Unitary Complexity and the Uhlmann Transformation Problem

John Bostanci
Department of Computer Science, Columbia University, New York, NY, USA

Yuval Efron
Department of Computer Science, Columbia University, New York, NY, USA

Tony Metger
Institute for Theoretical Physics, ETH Zurich, Switzerland

Alexander Poremba
Department of Computer Science and Department of Physics, Boston University, MA, USA

Luowen Qian
CIS Lab, NTT Research, Sunnyvale, CA, USA

Henry Yuen
Department of Computer Science, Columbia University, New York, NY, USA

Abstract

State transformation problems such as compressing quantum information or breaking quantum commitments are fundamental quantum tasks. However, their computational difficulty cannot easily be characterized using traditional complexity theory, which focuses on tasks with classical inputs and outputs.

To study the complexity of such state transformation tasks, we introduce a framework for *unitary synthesis problems*, including notions of reductions and unitary complexity classes. We use this framework to study the complexity of transforming one entangled state into another via local operations. We formalize this as the *Uhlmann Transformation Problem*, an algorithmic version of Uhlmann's theorem. Then, we prove structural results relating the complexity of the Uhlmann Transformation Problem, polynomial space quantum computation, and zero knowledge protocols.

The Uhlmann Transformation Problem allows us to characterize the complexity of a variety of tasks in quantum information processing, including decoding noisy quantum channels, breaking falsifiable quantum cryptographic assumptions, implementing optimal prover strategies in quantum interactive proofs, and decoding the Hawking radiation of black holes. Our framework for unitary complexity thus provides new avenues for studying the computational complexity of many natural quantum information processing tasks.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Uhlmann's theorem, unitary complexity theory

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.24

Related Version *Full Version*: `https://arxiv.org/abs/2306.13073` [13]

Funding JB and HY are supported by AFOSR award FA9550-21-1-0040, NSF CAREER award CCF-2144219, and the Sloan Foundation. TM acknowledges support from SNSF Project Grant No. 200021_188541 and AFOSR-Grant No. FA9550-19-1-0202. AP is partially supported by AFOSR YIP (award number FA9550-16-1-0495), the Institute for Quantum Information and Matter (an NSF Physics Frontiers Center; NSF Grant PHY-1733907) and by a grant from the Simons Foundation (828076, TV). LQ is supported by DARPA under Agreement No. HR00112020023. We thank the Simons Institute for the Theory of Computing, where some of this work was conducted.

Acknowledgements We thank Anurag Anshu, Lijie Chen, Andrea Coladangelo, Sam Gunn, Yunchao Liu, Joe Renes, Renato Renner, and Mehrdad Tahmasbi for helpful discussions. We thank Fred Dupuis for his help with understanding the decoupling results in his thesis. We are especially grateful to William Kretschmer, Fermi Ma, and John Wright for their thorough discussions on the first draft of the paper. We also thank Tomoyuki Morimae for his helpful comments on a preliminary version of this work. We thank anonymous conference reviewers for their helpful feedback.

1 Introduction

Complexity theory studies the resources required to solve computational problems. Quantum complexity has traditionally studied the quantum resources required to solve *classical* computational problems, i.e., problems that have classical inputs and outputs. However, quantum mechanics also introduces a new kind of computational problem: preparing and transforming quantum states. The goal of this paper is to initiate the formal complexity-theoretic study of such *quantum state transformation problems*. To this end, we extend the language of traditional complexity theory to encompass state transformation problems – we call the resulting framework *unitary complexity theory*.

The idea that the complexity of inherently quantum problems cannot easily be reduced to the complexity of classical problems has already been explored in prior works [30, 2, 6]. In recent years, oracle separations [34, 35, 38] have demonstrated that the complexity of breaking certain quantum cryptographic primitives is independent of the complexity of the decisional complexity classes NP or QMA; in other words, even if $\mathsf{P} = \mathsf{NP}$, certain quantum cryptographic primitives could still remain secure. In fact, [38] gives preliminary evidence that the ability to solve *any* decision problem (even undecidable ones!) would not help with breaking quantum cryptography. Unitary complexity theory allows us to re-establish the link between complexity theory and cryptography in the quantum world.

Beyond this cryptographic motivation, unitary complexity allows us to relate the computational resources required for seemingly unrelated state transformation tasks. In this paper, we focus on tasks involving *Uhlmann transformations*. The name stems from Uhlmann's theorem, [55] a fundamental result in quantum information theory that quantifies how well a bipartite pure state $|C\rangle$ can be mapped to another bipartite pure state $|D\rangle$ by only acting on a subsystem: letting ρ and σ denote the reduced density matrices on the first subsystem of $|C\rangle$ and $|D\rangle$, respectively, Uhlmann's theorem states that

$$\mathrm{F}(\rho,\sigma) = \max_{U} |\langle D| \mathbb{1} \otimes U |C\rangle|^2 , \tag{1.1}$$

where $\mathrm{F}(\rho,\sigma)$ denotes the fidelity function and the maximization is over all unitary transformations acting on the second subsystem. We call a unitary U achieving equality in Equation (1.1) an *Uhlmann transformation*.[1]

Uhlmann transformations are ubiquitous in quantum information processing. Some examples include:

- **Quantum Shannon theory.** Quantum Shannon theory is the study of the fundamental limits of quantum communication over noisy and noiseless channels. Protocols for a myriad of tasks such as state redistribution, entanglement distillation, and quantum communication over a noisy quantum channel all require performing Uhlmann transformations [24, 5, 12, 10].

[1] Such Uhlmann transformations are unique only if $|C\rangle, |D\rangle$ have full Schmidt rank.

- **Quantum cryptography.** While it is known that quantum commitment schemes with information-theoretic security are impossible [39, 37], they are possible under computational assumptions. Recent oracle separations suggest that their security can be based on weaker assumptions than what is needed classically and that the existence of inherently quantum cryptographic primitives may be independent from assumptions in traditional complexity [34, 8, 41, 35, 38]. It can be seen from the impossibility results of Mayers–Lo–Chau [39, 37] that the security of a quantum commitment scheme relies on the hardness of performing certain Uhlmann transformations.
- **Quantum gravity.** Attempts to unite quantum mechanics with general relativity have given rise to apparent paradoxes of whether black holes preserve information or not [23]. Recently, physicists have provided intriguing arguments based on *computational complexity* as possible resolutions to these paradoxes [22]. These arguments claim that distilling entanglement from the emitted Hawking radiation of a black hole is computationally infeasible – this can be equivalently phrased as a statement about the hardness of an Uhlmann transformation [22, 15].
- **Quantum complexity theory.** The $\mathsf{QIP} = \mathsf{PSPACE}$ theorem [28] gives a characterization of the power of (single-prover) quantum interactive proofs. Kitaev and Watrous [33] showed that optimal prover strategies in these interactive proofs boil down to applying Uhlmann transformations at each round.

The fact that Uhlmann transformations appear in these various quantum tasks suggests that they might be related. Can we formalize these relationships and show precise reductions, similarly to how e.g. the theory of NP completeness relates disparate classical computational problems? Can we formalize Uhlmann transformations as a computational problem that is, in some sense, *complete* for these various tasks, similarly to how 3-SAT provides a simple complete problem that elegantly captures the hardness of NP-complete problems? Can we provide complexity-theoretic evidence for the hardness of Uhlmann transformations? What computational restrictions does this place on our ability to e.g. achieve optimal communication rates in quantum Shannon theory?

The goal of this paper is to study such questions formally. Our first main contribution is to provide a general formal framework for reasoning about unitary complexity. This involves extending many of the traditional notions of complexity theory, such as reductions, complexity classes, complete problems, etc. to quantum state transformations and requires us to deal with many subtleties that arise in the unitary setting. Our second main contribution is to analyze the complexity of the Uhlmann Transformation Problem within this framework. This in turn allows us to show relationships between unitary complexity classes such as showing that (average case versions of) the classes $\mathsf{unitaryPSPACE}$ and $\mathsf{unitaryQIP}$ are equal. Finally, we show how the Uhlmann transformation problem plays a central role in connecting the complexity of many natural tasks in quantum information processing. For example, we establish reductions and equivalences between Uhlmann transformation problem and the security of quantum commitment schemes, falsifiable quantum cryptographic assumptions, quantum state compression, and more.

2 A fully quantum complexity theory

In [49] Rosenthal and Yuen initiated the study of complexity classes for *state synthesis* and *unitary synthesis* problems. A state synthesis problem is a sequence $(\rho_x)_{x\in\{0,1\}^*}$ of quantum states. A *state complexity class* is a collection of state synthesis problems that captures the computational resources needed to synthesize (i.e., generate) the states. For example, [49]

defined the class statePSPACE as the set of all state sequences $(\rho_x)_{x\in\{0,1\}^*}$ for which there is a polynomial-space (but possibly exponential-time) quantum algorithm A that, on input x, outputs an approximation to the state ρ_x.

Unitary complexity classes, which are the focus of this work, describe the computational resources needed to perform state *transformations*, formalized as *unitary synthesis problems*. A unitary synthesis problem is a sequence of unitary[2] operators $(U_x)_{x\in\{0,1\}^*}$ and a unitary complexity class is a collection of unitary synthesis problems. For example the class unitaryBQP is the set of all sequences of unitary operators $(U_x)_{x\in\{0,1\}^*}$ where there is a polynomial-time quantum algorithm A that, given an *instance* $x \in \{0,1\}^*$ and a quantum system B as input, (approximately) applies U_x to system B. As a simple example, any sequence of unitaries $(U_x)_x$ where x is simply (an explicit encoding of) a sequence of quantum gates that implement the unitary is obviously in unitaryBQP, since given x, the algorithm A can just execute the circuit specified by x in time polynomial in the length of x. On the other hand, x could also specify a unitary in a sequence in a more implicit way (e.g. by circuits for two quantum states between which U_x is meant to be the Uhlmann transformation), in which case the sequence $(U_x)_x$ could be harder to implement.

The reason we say that the algorithm A is given a *system* instead of a *state* is to emphasize that the state of the system is not known to the algorithm ahead of time, and in fact the system may be part of a larger entangled state. Thus the algorithm has to coherently apply the transformation U_x to the given system, maintaining any entanglement with an external system. This makes unitary synthesis problems fundamentally different, and in many cases harder to analyse, than state synthesis problems.

Traditional complexity classes like P, NP, and BQP have proven to be powerful ways of organizing and comparing the difficulty of different decision problems. In a similar way, state and unitary complexity classes are useful for studying the complexity of quantum states and of quantum state transformations. We can then ask about the existence of complete problems, reductions, inclusions, separations, closure properties, and more. Importantly, state and unitary complexity classes provide a useful language to articulate questions and conjectures about the computational hardness of inherently quantum problems. For example, we can ask whether unitaryPSPACE is contained in $\text{unitaryBQP}^{\text{PSPACE}}$ – in other words, can polynomial-space-computable unitary transformations be also computed by a polynomial-time quantum computer that is given oracle access to a PSPACE decision oracle?[3]

Unitary synthesis problems, classes, and reductions

We begin by giving general definitions for unitary synthesis problems and a number of useful unitary complexity classes, e.g. unitaryBQP and unitaryPSPACE. We then define a notion of *reductions* between unitary synthesis problems. Roughly speaking, we say that a unitary synthesis problem $\mathcal{U} = (U_x)_x$ polynomial-time reduces to $\mathcal{V} = (V_x)_x$ if an efficient algorithm for implementing $\mathcal{V}$ implies an efficient algorithm for implementing $\mathcal{U}$.

Next, we define *distributional* unitary complexity classes that capture the *average case complexity* of solving a unitary synthesis problem. Here, the unitary only needs to be implemented on an input state *randomly chosen* from some distribution $\mathcal{D}$ which is known ahead of time. This is a natural generalisation of traditional average-case complexity

[2] In our formal definition of unitary synthesis problems (see [13]), the U_x's are technically partial isometries, which is a promise version of unitaries, but we gloss over the distinction for now.

[3] This is an open question, and is related to the "Unitary Synthesis Problem" raised by Aaronson and Kuperberg [4].

statements to the unitary setting. This notion turns out to be particularly natural in the context of entanglement transformation problems because it is closely related to implementing the unitary on part of an entangled state $|\psi\rangle$.

The notion of average case complexity turns out to be central to our paper: nearly all of our results are about average-case unitary complexity classes and the average-case complexity of the Uhlmann Transformation Problem. Thus the unitary complexity classes we mainly deal with will be avgUnitaryBQP and avgUnitaryPSPACE, which informally mean sequences of unitaries that can be implemented by time-efficient and space-efficient quantum algorithms, respectively, and where the implementation error is measured with respect to inputs drawn from a fixed distribution over quantum states.

Interactive proofs for unitary synthesis

We then explore models of *interactive proofs* for unitary synthesis problems. Roughly speaking, in an interactive proof for a unitary synthesis problem $\mathcal{U} = (U_x)_x$, a polynomial-time verifier receives an instance x and a quantum system B as input, and interacts with an all-powerful but untrusted prover to try to apply U_x to system B. As usual in interactive proofs, the main challenge is that the verifier does not trust the prover, so the protocol has to test whether the prover actually behaves as intended. We formalize this with the complexity classes unitaryQIP and avgUnitaryQIP, which capture unitary synthesis problem that can be verifiably implemented in this interactive model. This generalizes the interactive state synthesis model studied by [49, 40].[4] The primary difference between the state synthesis and unitary synthesis models is that in the former, the verifier starts with a fixed input state (say, the all zeroes state), while in the latter the verifier receives a quantum system B in an unknown state that has to be transformed by U_x.

Zero-knowledge unitary synthesis

In the context of interactive protocols, we also introduce a notion of *zero-knowledge protocols* for unitary synthesis problems. Roughly speaking, a protocol is zero-knowledge if the interaction between the verifier and prover can be efficiently reproduced by an algorithm (called the *simulator*) that does not interact with the prover at all. This way, the verifier can be thought of as having learned no additional knowledge from the interaction aside from the fact that the task was solved [21]. The counterintuitive concept of zero-knowledge proofs has been one of the most consequential discoveries in complexity theory and cryptography.

Motivated by this, we introduce the unitary complexity class avgUnitaryHVSZK,[5] which is a unitary synthesis analogue of the decision class HVQSZK in traditional complexity theory [57], which captures the concept of *honest-verifier quantum zero-knowledge proofs*. Interestingly, for reasons that we explain in more detail in [13], the average-case aspect of avgUnitaryHVSZK appears to be necessary to obtain a nontrivial definition of zero-knowledge in the unitary synthesis setting.

Just like there is a zoo of traditional complexity classes [3], we expect that many unitary complexity classes can also be meaningfully defined and explored. In this paper we focus on the ones that turn out to be tightly related to the Uhlmann Transformation Problem. We discuss these relationships next.

[4] The class unitaryQIP was also briefly discussed informally by Rosenthal and Yuen [49].

[5] The "HV" modifier signifies that the zero-knowledge property is only required to hold with respect to verifiers that honestly follow the protocol, and the "S" in "SZK" signifies that it is *statistical* zero-knowledge.

▶ Remark 1. For simplicity, in the introduction we present informal statements of our results that gloss over some technical details that would otherwise complicate the result statement. For example, we do not distinguish between unitary synthesis problems and distributional versions of them. Formal results can be found in the full version of the paper [13].

3 Structural results about the Uhlmann Transformation Problem

Equipped with the proper language to talk about unitary synthesis problems, we introduce the Uhlmann Transformation Problem (see [13, Part II] for the formal statements). We define the unitary synthesis problem UHLMANN to be the sequence $(U_x)_{x\in\{0,1\}^*}$ where we interpret an instance x as an explicit encoding (as a list of gates) of a pair of quantum circuits (C, D) such that C and D, on the all-zeroes input, output pure bipartite states $|C\rangle, |D\rangle$ on the same number of qubits, and U_x is an associated Uhlmann transformation mapping $|C\rangle$ to $|D\rangle$ by acting on a local system. Usually, we will assume that C and D output $2n$ qubits (for some n specified as part of x) and the Uhlmann transformation acts on the last n qubits. If x does not specify such a pair, then an algorithm implementing the unitary synthesis problem is allowed to behave arbitrarily on such x; this is formally captured by allowing partial isometries as part of unitary synthesis problems in [13, Section 3].

Furthermore, for a parameter $0 \le \kappa \le 1$ we define the problem UHLMANN_κ, which is the same as UHLMANN, except that it is restricted to instances corresponding to states $|C\rangle, |D\rangle$ where the fidelity between the reduced density matrices ρ, σ of $|C\rangle, |D\rangle$ respectively on the first subsystem is at least κ; recall by Uhlmann's theorem that κ lower bounds how much overlap $|C\rangle$ can achieve with $|D\rangle$ by a local transformation. By definition, UHLMANN_κ instances are at least as hard as $\text{UHLMANN}_{\kappa'}$ instances when $\kappa \le \kappa'$. We provide formal definitions of UHLMANN, UHLMANN_κ, and their distributional versions in [13, Section 5].

Zero-knowledge and the Uhlmann Transformation Problem

We show that the Uhlmann Transformation Problem (with fidelity parameter $\kappa = 1$) *characterizes* the complexity of the unitary complexity class avgUnitaryHVPZK, which is the unitary synthesis version of the decision classes PZK and HVQPZK [56]. Here, PZK stands for "perfect zero knowledge", and refers to the special case of statistical zero-knowledge where the simulator can *perfectly* reproduce the view of the verifier.

▶ **Theorem 2** (Informal)**.** *UHLMANN_1 is complete for* avgUnitaryHVPZK *under polynomial-time reductions.*

This is formally stated and proved in [13, Section 6.1]. To show completeness we have to prove two directions. The first direction is to show that every (distributional) unitary synthesis problem in avgUnitaryHVPZK polynomial-time reduces to (the distributional version of) UHLMANN_1. This uses a characterization of quantum interactive protocols due to Kitaev and Watrous [33].

The second direction is to show that UHLMANN_1 is in avgUnitaryHVPZK by exhibiting an (honest-verifier) zero-knowledge protocol to solve the Uhlmann Transformation Problem. Our protocol is rather simple: in the average case setting, we assume that the verifier receives the last n qubits of the state $|C\rangle = C|0^{2n}\rangle$, and the other half is inaccessible. Its goal is to transform, with the help of a prover, the global state $|C\rangle$ to $|D\rangle$ by only acting on the last n qubits that it received as input. To this end, the verifier generates a "test" copy of $|C\rangle$ on its own, which it can do because C is a polynomial-size circuit. The verifier then sends to the prover two registers of n qubits; one of them is the first half of the test copy and one of

them (call it A) holds the "true" input state. The two registers are randomly shuffled. The prover is supposed to apply the Uhlmann transformation U to both registers and send them back. The verifier checks whether the "test" copy of $|C\rangle$ has been transformed to $|D\rangle$ by applying the inverse circuit $D^\dagger$ to the test copy and checking if all qubits are zero. If so, it accepts and outputs the register A, otherwise the verifier rejects.

If the prover is behaving as intended, then both the test copy and the "true" copy of $|C\rangle$ are transformed to $|D\rangle$. Furthermore, the prover cannot tell which of its two registers corresponds to the test copy, and thus if it wants to pass the verification with high probability, it has to apply the correct Uhlmann transformation on both registers. This shows that the protocol satisfies the completeness and soundness properties of an interactive proof. The zero-knowledge property is also straightforward: if both the verifier and prover are acting according to the protocol, then before the verifier's first message to the prover, the reduced state of the verifier is $|C\rangle\langle C| \otimes \rho$ (where ρ is the reduced density matrix of $|C\rangle$), and at the end of the protocol, the verifier's state is $|D\rangle\langle D| \otimes U\rho U^\dagger$. Both states can be produced in polynomial time.

One may ask: if the simulator can efficiently compute the state $U\rho U^\dagger$ without the help of the prover, does that mean the Uhlmann transformation U can be implemented in polynomial time? The answer is no, since the simulator only has to prepare the appropriate reduced state (i.e. essentially solve a state synthesis task), which is easy since the starting and ending states of the protocol are efficiently computable; in particular, $U\rho U^\dagger$ is (approximately) the reduced state of $|D\rangle$, which is easy to prepare. In contrast, the verifier has to implement the Uhlmann transformation on a *specific* set of qubits that are entangled with a *specific* external register, i.e. it has to perform a state transformation task that preserves coherence with the purifying register. This again highlights the distinction between state and unitary synthesis tasks.

A complete problem for avgUnitaryHVSZK?

It is natural to wonder about the complexity of $\textsc{Uhlmann}_\kappa$ for fidelity promise $\kappa < 1$. In other words, the reduced density matrices of the two states $|C\rangle, |D\rangle$ are not exactly equal. A reasonable conjecture is that $\textsc{Uhlmann}_\kappa$ (for non-negligible κ, say), is complete for $\mathsf{avgUnitaryHVSZK}$. This would correspond to the famous classical complexity result that the problem of distinguishing between whether two probability distributions (represented via sampling circuits) are close or far in trace distance is a SZK-complete problem [50].

In [13, Section 6.3] we argue that this conjecture is true assuming that a unitary version of the *polarization lemma* holds, which was instrumental for the SZK-completeness result of Sahai and Vadhan [50]. The unitary polarization lemma, if true, would state that $\textsc{Uhlmann}_\kappa$ polynomial-time reduces to $\textsc{Uhlmann}_{1-2^{-\mathrm{poly}(n)}}$ for all inverse polynomial κ.

The succinct Uhlmann Transformation Problem

We also define a *succinct* version of the Uhlmann Transformation Problem (denoted by $\textsc{SuccinctUhlmann}$), where the string x encodes a pair $(\hat{C}, \hat{D})$ of *succinct descriptions* of quantum circuits C, D. By this we mean that $\hat{C}$ (resp. $\hat{D}$) is a classical circuit that, given a number $i \in \mathbb{N}$ written in binary, outputs the i'th gate in the quantum circuit C (resp. D). Thus the circuits C, D in general can have *exponential* depth (in the length of the instance string x) and generate states $|C\rangle, |D\rangle$ that are unlikely to be synthesizable in polynomial time. Thus the task of synthesizing the Uhlmann transformation U that maps $|C\rangle$ to a state with maximum overlap with $|D\rangle$, intuitively, should be much harder than the non-succinct version. We confirm this intuition with the following result:

▶ **Theorem 3** (Informal). *SuccinctUhlmann is complete for* avgUnitaryPSPACE *under polynomial-time reductions.*

The class avgUnitaryPSPACE corresponds to distributional unitary synthesis problems that can be solved using a polynomial-space (but potentially exponential-depth) quantum algorithm. The fact that SuccinctUhlmann ∈ avgUnitaryPSPACE was already proved by Metger and Yuen [40], who used this to show that optimal prover strategies for quantum interactive proofs can be implemented in avgUnitaryPSPACE.[6] The fact that avgUnitaryPSPACE reduces to SuccinctUhlmann is because solving a distributional unitary synthesis problem $(U_x)_x$ in avgUnitaryPSPACE is equivalent to applying a local unitary that transforms an entangled state $|\psi_x\rangle$ representing the distribution to $(\mathbb{1} \otimes U_x)|\psi_x\rangle$. This is nothing but an instance of the SuccinctUhlmann transformation problem. We refer to the proof of [13, Lemma 7.5] for details.

We show another completeness result for SuccinctUhlmann:

▶ **Theorem 4** (Informal). *SuccinctUhlmann is complete for* avgUnitaryQIP *under polynomial-time reductions.*

Here, the class avgUnitaryQIP is like avgUnitaryHVPZK except there is no requirement that the protocol between the honest verifier and prover can be efficiently simulated. The proof of Theorem 4 starts similarly to the proof of the avgUnitaryHVPZK-completeness of Uhlmann, but requires additional ingredients, such as the state synthesis protocol of [49, 48] and the ability to simulate reflections about a state, given copies of the state [29]. We prove this by showing that SuccinctUhlmann is contained in avgUnitaryQIP, avgUnitaryQIP ⊆ avgUnitaryPSPACE, and then argue that avgUnitaryPSPACE is polynomial-time reducible to SuccinctUhlmann (see [13] for details).

Theorems 3 and 4 imply the following unitary complexity analogue of the QIP = PSPACE theorem [28] and the stateQIP = statePSPACE theorem [49, 40]:

▶ **Theorem 5.** avgUnitaryQIP = avgUnitaryPSPACE.

This partially answers an open question of [49, 40], who asked whether unitaryQIP = unitaryPSPACE (although they did not formalize this question to the same level as we do here).

4 Centrality of the Uhlmann Transformation Problem

In [13, Part III], we relate the Uhlmann Transformation Problem to quantum information processing tasks in a variety of areas: quantum cryptography, quantum Shannon theory, and high energy physics. We show that the computational complexity of a number of these tasks is in fact essentially *equivalent* to the hardness of Uhlmann. For some other problems we show that they are efficiently reducible to Uhlmann or SuccinctUhlmann. Although some of these connections have been already observed in prior work, we believe that the framework of unitary complexity theory formalizes and clarifies the relationships between these different problems.

We proceed to give a high level overview of our applications of the Uhlmann Transformation Problem and unitary complexity theory.

[6] This was phrased in a different way in their paper, as avgUnitaryPSPACE was not yet defined.

4.1 Quantum cryptography

We show that the Uhlmann Transformation Problem is deeply intertwined with the security of quantum cryptography. First, we show the security of quantum commitment schemes is *equivalent* to the average-case hardness of the Uhlmann Transformation Problem.

Quantum commitments

A bit commitment scheme is a fundamental cryptographic primitive that allows two parties (called a *sender* and *receiver*) to engage in a two-phase communication protocol: in the first phase (the "commit phase"), the sender sends a commitment (i.e. some string) to a bit b to the receiver; the *hiding* property of a bit commitment scheme ensures that the receiver cannot decide the value of b from this commitment string alone. In the second phase (the "reveal phase"), the sender sends another string to the receiver that allows the receiver to compute the value of b; the *binding* property of commitments ensures that the sender can only reveal the correct value of b, i.e. if the sender sent a reveal string that was meant to convince the receiver it had committed to a different value of b, the receiver would detect this.

Commitment schemes – even quantum ones – require efficiency constraints on the adversary [39, 37]; at least one of the hiding or binding properties must be computational. In classical cryptography, commitment schemes can be constructed from one-way functions [43], but recent works suggest the possibility of basing quantum commitment schemes on weaker, inherently quantum assumptions such as the existence of pseudorandom states [34, 8, 41, 35] or EFI pairs [16].

The following theorem shows that the existence of secure quantum commitment schemes is essentially equivalent to Uhlmann being hard on average. Roughly speaking, hardness on average means that there is an efficiently sampleable distribution over pairs of quantum circuits (C, D) such that all polynomial-time algorithms fail to implement the Uhlmann transformation corresponding to $(|C\rangle, |D\rangle)$ with non-negligible probability over the sampling of (C, D).

▶ **Theorem 6** (Informal). *$\textsc{Uhlmann}_{1-\epsilon}$ for some negligible ϵ is hard on average if and only if secure quantum commitments exist.*

This theorem is formally stated and proved as [13, Theorem 8.10]. This formalizes a connection between Uhlmann transformations and quantum commitments that was suggested by Yan in his in-depth study of properties of quantum bit commitments [59]. The necessity for the hardness of Uhlmann is implicit in the original impossibility proofs of information-theoretic security for commitments [39, 37]; the sufficiency is due to the fact that *non-interactive quantum commitments* can be constructed from hard Uhlmann instances.

Given the close connection between zero knowledge protocols for unitary synthesis and the Uhlmann Transformation Problem, we also prove the following:

▶ **Theorem 7** (Informal). *If there is a hard distribution of instances for* avgUnitaryHVSZK*, then secure quantum commitments exist.*

We note that this would follow as an immediate corollary if we were able to prove that $\textsc{Uhlmann}_{1-\epsilon}$ is a complete problem for avgUnitaryHVSZK; however as mentioned previously this remains a conjecture. We instead prove Theorem 7 directly by showing that hard-on-average problems in avgUnitaryHVSZK implies $\textsc{Uhlmann}_{1-\epsilon}$ is hard on average.

This is analogous to the classical result of Ostrovsky [44] who showed that if the classical complexity class SZK is hard on average, then one-way functions (and thus secure bit commitments [42]) exist. This is formally stated and proved as [13, Theorem 8.11].

Minimal assumptions in quantum cryptography

In classical cryptography, the existence of one-way functions is considered a *minimal assumption* in the sense that the security of virtually all (classical) cryptography implies it [27, 26]. It is a fascinating open question of what is the minimal assumption (if there exists one) in quantum cryptography; as of writing the leading contender for the minimal quantum cryptographic assumption is the existence of quantum commitments, meaning that many quantum cryptographic primitives can be shown to imply the existence of quantum commitments [16, 32]. If quantum commitments are indeed minimal (mirroring the setting of classical cryptography), then this would show that the hardness of the Uhlmann Transformation Problem is necessary for computationally secure quantum cryptography.

Breaking falsifiable quantum cryptographic assumptions

While we don't know yet if the hardness of the Uhlmann Transformation Problem is necessary for computational quantum cryptography, we show that the hardness of the *succinct* Uhlmann Transformation Problem is necessary for the security of a wide class of quantum cryptographic primitives. We consider the general notion of a *falsifiable quantum cryptographic assumption*, which can be seen as a quantum analogue of the notion of a falsifiable assumption considered by Naor [43] as well as Gentry and Wichs [20]. Our notion of a falsifiable quantum cryptographic assumption captures almost any reasonable definition of security in quantum cryptography which can be phrased in terms of an interactive *security game* between an adversary and a challenger. We show the following generic upper bound on the complexity of breaking falsifiable quantum cryptographic assumptions (see [13, Theorem 8.15] for the formal statement):

▶ **Theorem 8** (Informal). *A falsifiable quantum cryptographic assumption is either information-theoretically secure, or the task of breaking security reduces to* SuccinctUhlmann.

Since SuccinctUhlmann is complete for avgUnitaryPSPACE (Theorem 3), this means that avgUnitaryBQP $\neq$ avgUnitaryPSPACE is a necessary complexity-theoretic assumption for computational quantum cryptography. This suggests that unitary complexity provides the appropriate framework to establish a close link between complexity theory and quantum cryptography, as recent work [34, 8, 41, 35, 38] has shown that traditional complexity theoretic assumptions are not always linked to quantum cryptography in the way one would expect.

4.2 Quantum Shannon theory applications

Quantum Shannon theory studies the achievability and limits of quantum communication tasks (see [58, 31, 47] for a comprehensive overview). While the information-theoretic aspects of quantum communication tasks are well-understood, the complexity of implementing these protocols has received remarkably little attention. Here, we study the computational complexity of some fundamental tasks in quantum Shannon theory, namely noisy channel decoding and compression of quantum states using our framework for unitary complexity and our results on the Uhlmann transformation problem.[7]

[7] We also note that in independent work after the publication of our results, Arnon-Friedman, Brakerski, and Vidick have investigated the computational aspects of entanglement distillation [11], showing that in general entanglement distillation is computationally infeasible assuming quantum commitments exist. It would be interesting to connect their results to our framework for unitary complexity to build up a more rigorous theory of the complexity of quantum Shannon tasks.

Decodable channel problem

Consider a quantum channel $\mathcal{N}$ that maps a register A to a register B. Suppose that the channel $\mathcal{N}$ is *decodable*, meaning that it is possible to information-theoretically (approximately) recover the information sent through the channel; i.e., there exists a decoding channel $\mathcal{D}$ mapping register B back to register A such that $\mathcal{D}_{\mathsf{B}\to\mathsf{A}'}\left(\mathcal{N}_{\mathsf{A}\to\mathsf{B}}(\Phi_{\mathsf{AR}})\right) \approx \Phi_{\mathsf{A}'\mathsf{R}}$, where $|\Phi\rangle_{\mathsf{AR}}$ is the maximally entangled state. Note that the register R is not touched.

Important examples of decodable channels come from coding schemes for noisy quantum channels: suppose $\mathcal{K}$ is a noisy quantum channel that has capacity C (meaning it is possible to (asymptotically) transmit C qubits through $\mathcal{K}$). Let $\mathcal{E}$ denote a channel that takes C qubits and maps it to an input to $\mathcal{K}$. For example, we can think of $\mathcal{E}$ as an encoder for a quantum error-correcting code. If $\mathcal{E}$ is a good encoding map, the composite channel $\mathcal{N} : \rho \mapsto \mathcal{K}(\mathcal{E}(\rho))$ is decodable.

We define the *Decodable Channel Problem*: given as input a circuit description of a channel $\mathcal{N}$ that maps register A to register B and furthermore is promised to be decodable, and given the register B of the state $(\mathcal{N}\otimes\mathbb{1})(\Phi_{\mathsf{AR}})$, decode and output a register $\mathsf{A}' \equiv A$ such that the final joint state of $\mathsf{A}'\mathsf{R}$ is close to $|\Phi\rangle$. Although it is information-theoretically possible to decode the output of $\mathcal{N}$, it may be computationally intractable to do so. In fact, we can characterize the complexity of the Decodable Channel Problem:

▶ **Theorem 9** (Informal). *The Decodable Channel Problem can be solved in polynomial-time up to inverse polynomial error if and only if* Uhlmann *can be solved in polynomial-time up to inverse polynomial error.*

This theorem is formally stated and proved as [13, Theorem 9.6]; since we do not expect that Uhlmann is solvable in polynomial-time, this suggests that the Decodable Channel Problem is computationally hard in general. The main idea behind the upper bound (Decodable Channel Problem is easy if Uhlmann is easy) is that a channel $\mathcal{N}$ is decodable if and only if the output of the *complementary channel*[8] $\mathcal{N}^c$, when given register A of the maximally entangled state $|\Phi\rangle_{\mathsf{AR}}$, is approximately unentangled with register R. Thus by Uhlmann's theorem there exists an Uhlmann transformation acting on the output of the channel $\mathcal{N}$ that recovers the maximally entangled state. If Uhlmann $\in$ avgUnitaryBQP, then this transformation can be performed efficiently.

The proof of the lower bound (Decodable Channel Problem is hard if Uhlmann is hard) draws inspiration from quantum commitments. As discussed earlier, the hardness of Uhlmann essentially implies the existence of secure quantum commitments, and in particular one where the hiding property is computational. From this, we can construct a hard instance of the Decodable Channel Problem: consider a channel $\mathcal{N}$ that takes as input a single bit $|b\rangle$, and then outputs the commitment register of the commitment to bit b (and discards the reveal register). The ability to decode this "commitment channel" implies the ability to break the hiding property of the underlying commitment scheme, and therefore decoding must be computationally hard.

Compression of quantum information

Another fundamental task in information theory – both classical and quantum – is compression of data. Shannon's source coding theorem shows that the Shannon entropy of a random variable X characterizes the rate at which many independent copies of X can be

[8] The output of the complementary channel can be thought of as the qubits that a purification (formally, a Stinepring dilation) of the channel $\mathcal{N}$ discards to the environment.

compressed [52]. Similarly, Schumacher proved that the von Neumann entropy of a density matrix ρ characterizes the rate at which many independent copies of ρ can be (coherently) compressed [51].

We consider the *one-shot* version of the information compression task, where one is given just one copy of a density matrix ρ (rather than many copies) and the goal is to compress it to as few qubits as possible while being able to recover the original state within some error. In the one-shot setting the von Neumann entropy no longer characterizes the optimal compression of ρ; instead this is given by a one-shot entropic quantity known as the *smoothed max-entropy* [54]. What is the computational effort required to perform near-optimal one-shot compression of quantum states? Our next result gives upper and lower bounds for the computational complexity of this task:

▶ **Theorem 10** (Informal). *Quantum states can be optimally compressed to their smoothed max entropy in polynomial-time if* $\textsc{Uhlmann}_{1-\epsilon} \in \mathsf{avgUnitaryBQP}$ *for some negligible* ϵ*. Furthermore, if stretch pseudorandom state generators exist, then optimal compression of quantum states cannot be done in polynomial time.*

The upper bound (i.e., compression is easy if Uhlmann is easy) is proved using a powerful technique in quantum information theory known as *decoupling* [19]. The hardness result for compression is proved using a variant of *pseudorandom states*, a cryptographic primitive that is a quantum analogue of pseudorandom generators [29].

4.2.1 Black-hole radiation decoding

In recent years, quantum information and quantum complexity have provided a new lens on long-standing questions surrounding the quantum-mechanical description of black holes. [46, 7, 22, 17, 53, 14, 60]. We consider applications of the Uhlmann Transformation Problem to computational tasks arising from this research.

In particular, we consider the Harlow-Hayden *black hole radiation decoding task* [22], which is defined as follows. We are given as input a circuit description of a tripartite state $|\psi\rangle_{\mathsf{BHR}}$ that represents the global pure state of a single qubit (register B), the interior of a black hole (register H), and the Hawking radiation that has been emitted by the black hole (register R). Moreover, we are promised that it is possible to *decode* from the emitted radiation R a single qubit A that forms a maximally entangled state $|\mathrm{EPR}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ with register B. The task is to perform this decoding when given register R of a system in the state $|\psi\rangle$.

Harlow and Hayden [22] showed that the decoding task is computationally intractable assuming that $\mathsf{SZK} \not\subseteq \mathsf{BQP}$. However, precisely characterizing the task's complexity (i.e., providing an equivalence rather than a one-way implication) appears to require the notions of a fully quantum complexity theory. Brakerski recently showed that this task is equivalent to breaking the security of a quantum cryptographic primitive known as EFI pairs [15]. We reformulate this equivalence in our unitary complexity framework to show that black hole radiation decoding (as formalised above) can be solved in polynomial-time if and only if $\textsc{Uhlmann} \in \mathsf{avgUnitaryBQP}$.

5 Summary and future directions

Computational tasks with quantum inputs and/or outputs are ubiquitous throughout quantum information processing. The traditional framework of complexity theory, which is focused on computational tasks with classical inputs and outputs, cannot naturally capture the complexity of these "fully quantum" tasks.

In this paper we introduce a framework to reason about the computational complexity of unitary synthesis problems. We then use this framework to study Uhlmann's theorem through an algorithmic lens, i.e. to study the complexity of Uhlmann transformations. We prove that variants of the Uhlmann Transformation Problem are complete for some unitary complexity classes, and then explore relationships between the Uhlmann Transformation Problem and computational tasks in quantum cryptography, quantum Shannon theory, and high energy physics.

The study of the complexity of state transformation tasks is a very new field and we hope that our formal framework of unitary complexity theory and our findings about the Uhlmann Transformation Problem provide a useful starting point for a rich theory of the complexity of "fully quantum" problems. Many questions in this direction have yet to be explored. Throughout this paper, we have included many concrete open problems, which we hope will spark future research in this new direction in complexity theory. Additionally, our work suggests some high-level, open-ended future directions to explore:

Populating the zoo

An important source of the richness of computational complexity theory is the variety of computational problems that are studied. For example, the class NP is so interesting because it contains many complete problems that are naturally studied across the sciences [45], and the theory of NP-completeness gives a unified way to relate them to each other.

Similarly, a fully quantum complexity theory should have its own zoo of problems drawn from a diverse range of areas. We have shown that core computational problems in quantum cryptography, quantum Shannon theory, and high energy physics can be related to each other through the language of unitary complexity theory. What are other natural problems in e.g. quantum error-correction, quantum metrology, quantum chemistry, or condensed matter physics, and what can we say about their computational complexity?

The crypto angle

Complexity and cryptography are intimately intertwined. Operational tasks in cryptography have motivated models and concepts that have proved indispensible in complexity theory (such as pseudorandomness and zero-knowledge proofs), and conversely complexity theory has provided a rigorous theoretical foundation to study cryptographic hardness assumptions.

We believe that there can be a similarly symbiotic relationship between quantum cryptography and a fully quantum complexity theory. Recent quantum cryptographic primitives such as quantum pseudorandom states [29] or one-way state generators [41] are unique to the quantum setting, and the relationships between them are barely understood. For example, an outstanding question is whether there is a meaningful *minimal hardness assumption* in quantum cryptography, just like one-way functions are in classical cryptography. Can a fully quantum complexity theory help answer this question about minimal quantum cryptographic assumptions, or at least provide some guidance? For example, there are many beautiful connections between one-way functions, average-case complexity, and Kolomogorov complexity [27, 26, 36]. Do analogous results hold in the fully quantum setting?

The learning theory angle

Quantum learning theory has also seen rapid development, particularly on the topic of quantum state learning [1, 25, 18, 9]. Learning quantum states or quantum processes can most naturally be formulated as tasks with quantum inputs. Traditionally these tasks have

been studied in the information-theoretic setting, where sample complexity is usually the main measure of interest. However we can also study the computational difficulty of learning quantum objects. What does a complexity theory of quantum learning look like?

Traditional versus fully quantum complexity theory

While traditional complexity theory appears to have difficulty reasoning about fully quantum tasks, can we obtain *formal* evidence that the two theories are, in a sense, independent of each other? For example, can we show that $\mathsf{P} = \mathsf{PSPACE}$ does not imply $\mathsf{unitaryBQP} = \mathsf{unitaryPSPACE}$? One would likely have to show this in a *relativized* setting, i.e., exhibit an oracle O relative to which $\mathsf{P}^O = \mathsf{PSPACE}^O$ but $\mathsf{unitaryBQP}^O \neq \mathsf{unitaryPSPACE}^O$. Another way would be to settle Aaronson and Kuperberg's "Unitary Synthesis Problem" [4] in the negative; see [38] for progress on this. Such results would give compelling evidence that the reasons for the hardness of unitary transformations are intrinsically different than the reasons for the hardness of a Boolean function. More generally, what are other ways of separating traditional from fully quantum complexity theory?

References

1 Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, 2007. doi:10.1098/rspa.2007.0113.

2 Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016. arXiv:1607.05256.

3 Scott Aaronson. The Complexity Zoo. https://complexityzoo.net/Complexity_Zoo, 2023. Accessed: 2023-04-01.

4 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. doi:10.4086/toc.2007.v003a007.

5 Anura Abeyesinghe, Igor Devetak, Patrick Hayden, and Andreas Winter. The mother of all protocols: Restructuring quantum information's family tree. *Proceedings: Mathematical, Physical and Engineering Sciences*, 465(2108):2537–2563, 2009. doi:10.1098/rspa.2009.0202.

6 Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. Quantum algorithmic measurement. *Nature Communications*, 13(1):887, 2022. doi:10.1038/s41467-021-27922-0.

7 Ahmed Almheiri, Donald Marolf, Joseph Polchinski, and James Sully. Black holes: complementarity or firewalls? *Journal of High Energy Physics*, 2013(2):62, February 2013. doi:10.1007/JHEP02(2013)062.

8 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. doi:10.1007/978-3-031-15802-5_8.

9 Anurag Anshu and Srinivasan Arunachalam. A survey on the complexity of learning quantum states. *Nature Reviews Physics*, 6(1):59–69, January 2024. doi:10.1038/s42254-023-00662-4.

10 Anurag Anshu, Rahul Jain, and Naqueeb Ahmad Warsi. A one-shot achievability result for quantum state redistribution. *IEEE Transactions on Information Theory*, 64(3):1425–1435, 2018. doi:10.1109/TIT.2017.2776112.

11 Rotem Arnon-Friedman, Zvika Brakerski, and Thomas Vidick. Computational entanglement theory, 2023. arXiv:2310.02783.

12 Mario Berta, Matthias Christandl, and Renato Renner. The quantum reverse shannon theorem based on one-shot information theory. *Communications in Mathematical Physics*, 306(3):579–615, 2011. doi:10.1007/s00220-011-1309-7.

13 John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the uhlmann transformation problem. *arXiv preprint arXiv:2306.13073*, 2023. doi:10.48550/arXiv.2306.13073.

14 Adam Bouland, Bill Fefferman, and Umesh Vazirani. Computational Pseudorandomness, the Wormhole Growth Paradox, and Constraints on the AdS/CFT Duality (Abstract). In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2020.63.

15 Zvika Brakerski. Black-hole radiation decoding is quantum cryptography. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 37–65, Cham, 2023. Springer Nature Switzerland. doi:10.1007/978-3-031-38554-4_2.

16 Zvika Brakerski, Ran Canetti, and Luowen Qian. On the Computational Hardness Needed for Quantum Cryptography. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2023.24.

17 Adam R Brown, Daniel A Roberts, Leonard Susskind, Brian Swingle, and Ying Zhao. Complexity, action, and black holes. *Physical Review D*, 93:086006, April 2016. doi:10.1103/PhysRevD.93.086006.

18 Costin Bădescu and Ryan O'Donnell. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1398–1411, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3406325.3451109.

19 Frédéric Dupuis. *The decoupling approach to quantum information theory*. PhD thesis, Université de Montréal, 2009. doi:10.71781/10720.

20 Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 99–108, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993636.1993651.

21 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. doi:10.1137/0218012.

22 Daniel Harlow and Patrick Hayden. Quantum computation vs. firewalls. *Journal of High Energy Physics*, 2013(6):85, 2013. doi:10.1007/JHEP06(2013)085.

23 Stephen W Hawking. Breakdown of predictability in gravitational collapse. *Physical Review D*, 14:2460–2473, 1976. doi:10.1103/PhysRevD.14.2460.

24 Patrick Hayden, Michał Horodecki, Andreas Winter, and Jon Yard. A decoupling approach to the quantum capacity. *Open Systems & Information Dynamics*, 15(01):7–19, 2008. doi:10.1142/S1230161208000043.

25 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. doi:10.1038/s41567-020-0932-7.

26 Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 134–147, 1995. doi:10.1109/SCT.1995.514853.

27 Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography. In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, 1989. doi:10.1109/SFCS.1989.63483.

28 Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Journal of the ACM*, 58(6), December 2011. doi:10.1145/2049697.2049704.

29 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. doi:10.1007/978-3-319-96878-0_5.

30 Elham Kashefi and Carolina Moura Alves. On the complexity of quantum languages, 2004. arXiv:quant-ph/0404062.

31 Sumeet Khatri and Mark M. Wilde. Principles of quantum communication theory: A modern approach, 2020. arXiv:2011.04672.

32 Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness, 2023. doi:10.48550/arXiv.2310.11526.

33 Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 608–617, New York, NY, USA, 2000. Association for Computing Machinery. doi:10.1145/335305.335387.

34 William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2021.2.

35 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*, STOC 2023, pages 1589–1602, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585225.

36 Yanyi Liu and Rafael Pass. On one-way functions and kolmogorov complexity. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1243–1254, 2020. doi:10.1109/FOCS46700.2020.00118.

37 Hoi-Kwong Lo and H.F. Chau. Why quantum bit commitment and ideal quantum coin tossing are impossible. *Physica D: Nonlinear Phenomena*, 120(1):177–187, 1998. Proceedings of the Fourth Workshop on Physics and Consumption. doi:10.1016/S0167-2789(98)00053-0.

38 Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography, 2023. doi:10.48550/arXiv.2310.08870.

39 Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997. doi:10.1103/PhysRevLett.78.3414.

40 Tony Metger and Henry Yuen. stateQIP = statePSPACE. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1349–1356, Los Alamitos, CA, USA, November 2023. IEEE Computer Society. doi:10.1109/FOCS57990.2023.00082.

41 Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022*, pages 269–295, Cham, 2022. Springer Nature Switzerland. doi:10.1007/978-3-031-15802-5_10.

42 Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4:151–158, 1991. doi:10.1007/BF00196774.

43 Moni Naor. On cryptographic assumptions and challenges. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, pages 96–109, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. doi:10.1007/978-3-540-45146-4_6.

44 Rafail Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *SCT*, pages 133–138. Citeseer, 1991. doi:10.1109/SCT.1991.160253.

45 Christos H. Papadimitriou. NP-completeness: A retrospective. In Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Automata, Languages and Programming*, pages 2–6, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg. doi:10.1007/3-540-63165-8_160.

46 John Preskill. Do black holes destroy information? In *Proceedings of the International Symposium on Black Holes, Membranes, Wormholes and Superstrings*, pages 22–39. World Scientific, 1992. doi:10.1142/9789814536752.

47 Joseph M. Renes. *Quantum Information Theory*. De Gruyter Oldenbourg, Berlin, Boston, 2022. doi:10.1515/9783110570250.

48 Gregory Rosenthal. Efficient quantum state synthesis with one query. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2508–2534. SIAM, 2024. doi:10.1137/1.9781611977912.89.

49 Gregory Rosenthal and Henry Yuen. Interactive Proofs for Synthesizing Quantum States and Unitaries. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.112.

50 Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):196–249, March 2003. doi:10.1145/636865.636868.

51 Benjamin Schumacher. Quantum coding. *Physical Review A*, 51:2738–2747, April 1995. doi:10.1103/PhysRevA.51.2738.

52 Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. doi:10.1002/j.1538-7305.1948.tb01338.x.

53 Leonard Susskind. Computational complexity and black hole horizons. *Fortschritte der Physik*, 64(1):24–43, 2016. doi:10.1002/prop.201500092.

54 Marco Tomamichel. A framework for non-asymptotic quantum information theory, 2013. arXiv:1203.2142.

55 Armin Uhlmann. The "transition probability" in the state space of a $*$-algebra. *Reports on Mathematical Physics*, 9(2):273–279, 1976. doi:10.1016/0034-4877(76)90060-4.

56 John Watrous. Limits on the power of quantum statistical zero-knowledge. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 459–468, 2002. doi:10.1109/SFCS.2002.1181970.

57 John Watrous. Zero-knowledge against quantum attacks. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pages 296–305, New York, NY, USA, 2006. Association for Computing Machinery. doi:10.1145/1132516.1132560.

58 Mark M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2 edition, 2017. doi:10.1017/CBO9781139525343.

59 Jun Yan. General properties of quantum bit commitments (extended abstract). In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022*, pages 628–657, Cham, 2022. Springer Nature Switzerland. doi:10.1007/978-3-031-22972-5_22.

60 Lisa Yang and Netta Engelhardt. The complexity of learning (pseudo)random dynamics of black holes and other chaotic systems, 2023. arXiv:2302.11013.