance error of at most $\epsilon 2^{-i}$. The trace distance between the outputs of $\mathcal{A}^{\mathfrak{ES}}(\epsilon, n)$ and $\mathcal{A}^{\mathfrak{ES}}(0, n)$ is therefore bounded by $\sum_{i=1}^{\infty} \epsilon 2^{-i} = \epsilon$. It is thus sufficient to establish the theorem for $\mathfrak{ES}(0, n)$.

For any fixed q, there exists a stateful machine $\hat{\mathfrak{ES}}(0, q, n)$ which is perfectly indistinguishable from $\mathfrak{IS}(n)$ to all adversaries who make a maximum total number q of queries. The Init procedure of $\hat{\mathfrak{ES}}(0, q, n)$ samples a random element U_i from an exact unitary $2q$-design $D^{2q} = \{U_i\}_{i\in I}$. Queries to Gen are answered with a copy of $U_i|0\rangle$, and $\mathsf{Reflect}$ is implemented by applying $\mathbb{1} - 2U_i|0\rangle\langle 0|U_i^\dagger$. It will be helpful to express $\hat{\mathfrak{ES}}(0, q, n)$ in an equivalent isometric form. In this form, the initial oracle state is $|\eta\rangle = |I|^{-1/2}\sum_{i\in I}|i\rangle_{\hat{B}}$. Gen queries are answered using the $\hat{B}$-controlled isometry

$$\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}} = \sum_{i\in I}|i\rangle\langle i|_{\hat{B}} \otimes U_i|0\rangle_{A_{t+1}}\,. \tag{20}$$

$\mathsf{Reflect}$ queries are answered by

$$\hat{V}^{\mathsf{Reflect}}_{\hat{B}A^*\to\hat{B}A^*} = \mathbb{1} - 2\sum_{i\in I}|i\rangle\langle i|_{\hat{B}} \otimes U_i|0\rangle\langle 0|_{A^*}U_i^\dagger \tag{21}$$

$$= \mathbb{1} - 2\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A^*}\left(\hat{V}^{t\to t+1}\right)^\dagger_{\hat{B}A^*\to\hat{B}}\,. \tag{22}$$

Now suppose $\mathcal{A}$ is an arbitrary (i.e., not bounded-query) algorithm making only Gen queries. We will show that after q queries, the oracles $\mathfrak{ES}(0, n)$ and $\hat{\mathfrak{ES}}(0, q, n)$ are equivalent, and that this holds for all q. We emphasize that $\mathfrak{ES}(0, n)$ does not depend on q; we can thus apply the equivalence for the appropriate total query count q_{total} after $\mathcal{A}$ has produced its final state, even if q_{total} is determined only at runtime. It will follow that $\mathfrak{ES}(0, n)$ is equivalent to $\mathfrak{IS}(n)$.

To show the equivalence betwen $\mathfrak{ES}(0, n)$ and $\hat{\mathfrak{ES}}(0, q, n)$, we will demonstrate a partial isometry $V^{\mathrm{switch},t}$ that transforms registers B^t of $\mathfrak{ES}(0, n)$ (after t Gen queries and no $\mathsf{Reflect}$ queries) into the register $\hat{B}$ of $\hat{\mathfrak{ES}}(0, q, n)$, in such a way that the corresponding global states on A^tB^t and $A^t\hat{B}$ are mapped to each other. The isometry is partial because its domain is the symmetric subspace of $\mathbb{C}^{2^n\otimes t}$. It is defined as follows:

$$V^{\mathrm{switch},t}_{B^t\to\hat{B}} = \sqrt{\frac{d_{\mathrm{Sym}^t\mathbb{C}^d 2^n}}{|I|}}\sum_{i\in I}\left(\langle 0|U_i^T\right)^{\otimes t}_{B^t} \otimes |i\rangle_{\hat{B}}\,. \tag{23}$$

To verify that this is indeed the desired isometry, we calculate:

$$\left(\langle 0|U_i^T\right)_{B^t}^{\otimes t}|\phi^+_{\mathrm{Sym}}\rangle_{A^tB^t} = \sqrt{\frac{2^{nt}}{d_{\mathrm{Sym}^t\mathbb{C}^{2^n}}}}\left(\langle 0|U_i^T\right)_{B^t}^{\otimes t}\Pi^{\mathrm{Sym}}_{B^t}|\phi^+\rangle_{A^tB^t} \tag{24}$$

$$= \sqrt{\frac{2^{nt}}{d_{\mathrm{Sym}^t\mathbb{C}^{2^n}}}}\left(\langle 0|U_i^T\right)_{B^t}^{\otimes t}|\phi^+\rangle_{A^tB^t} \tag{25}$$

$$= \sqrt{\frac{2^{nt}}{d_{\mathrm{Sym}^t\mathbb{C}^{2^n}}}}\left(\langle 0|\right)_{B^t}^{\otimes t}\otimes\left(U_i\right)_{A^t}^{\otimes t}|\phi^+\rangle_{A^tB^t} \tag{26}$$

$$= \sqrt{\frac{1}{d_{\mathrm{Sym}^t\mathbb{C}^{2^n}}}}\left(U_i|0\rangle\right)_{A^t}^{\otimes t}. \tag{27}$$

Here we have used the fact that $\left(\langle 0|U_i^T\right)^{\otimes t}$ is in the symmetric subspace in the second equality, and the third and forth equality are applications of the Mirror Lemma (Lemma 1) with $d = d' = 2^{nt}$, and $d = 1,\ d' = 2^{nt}$, respectively.

We have hence proven the exact correctness of $\mathfrak{ES}(0,n)$ without the Reflect interface. Note that the global state after t queries to $\mathfrak{ES}(0,n).\mathsf{Gen}$ is the maximally entangled state of two copies of the t-fold symmetric subspace; of course, this is only true up to actions performed by the adversary, but those trivially commute with maps applied only to the oracle registers. As the global state is in the domain of $V^{\mathrm{switch},t}_{B^t\to\hat{B}}$, we obtain the equation

$$\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}}V^{\mathrm{switch},t}_{B^t\to\hat{B}} = V^{\mathrm{switch},t+1}_{B^{t+1}\to\hat{B}}V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}}. \tag{28}$$

More precisely, we observe that the two sides of the above have the same effect on the global state, and then conclude that they must be the same operator by the Choi-Jamoiłkowski isomorphism.

Recalling that $V^{\mathrm{switch},t}$ is partial with the symmetric subspace as its domain, we see that Equation (28) is equivalent to

$$\left(V^{\mathrm{switch},t+1}_{B^{t+1}\to\hat{B}}\right)^\dagger\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}}V^{\mathrm{switch},t}_{B^t\to\hat{B}} = \Pi^{\mathrm{Sym}^{t+1}\mathbb{C}^{2^n}}_{B^{t+1}}V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}} \tag{29}$$

$$= V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}}\Pi^{\mathrm{Sym}^t\mathbb{C}^{2^n}}_{B^t}. \tag{30}$$

By taking the above equality times its adjoint, we arrive at

$$\left(V^{\mathrm{switch},t}_{B^t\to\hat{B}}\right)^\dagger\left(\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}}\right)^\dagger V^{\mathrm{switch},t+1}_{B^{t+1}\to\hat{B}}\left(V^{\mathrm{switch},t+1}_{B^{t+1}\to\hat{B}}\right)^\dagger\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}}V^{\mathrm{switch},t}_{B^t\to\hat{B}}$$
$$= \Pi^{\mathrm{Sym}^t\mathbb{C}^{2^n}}_{B^t}\left(V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}}\right)^\dagger V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}}\Pi^{\mathrm{Sym}^t\mathbb{C}^{2^n}}_{B^t}. \tag{31}$$

By Equation (28), the range of $\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}} V^{\mathrm{switch},t}_{B^t\to\hat{B}}$ is contained in the range of $V^{\mathrm{switch},t+1}_{B^{t+1}\to\hat{B}} \otimes \mathbb{1}_{A_{t+1}}$. We can thus simplify as follows:

$$\begin{aligned}\left(V^{\mathrm{switch},t}_{B^t\to\hat{B}}\right)^\dagger \left(\hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}}\right)^\dagger \hat{V}^{t\to t+1}_{\hat{B}\to\hat{B}A_{t+1}} V^{\mathrm{switch},t}_{B^t\to\hat{B}}&\\ =\Pi^{\mathrm{Sym}^t\mathbb{C}^{2^n}}_{B^t} \left(V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}}\right)^\dagger V^{t\to t+1}_{B^t\to B^{t+1}A_{t+1}} \Pi^{\mathrm{Sym}^t\mathbb{C}^{2^n}}_{B^t}&. \qquad (32)\end{aligned}$$

Now observe that both sides of the above consist of a projection operator "sandwiched" by some operation. These two projection operators are precisely the projectors which define the reflection operators of $\mathfrak{E}\hat{\mathfrak{S}}(0,q,n)$ (on the left-hand side) and $\mathfrak{ES}(0,n)$ (on the right-hand side.) We thus see that Equation (32) shows that applying $\mathfrak{ES}(0,n).\mathsf{Reflect}$ is the same as switching to $\mathfrak{E}\hat{\mathfrak{S}}(0,q,n)$, applying $\mathfrak{E}\hat{\mathfrak{S}}(0,q,n).\mathsf{Reflect}$, and then switching back to $\mathfrak{ES}(0,n)$. The same holds for the controlled versions $\mathfrak{ES}(0,n).\mathsf{CReflect}$ and $\mathfrak{E}\hat{\mathfrak{S}}(0,n).\mathsf{CReflect}$.

This completes the proof of the exact equality between the stateful machines $\mathfrak{IS}(n)$ and $\mathfrak{ES}(0,n)$. The approximate case follows as argued above. □

It turns out that if we have an a priori bound of the form $q = O(\sqrt{2^n\epsilon})$ on the number of queries that will be made to our state sampler, in relation to the number of qubits n and the desired accuracy ϵ, there is also an alternative protocol, due to Zvika Brakerski. The approach is based on Zhandry's compressed oracle technique and the work by Ji, Liu and Song. In [18] and in [10] one can find the following theorem for the two mentioned phase variants, respectively.

Theorem 8 (Lemma 1 in [18], respectively Theorem 1.2 in [10]). *Let $H : \{0,1\}^n \to \{0,1\}^n$ be a random function. Then k copies of the n-qubit quantum state*

$$|\psi^H\rangle = 2^{-n/2} \sum_{x\in\{0,1\}^n} \omega^{H(x)}|x\rangle \qquad (33)$$

are statistically indistinguishable from k copies of a Haar random quantum state up to error $O(k^2/2^n)$, for $\omega = e^{\frac{2\pi i}{2^n}}$, respectively $\omega = -1$.

Let now $\mathfrak{EF}(n,n)$ be the stateful machine with interfaces Init and Query simulating a random function from n bits to n bits that was given in [31]. Then we get the following

Corollary 3. *Let $\mathfrak{ES}'(n)$ be the following stateful machine:*

- $\mathfrak{ES}'(n).\mathsf{Init}$ *is equal to* $\mathfrak{EF}(n,n).\mathsf{Init}$.
- $\mathfrak{ES}'(n).\mathsf{Gen}$ *produces a copy of $|\psi^H\rangle$, simulating H using a single query to* $\mathfrak{EF}(n,n).\mathsf{Query}$.
- $\mathfrak{ES}'(n).\mathbb{C}\mathsf{CReflect}$ *implements the controlled reflection about $|\psi^H\rangle$, simulating H using two queries to* $\mathfrak{EF}(n,n).\mathsf{Query}$.

For all oracle algorithms $\mathcal{A}$ making q that make q queries and that can depend on n in an arbitrary way,

$$\left|\Pr\left[\mathcal{A}^{\mathfrak{IS}(n)}=1\right] - \Pr\left[\mathcal{A}^{\mathfrak{ES}'(n)}=1\right]\right| \le O(q^2/2^n)\,. \qquad (34)$$

4 Simulating a Haar-random unitary oracle

4.1 The problem, and our approach

We begin by defining the ideal object we'd like to emulate. This ideal object samples a Haar-random unitary U, and then answers two types of queries: queries to U, and queries to its inverse $U^\dagger$.

Construction 3 (Ideal unitary sampler) *Let n be a positive integer. The ideal unitary sampler is an ideal machine $\mathfrak{IU}(n)$ with interfaces* $(\mathsf{Init}, \mathsf{Eval}, \mathsf{Invert})$*, defined as follows.*

1. $\mathfrak{IU}(n).\mathsf{Init}$: *takes no input; samples a description $\tilde{U}$ of a Haar-random n-qubit unitary operator U.*
2. $\mathfrak{IU}(n).\mathsf{Eval}$: *takes n-qubit register as input, applies U and responds with the output;*
3. $\mathfrak{IU}(n).\mathsf{Invert}$: *takes n-qubit register as input, applies U^{-1} and responds with the output.*

Below, we construct a stateful machine that runs in polynomial *space* (and the runtime of which we don't characterize), and that is indistinguishable from $\mathfrak{IU}(n)$ for arbitrary query algorithms.

Our approach. It turns out that the solution of a much easier task comes to our help, namely simulating a Haar random unitary for an algorithm that makes an *a priori* polynomially bounded number t of queries. In this case we can just pick a unitary t-design, sample an element from it and answer the up to t queries using this element. As in the proof of Theorem 7, we can also construct an isometric stateful machine version of this strategy: Instead of sampling a random element from the t-design, we can prepare a quantum register in a superposition, e.g. over the index set of the t-design (Init), and then apply the t-design element (Eval) or its inverse (Invert) controlled on that register.

Now consider an algorithm that makes t parallel queries to a Haar random unitary (for ease of exposition let us assume here that the algorithm makes no inverse queries). The effect of these t parallel queries is just the application of the t-twirling channel (or the mixed twirling channel defined in Equation (10)) to the t input registers. The t-design-based isometric stateful machine simulates this t-twirling channel faithfully. What is more, it applies a Stinespring dilation of the t-twirling channel, the dilating register being the one created by initialization.

Now suppose we have answered t queries using the t-design-based machine, and are now asked to answer another, still parallel, query. Of course we cannot, in general, just answer it using the t-design, as its guarantees only hold for t applications of the unitary. But all Stinespring dilations of a quantum channel are equivalent in the sense that there exists a (possibly partial) isometry acting on the dilating register of one given dilation, that transforms it into another given dilation. So we can just apply an isometry that transforms our t-design based Stinespring dilation into a $t+1$-design based one, and subsequently answer the $t+1$st query using a controlled unitary.

4.2 Construction and proof

We continue to describe a stateful machine that simulates $\mathfrak{IU}(n)$ exactly and has a state register of size polynomial in n and the total number of queries q that an algorithm makes to its Eval and Invert interfaces. The existence of the required unitary t-designs is due to Corollary 2.

We recall our conventions for dealing with many copies of fixed-sized registers. We let A denote an n-qubit register, we let A_j denote indexed copies of A, and we let A^t denote $A_1 A_2 \cdots A_t$. In this case, the various copies of A will be the input registers of the adversary, on which the simulator will act. The oracle will now hold a single register $\hat{B}_t$ whose size will grow with the number of queries t. This register holds an index of an element in a t-design.

For the construction below, we need the following quantum states and operators. For a positive integer n, choose a family of n-qubit unitary designs $\{D_t\}_{t\in\mathbb{N}}$, where $D_t = \{U_{t,i}\}_{i\in I_t}$ is a unitary t-design. Let $\hat{B}_t$ be a register of dimension $|I_t|$ and define the uniform superposition over indices

$$|\eta_t\rangle_{\hat{B}_t} = \frac{1}{\sqrt{|I_t|}} \sum_{i\in I_t} |i\rangle_{\hat{B}_t}. \tag{35}$$

For nonnegative integers t, t', ℓ, define the unitaries

$$V^{(t,t',\ell)}_{A^{t'}\hat{B}_t} = \sum_{i\in I_t} (U_{t,i})^{\otimes \ell}_{A_1 A_2 \ldots A_\ell} \otimes \left(U^\dagger_{t,i}\right)^{\otimes t'-\ell}_{A_{\ell+1} A_{\ell+2} \ldots A_{t'}} \otimes |i\rangle\langle i|_{\hat{B}_t}. \tag{36}$$

These isometries perform the following: controlled on an index i of a t-design $U_{t,i}$, apply $U_{t,i}$ to ℓ registers and $U^\dagger_{t,i}$ to $t'-\ell$ registers. For us it will always be the case that $t' \leq t$, since otherwise the t-design property no longer makes the desired guarantees on the map V.

We also let $W^{(t,\ell)}_{\hat{B}_t \to \hat{B}_{t+1}}$ be an isometry such that

$$V^{(t+1,t,\ell)}_{A^t \hat{B}_{t+1}} |\eta_{t+1}\rangle_{\hat{B}_{t+1}} = W_{\hat{B}_t \to \hat{B}_{t+1}} V^{(t,t,\ell)}_{A^t \hat{B}_t} |\eta_t\rangle_{\hat{B}_t} \tag{37}$$

for $\ell = 0, ..., t$. The isometry W always exists, as all Stinespring dilations are isometrically equivalent, and both $V^{(t,t,\ell)}_{A^t\hat{B}_t}|\eta_t\rangle_{\hat{B}_t}$ and $V^{(t+1,t,\ell)}_{A^t\hat{B}_{t+1}}|\eta_{t+1}\rangle_{\hat{B}_{t+1}}$ are Stinespring dilations of the mixed twirling channel $\mathcal{T}^{(t,\ell)}$ by the t-design property.

We are now ready to define the space-efficient unitary sampler.

Construction 4 (Space-efficient unitary sampler) *Let n be a positive integer and $\{D_t\}_{t\in\mathbb{N}}$ a family of n-qubit unitary t-designs $D_t = \{U_{t,i}\}_{i\in I_t}$, with $|I_t| = 2^{\mathrm{poly}(n,t)}$. Define a stateful machine $\mathfrak{EU}(n,\epsilon)$ with interfaces* (Init, Eval, Invert) *as follows. The machine will maintain counters t_e (the number of* Eval *queries), t_i (the number of* Invert *queries), and $t := t_e + t_i$.*

1. $\mathfrak{EU}(n)$.Init : *Prepares the state $|\eta_1\rangle_{\hat{B}_1}$ and stores it.*
2. $\mathfrak{EU}(n)$.Eval :

- *If $t = 0$, apply $V^{(1,1,1)}_{A_1\hat{B}_1}$, where A_1 is the input register.*
- *If $t > 0$, apply $W^{(t,t_e)}_{\hat{B}_t \to \hat{B}_{t+1}}$ to the state register and subsequently apply $V^{t+1,1,1}_{A_{t+1}\hat{B}_{t+1}}$, where A_{t+1} is the input register.*

3. $\mathfrak{IU}(n).\mathsf{Invert}$:
   - *If $t = 0$, apply $V^{(1,1,0)}_{A_1\hat{B}_1}$, where A_1 is the input register.*
   - *If $t > 0$, apply $W^{(t,t_e)}_{\hat{B}_t \to \hat{B}_{t+1}}$ to the state register and subsequently apply $V^{t+1,1,0}_{A_{t+1}\hat{B}_{t+1}}$, where A_{t+1} is the input register.*

We want to show that the above sampler is indistinguishable from the ideal sampler to any oracle algorithm, in the following sense. Given a stateful machine $\mathcal{C} \in \{\mathfrak{IU}(n), \mathfrak{EU}(n,\epsilon)\}$ and a (not necessarily efficient) oracle algorithm $\mathcal{A}$, we define the process $b \leftarrow \mathcal{A}^{\mathcal{C}}$ as follows:

1. $\mathcal{C}.\mathsf{Init}$ is called;
2. $\mathcal{A}$ receives oracle access to $\mathcal{C}.\mathsf{Eval}$ and $\mathcal{C}.\mathsf{Invert}$;
3. $\mathcal{A}$ outputs a bit b.

Theorem 9. *For all oracle algorithms $\mathcal{A}$*

$$\Pr\left[\mathcal{A}^{\mathfrak{IU}(n)} = 1\right] = \Pr\left[\mathcal{A}^{\mathfrak{EU}(n,\epsilon)} = 1\right]. \tag{38}$$

Proof. We begin by proving the following claim by induction. The claim states that the theorem holds for adversaries who only make parallel queries.

Claim. For all $x \in \{0,1\}^t$, let $V^{(x)}_{A^t \to A^t\hat{B}_t}$ be the isometry that is implemented by making t parallel queries to $\mathfrak{EU}(n,\epsilon)$, where the i-th query is made to the Eval interface if $x_i = 1$ and to the Invert interface if $x_i = 0$. Let further $\sigma \in S_t$ be a permutation such that $\sigma.x = 11...100...0$, where the lower dot denotes the natural action of S_t on strings of length t. Then

$$V^{(x)}_{A^t \to A^t\hat{B}_t} = \sigma^{-1}_{A^t} V^{(t,t,\ell)}_{A^t\hat{B}_t} |\eta_t\rangle_{\hat{B}_t}, \tag{39}$$

where σ acts by permuting the t registers.

Proof. For $t = 1$, the claim trivially holds. Now suppose the claim holds for $t-1$. By definition of the Eval and Invert interfaces,

$$V^{(x)}_{A^t \to A^t\hat{B}_t} = V^{t,1,x_t}_{A_t\hat{B}_t} W^{(t,\ell)}_{\hat{B}_{t-1} \to \hat{B}_t} V^{(x_{[1;t-1]})}_{A^{t-1} \to A^{t-1}\hat{B}_{t-1}}, \tag{40}$$

where $x_{[a,b]} = x_a x_{a+1} ... x_b$. By the induction hypothesis, we have

$$V^{(x_{[1;t-1]})}_{A^{t-1} \to A^{t-1}\hat{B}_{t-1}} = \hat{\sigma}^{-1}_{A^{t-1}} V^{(t-1,t-1,\ell-x_t)}_{A^{t-1}\hat{B}_{t-1}} |\eta_{t-1}\rangle_{\hat{B}_{t-1}} \tag{41}$$

for an appropriate permutation $\hat{\sigma} \in S_{t-1}$. By the design property of D_j for $j = t, t-1$ and the definition of $W^{(t,\ell)}$ we obtain

$$\begin{aligned}
& \mathcal{T}_{D_{t-1}}^{(t-1,\ell-x_t)} = \mathcal{T}_{D_t}^{(t-1,\ell-x_t)} \\
\Leftrightarrow \quad & W_{\hat{B}_{t-1}\to\hat{B}_t}^{(t-1,\ell)} V_{A^{t-1}\hat{B}_{t-1}}^{(t-1,t-1,\ell-x_t)} |\eta_{t-1}\rangle_{\hat{B}_{t-1}} = V_{A^{t-1}\hat{B}_t}^{(t,t-1,\ell-x_t)} |\eta_{t-1}\rangle_{\hat{B}_t} \\
\Leftrightarrow \quad & W_{\hat{B}_{t-1}\to\hat{B}_t}^{(t,\ell)} \hat{\sigma}_{A^{t-1}}^{-1} V_{A^{t-1}\hat{B}_{t-1}}^{(t-1,t-1,\ell-x_t)} |\eta_{t-1}\rangle_{\hat{B}_{t-1}} = \hat{\sigma}_{A^{t-1}}^{-1} V_{A^{t-1}\hat{B}_t}^{(t,t-1,\ell-x_t)} |\eta_{t-1}\rangle_{\hat{B}_t}.
\end{aligned} \tag{42}$$

Here we have used the fact that the permutation and $W^{(t-1,\ell)}$ commute because they act on disjoint sets of registers. Putting Equations (40), (41) and (42) together, it follows that

$$V_{A^t\to A^t\hat{B}_t}^{(x)} = V_{A_t\hat{B}_t}^{t,1,x_t} \hat{\sigma}_{A^{t-1}}^{-1} V_{A^{t-1}\hat{B}_t}^{(t,t-1,\ell-x_t)} |\eta_t\rangle_{\hat{B}_t}. \tag{43}$$

But clearly

$$V_{A_t\hat{B}_t}^{t,1,x_t} \hat{\sigma}_{A^{t-1}}^{-1} V_{A^{t-1}\hat{B}_t}^{(t,t-1,\ell-x_t)} = \sigma_{A^t}^{-1} V_{A^t\hat{B}_t}^{(t,t,\ell)} \tag{44}$$

For an appropriate permutation σ that consists of applying $\hat{\sigma}$ and then sorting in x_t correctly.

The generalization to adaptive algorithms is done via *post-selection*: Given an algorithm $\mathcal{A}$ with some oracles $O_1, O_2, ..., O_k$, consider non-adaptive algorithm $\tilde{\mathcal{A}}$ that first queries the oracles a sufficient number of times, each of the queries being made with the first half of a maximally entangled state as input. Subsequently the adaptive adversary is run, answering the queries by performing the sender's part of the standard quantum teleportation with the input playing the role of the state to be teleported, and the second half of one of the maximally entangled states playing the role of the sender's half of the entangled resource state for teleportation. Conditioned on the event that all the Pauli corrections in all the teleportation protocols are equal to the identity, the output of $\tilde{\mathcal{A}}$ is equal to the output of $\mathcal{A}$.

Now consider the case where $k = 2$ and O_1 and O_2 are the Eval and Invert interfaces of $\mathfrak{EU}(n, 0)$, or $\mathfrak{IU}(n)$. As the output of $\tilde{\mathcal{A}}$ is *exactly* the same in the two cases, the same holds for the version of $\tilde{\mathcal{A}}$ where we condition, on the outcome that all the Pauli corrections in all the teleportation protocols are equal to the identity, which proves the theorem. □

Using Corollary 2 and the above, we get the following upper bound on the space complexity of lazy sampling Haar random unitaries.

Corollary 4. *The space complexity S of simulating $\mathfrak{IU}(n)$ as a function of n and the number of queries q is bounded from above by the logarithm of number of elements in any family of exact n-qubit unitary q-designs, and hence*

$$S(n, q) \leq 2q(2n + \log e) + O(\log q)\,. \tag{45}$$

Proof. According to Corollary 2, There exists an exact unitary q-design such that $2q \log\left(\frac{e(2^{2n}+q-1)}{q}\right) \leq 2q(2n + \log e)$ qubits suffice to coherently store the index of an element from it. The only additional information that $\mathfrak{EU}(n)$ needs to store is how many direct and inverse queries have been answered, which can be done using $\log q$ bits.

□

Our results suggest two possible approaches to devise a time-efficient lazy sampler for Haar random unitaries. The most promising one is to use the same approach as for the state sampler and explicitly constructing the update isometry, possibly using explicit bases for the irreducible representations of $U(2^n)$, or using the Schur transform [6]. The other one would be to use the t-design update method described above, but using efficient approximate t-designs, e.g. the ones constructed in [11]. This would, however, likely require a generalization of the Stinespring dilation continuity result from [20] to so-called quantum combs [12]. In addition, we would need to show that the transition isometries, i.e. the approximate analogue of the isometries $W^{(t,\ell)}$ from Construction 4, are efficiently implementable. We leave the exploration of these approaches for future work.

5 Application: untraceable quantum money

5.1 Untraceable quantum money

Our definition of quantum money deviates somewhat from others in the literature [1,18]. We allow the bank to maintain an internal quantum register, we do not require that the money states are pure, and we allow adversaries to apply arbitrary (i.e., not necessarily efficiently implementable) channels.

Definition 2 (Quantum money). *A quantum money scheme is a family of stateful machines $\mathfrak{M}$ indexed by a security parameter λ, and having two interfaces:*

1. Mint*: receives no input, outputs an n-qubit register;*
2. Ver*: receives an n-qubit register as input, outputs an n-qubit register together with a flag* $\{\mathsf{acc}, \mathsf{rej}\}$*,*

satisfying the following two properties:

- *correctness:* $\|\mathsf{Ver} \circ \mathsf{Mint} - \mathbb{1} \otimes |\mathsf{acc}\rangle\langle\mathsf{acc}|\| \leq \mathsf{negl}(\lambda)$*;*[8]
- *unforgeability: for all channels Λ with oracle, and all $k \geq 0$,*

$$\Pr\left[\mathsf{acc}^{k+1} \leftarrow {}_{\text{flag}}|\mathsf{Ver}^{\otimes k+1} \circ \Lambda^{\mathsf{Ver}} \circ \mathsf{Mint}^{\otimes k}\right] \leq \mathsf{negl}(\lambda)\,,$$

where ${}_{\text{flag}}|$ denotes discarding all registers except Ver *flags.*

[8] Note that it is understood that this inequality should hold no matter which interfaces have been called in between the relevant Mint and Ver calls

It is implicit in the definition that n is a fixed polynomial function of λ, and that all relevant algorithms are uniform in λ.

Next, we define untraceability for quantum money schemes.

Definition 3 (Untraceability game). *The untraceability game* $\mathsf{Untrace}_\lambda[\mathfrak{M}, \mathcal{A}]$ *between an adversary* $\mathcal{A}$ *and a quantum money scheme* $\mathfrak{M}$ *at security parameter* λ *proceeds as follows:*

1. *set up the trace:* $\mathcal{A}(1^\lambda)$ *receives oracle access to* Ver *and* Mint*, and outputs registers* M_1*,* M_2*,* $\ldots$*,* M_k *and a permutation* $\pi \in S_k$*;*
2. *permute and verify bills:* $b \leftarrow \{0,1\}$ *is sampled, and if* $b = 1$ *the registers* $M_1 \cdots M_k$ *are permuted by* π*.* Ver *is invoked on each* M_j*; the accepted registers are placed in a set* $\mathcal{M}$ *while the rest are discarded;*
3. *complete the trace:* $\mathcal{A}$ *receives* $\mathcal{M}$ *and the entire internal state of* $\mathfrak{M}$*, and outputs a guess* $b' \in \{0,1\}$*.*

The output of $\mathsf{Untrace}_\lambda[\mathfrak{M}, \mathcal{A}]$ *is* $\delta_{bb'}$*; in the case* $b = b'$*, we say that* $\mathcal{A}$ *wins.*

Definition 4 (Untraceable quantum money). *A quantum money scheme* $\mathfrak{M}$ *is untraceable if, for every algorithm* $\mathcal{A}$*,*

$$\Pr\left[1 \leftarrow \mathsf{Untrace}_\lambda[\mathfrak{M}, \mathcal{A}]\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda)\,.$$

The intuition behind the definition is as follows. In general, one might consider a complicated scenario involving many honest players and many adversaries, where the goal of the adversaries is to trace the movement of at least one bill in transactions involving at least one honest player. Tracing in transactions involving only adversaries is of course trivial. The first natural simplification is to view all the adversaries as a single adversarial party; if that party cannot trace, then neither can any individual adversary. Next, we assume that honest players will verify any bills they receive immediately; obviously, if they do not do this, and then participate in transactions with the adversary, then tracing is again trivial. We thus arrive at the situation described in the game: the adversary is first allowed to create candidate bills arbitrarily, including storing information about them and entangling them with additional registers, before handing them to honest players who may or may not perform some transactions; the goal of the adversary is to decide which is the case, with the help of the bank. Note that one round of this experiment is sufficient in the security game, as an adversary can always use the Ver and Mint oracles to simulate additional rounds.

One might reasonably ask if there are even stronger definitions of untraceability than the above. Given its relationship to the ideal state sampler, we believe that Haar money, defined below, should satisfy almost any notion of untraceability, including composable notions. We also remark that, based on the structure of the state simulator, which maintains an overall pure state supported on two copies of the symmetric subspace of banknote registers, it is straightforward to see that the scheme is also secure against an "honest but curious" or "specious" [26,15] bank. We leave the formalization of these added security guarantees to future work.

5.2 Haar money

Next, we show how the lazy state sampler (Construction 2) yields untraceable quantum money. The construction follows the idea of [18] sample a single (pseudo)random quantum state and hand out copies of it as banknotes.

Construction 5 (Haar money) *Let n be a positive integer and $\epsilon > 0$. The Haar scheme $\mathfrak{HM}(n, \epsilon)$ is defined as follows:*

- Mint*: on first invocation, instantiate* $\mathfrak{ES} := \mathfrak{ES}(n, \epsilon)$ *by running* $\mathfrak{ES}.\mathsf{Init}$*. On all invocations, output result of* $\mathfrak{ES}.\mathsf{Gen}$*;*
- Ver*: apply* $\mathfrak{ES}.\mathsf{Ver}$*; in the* acc *case, call* Mint *and output the result; in the* rej *case, output* 0^n*.*

We remark that, while Construction 2 does not explicitly include a Ver interface, one can easily be added by Lemma 5.

Proposition 3. *Haar money is an untraceable quantum money scheme.*

Proof. We need to show three properties: completeness, unforgeability, and untraceability. For the completeness and unforgeability properties, observe that Theorem 7 implies that the adversary's view is indistinguishable (up to negligible terms) if we replace the efficient state sampler $\mathfrak{ES}$ with the ideal $\mathfrak{IS}$. Once we've made that replacement, completeness follows from the definition of $\mathfrak{IS}.\mathsf{Gen}$ and $\mathfrak{IS}.\mathsf{Ver}$, and unforgeability follows from the complexity-theoretic no-cloning theorem [1].

For untraceability, it is of course true that $\mathfrak{IS}$ is obviously untraceable. However, we cannot simply invoke Theorem 7 to conclude the same about $\mathfrak{ES}$, since the adversary will receive the state of the bank at the end of the game. Instead, we argue as follows. Consider step 2 (permute and verify bills) in the untraceability game $\mathsf{Untrace}_\lambda[\mathfrak{HM}, \mathcal{A}]$. An equivalent way to perform this step is to (i.) verify all the registers first, (ii.) discard the ones that fail verification, and then (iii.) apply the permutation, conditioned on the challenge bit b. Steps (i.) and (ii.) are applied always and in particular do not depend on b. However, after (i.) and (ii.) have been applied, by the definition of $\mathfrak{ES}$ the joint state of the bank and all the $M_j \in \mathcal{M}$ (and indeed all verified bills in existence) is negligibly far from the state $|\phi^+_{\mathrm{Sym}}\rangle$, i.e., the maximally entangled state on the symmetric subspace. This state is clearly invariant under permutation of the money registers, and in particular under the permutation of the registers in $\mathcal{M}$ selected by the adversary. We emphasize that this invariance holds for the entire state (including the bank.) As the remainder of the game experiment is simply some channel applied to that state, and this channel does not depend on b, the result follows. □

While Haar money is an information-theoretically unforgeable and untraceable quantum money scheme, it is easy to see that the quantum money scheme devised in [18] is *computationally* unforgeable and untraceable.

References

1. Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 41–60. ACM, 2012.
2. Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Can you sign a quantum state. Cryptology ePrint Archive, Report 2018/1164, 2018. `https://eprint.iacr.org/2018/1164`.
3. Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Unforgeable quantum encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 489–519, Cham, 2018. Springer International Publishing.
4. Gorjan Alagic, Christian Majenz, and Alexander Russell. Efficient simulation of random states and random unitaries. *arXiv preprint arXiv:1910.05729*, 2019.
5. Andris Ambainis and Joseph Emerson. Quantum t-designs: T-wise independence in the quantum world. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity*, CCC '07, pages 129–140, Washington, DC, USA, 2007. IEEE Computer Society.
6. Dave Bacon, Isaac L. Chuang, and Aram W. Harrow. Efficient quantum circuits for schur and clebsch-gordan transforms. *Phys. Rev. Lett.*, 97:170502, Oct 2006.
7. Mihir Bellare and Phillip Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and communications security*, CCS '93, pages 62–73, New York, NY, USA, 1993. ACM.
8. Charles Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
9. Dan Boneh and Mark Zhandry. Quantum-secure message authentication codes. In *Advances in Cryptology - EUROCRYPT 2013*, pages 592–608. Springer, 2013.
10. Zvika Brakerski and Omri Shmueli. (pseudo) random quantum states with binary phase. *arXiv preprint arXiv:1906.10611*, 2019.
11. Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, 2016.
12. G. Chiribella, G. M. D'Ariano, and P. Perinotti. Quantum circuit architecture. *Phys. Rev. Lett.*, 101:060401, Aug 2008.
13. Matthias Christandl. *The structure of bipartite quantum states-Insights from group theory and cryptography*. PhD thesis, University of Cambridge, 2006.
14. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 356–383, Cham, 2019. Springer International Publishing.
15. Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In *Advances in Cryptology–CRYPTO 2010*, pages 685–706. Springer, 2010.
16. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.
17. Aram W. Harrow. The Church of the Symmetric Subspace. *arXiv e-prints*, page arXiv:1308.6595, Aug 2013.

18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 126–152, Cham, 2018. Springer International Publishing.
19. Daniel Kane. Small designs for path-connected spaces and path-connected homogeneous spaces. *Transactions of the American Mathematical Society*, 367(9):6387–6414, 2015.
20. Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner. The information-disturbance tradeoff and the continuity of stinespring's representation. *IEEE transactions on information theory*, 54(4):1708–1717, 2008.
21. Richard A Low. Pseudo-randomness and learning in quantum computation. *arXiv preprint arXiv:1006.5227*, 2010.
22. Christian Majenz. Entropy in Quantum Information Theory – Communication and Cryptography. *arXiv e-prints*, page arXiv:1810.10436, Oct 2018.
23. Dominic Mayers and Andrew Yao. Self testing quantum apparatus. *Quantum Info. Comput.*, 4(4):273–286, July 2004.
24. Michele Mosca and Douglas Stebila. Quantum coins. *Error-Correcting Codes, Finite Geometries and Cryptography*, 523:35–47, 2010.
25. Aidan Roy and A. J. Scott. Unitary designs and codes. *Designs, Codes and Cryptography*, 53(1):13–31, Oct 2009.
26. Louis Salvail, Christian Schaffner, and Miroslava Sotáková. On the power of two-party quantum cryptography. In *Proceedings of the 15th International Conference on the Theory and Application of Cryptology and Information Security: Advances in Cryptology*, ASIACRYPT '09, pages 70–87, Berlin, Heidelberg, 2009. Springer-Verlag.
27. W Forrest Stinespring. Positive functions on c*-algebras. *Proceedings of the American Mathematical Society*, 6(2):211–216, 1955.
28. Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, pages 135–152, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
29. John Watrous. Zero-Knowledge against Quantum Attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009.
30. Mark N. Wegman and J.Lawrence Carter. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, 22(3):265 – 279, 1981.
31. Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 239–268, Cham, 2019. Springer International Publishing.