Frontiers in Quantum Cryptography: New Functionalities, Primitives, and Foundations

John Bostanci (Columbia), Eli Goldin (NYU), Aparna Gupte (MIT), Seyoon Raghavan (MIT), Kabir Tomer (UIUC), Tina Zhang (MIT), Zvika Brakerski (Weizmann), Yael Kalai (MIT), Henry Yuen (Columbia)

April 20, 2025– April 25, 2025

1 Overview

There are deep ties between the fields of quantum computation and cryptography. Perhaps one of the most famous examples was the discovery of an efficient quantum algorithm for factoring [37], which posed a threat to cryptographic systems based on RSA. Recently, however, there have been intense efforts to explore the ways in which quantum computers give rise to new forms of cryptography, and the nature of assumptions required for this new cryptography. A striking example of this is the recent discovery that the existence of quantum cryptography might be independent of complexity theoretic separations like the P versus NP question, unlike classical cryptography.

The goal of this meeting was to further facilitate developments in quantum cryptography, and to bring together researchers working in many fields of theoretical computer science and cryptography to share their perspectives and techniques. Extended in person meetings are vital for achieving this goal, as they allow for extended discussions and collaborations that are often difficult to maintain virtually due to the vast distances that separate participants.

One major highlight of the workshop was the success of our plenary lectures, which were held in the mornings. These lectures were overviews of recent progress in a number of areas of quantum cryptography. By encompassing both a birds-eye view of challenges in the area, as well as descriptions of important techniques and tools used in recent breakthroughs, the plenary lectures established a welcoming atmosphere that encouraged collaboration throughout the workshop. The plenary sessions are summarized in Section 2.1.

A second major highlight of the workshop was a collection of breakout sessions led by a almost all of the participants in the workshop. These breakout sessions were focused and targeted, with the intention of gathering small groups of researchers to learn and advance state of the art techniques related to recent developments in quantum cryptography. By having smaller sessions, participants were encouraged to ask questions, propose ideas, and find connections with their own work. We describe the contents of the breakout sessions in Section 2.2.

In the remainder of the week, participants took part in informal discussions and meetings on pressing problems in quantum cryptography. Section 3 summarizes the broad topics of these discussions. Section 4 contains concluding remarks and comments on the workshop.

2 Workshop Highlights

2.1 Summary of Plenary Lectures

2.1.1 Barak Nehoran: Quantum Money and Unclonable Cryptography

Barak Nehoran gave a pair of talks summarizing the field of unclonable cryptography, with a focus on quantum money. The no-cloning principle states that there is no universal algorithm that can clone an unknown quantum state. Roughly speaking, unclonable cryptography attempts to imbue unclonable ensembles of quantum states with useful cryptographic properties. The first example of an unclonable cryptographic primitive is quantum money.

Quantum money describes a family of states (colloquially called bank notes) that, together with a classical serial number, can be verified but can not be counterfeited by any efficient adversary. When the verification can only be performed by a party with a secret key (i.e. the bank itself), we call the resulting scheme a *private-key quantum money* scheme, and when it can be performed without the use of a secret key, we call the resulting quantum money scheme a *public-key quantum money* scheme.

Quantum money was first introduced in 1970, with the work of Stephen Wiesner [38]. In his seminal paper, Wiesner proposed a construction of private-key quantum money that was information theoretically secure (as in, even computationally un-bounded adversaries can not counterfeit the bank notes). In this original quantum money scheme, a bank first generates a $2n$-bit key, interpreting the first n bits as a list of bases, either the standard basis, $\{|0\rangle, |1\rangle\}$, or the Hadamard basis, $\{|+\rangle, |-\rangle\}$, and the final n bits as being encoded in the corresponding basis (so, in the Hadamard basis $0 \mapsto |+\rangle$ and $1 \mapsto |-\rangle$). Minting a bank note involves giving out the encoding of the final n bits in the appropriate basis, as well as the measurement outcomes. The bank stores the basis information privately, which it can use to verify the validity of a bank note, as well as to check that it minted the bank note itself. This original quantum money scheme went on to inspire the quantum key-exchange protocol of Bennett and Brassard [10].

However, this construction of quantum money has a number of short-comings. First, the keys of the quantum money schemes are large ($\geq n$ bits), which means that the bank necessarily stores large amounts of information. It turns out that, unlike in the case of Wiesner's scheme, private-key quantum money schemes with short keys can not be information theoretically secure, but the exact computational assumptions required to construct private-key quantum money with short keys is not known. A larger barrier in the way of using Wiesner's scheme is the fact that verification can only be performed by the bank, and therefore any parties attempting to transfer quantum money will have to first send the bank note to the bank to be verified.

The solution to this second problem is a quantum cryptographic primitive known as public-key quantum money [10, 3]. A further strengthening of this primitive is known as public-key quantum lightning [41], where the security is now that no party (even the bank itself) can get two copies of a bank note that pass verification with the same serial number. Unlike private-key quantum money, which (even with short keys) can be constructed from standard cryptographic primitives like pseudo-random functions, constructing public-key quantum money and lightning has been significantly more challenging. Since 2009, at least 13 papers have proposed constructions of quantum money and/or lightning, and none of them have successfully reduced the security of quantum money or lightning to a "standard" cryptographic assumption. Barak talk about two promising constructions of public-key quantum money, the construction from knot invariants [21], and the construction relative to hidden subspace oracles [4]. Finally, Barak gave a summary of a recent construction of public-key quantum lightning from group actions [15], which generalizes a number of previous attempts to build quantum lightning from Abelian group actions [42], although these states have not yet been used to construct unclonable encryption or copy-protection because it is not clear how to encode classical information in them at the moment.

Finally, Prabhanjan spoke about the impact of unclonable cryptography on quantum learning theory, complexity theory, and non-local computation, as well as its central role in differentiating quantum cryptography from classical cryptography.

2.1.2 Prabhanjan Ananth: Quantum Copy-Protection

Prabhanajan Ananth gave a summary talk on the state of quantum copy protection and other unclonable cryptographic primitives. Unlike quantum money and lightning (which Prabhanjan refers to as "generation

1 primitives"), which only require verifiability, quantum copy-protection and other "generation 2 primitives" ask that unclonable quantum states can be "used" to perform some task, such as evaluating a classical function f. Many of these primitives have a similar security notion, which is described by the following unclonable security game.

1. A challenger samples a state ρ, and gives it to a party A (which they could use to answer some question).
2. A must send one quantum register to each of their friends, B and C.
3. After receiving the register from A, both B and C will receive challenges, and all three of them win if both B and C submit valid answers.

For example, in the case of quantum copy protection, the first state ρ will be the evaluation key $|\psi_f\rangle$ for some function f, the challenges to B and C will be inputs x_1 and x_2, and the parties win if B outputs $f(x_1)$ and C outputs $f(x_2)$. We can define the "trivial" success probability to be the max of the success probabilities in the events that A forwards their state ρ to one of B or C, and will typically define security of a primitive to be that the success probability of any quantum polynomial-time adversaries A, B, and C is not too much more than this trivial success probability. Unlike quantum money and lightning, there are many more details to consider when talking about security of unclonable primitives. For example, some issues raised are whether the two challenges are sampled independently or the same for B and C, or how the distribution over challenges affects the trivial success probability.

Then Prabhnajan focused on a number of specific unclonable primitives, starting with copy-protection. Functions that can be learned from their input-output behavior can never be copy-protected, as A can extract a description of f and send it to both B and C. However, for a number of families of functions, including point functions, pseudorandom functions, and primitives satisfying puncturing security, there have been works in the past 3 years describing copy-protection schemes from classical oracles.

In another primitives, unclonable encryption, the state ρ is an encryption of a uniformly random bit b, and the security game is that B and C should not both be able to output b after being given a decryption key. Unlike in copy-protection, unclonable encryption is possible (with weak security) information theoretically, and recent work [11] has even shown that getting $1/\mathsf{poly}(n)$-close to the trivial success probability is possible. Building off of this work, Prabhajnan presented a "best possible" unclonable encryption scheme, in that the scheme has better security than any other unclonable encryption scheme. The idea is to sample $2n$ many orthogonal Haar random states and have the encryption of 0 be a mixture of the first n, while an encryption to 1 is the mixture of the second n.

Then Prabhanjan spoke about popular techniques for achieving unclonable security, starting with the so-called BB84 states from [38, 10]. Prabhanjan showed how coset states directly generalize these BB84 states, and presented examples of their usage. Then, it was shown how to interpret a coset state as so-called group action state, similar to those that appear in [42]

2.1.3 Michael Walter: Compiled nonlocal games

Nonlocal games have been an object of study in quantum foundations since Bell's historical observation that some nonlocal games can be won with higher probability by entangled quantum players than by classical players. A robust machinery to analyse the value of quantum nonlocal games was also fleshed out in a series of works in quantum complexity theory which culminated with the proof that MIP* = RE. Even more recently, nonlocal games have proven to be useful in cryptography.

Traditional nonlocal games involve an efficient referee (or verifier) and two or more unbounded players (or provers) who cannot communicate with each other. The referee asks the players some questions, and based on their answers decides whether to accept or reject them. This model is somewhat non-standard by cryptographic standards, because of the no-communication assumption, which cannot naïvely be enforced by any computational assumption about the hardness of some computational problem. However, recent work by Kalai, Lombardi, Vaikuntanathan and Yang proposed a *compiler* which shows how to generically map any nonlocal game into a single-prover argument system by simulating the no-communication assumption using cryptography (specifically, homomorphic encryption).

Attempts to show that this proposed compilation procedure preserves the quantum value of the nonlocal game it started with have yielded several improvements to the assumptions necessary for existing cryptographic results, such as cryptographically sound verification of QMA instances and succinct arguments for QMA. Michael Walter talked about a new result which is not on the face of it useful for any additional cryptography, but which represents the first nontrivial bound on the quantum value of *any* compiled nonlocal game. The techniques are inspired by the classic proof that the NPA hierarchy for the value of a quantum nonlocal game converges. Because MIP* = RE, implying that quantum nonlocal games are undecidable, the NPA hierarchy's *rate* of convergence to the true value of the game cannot be quantified. Walter et al.'s proof inherits this, meaning that they show the compiled value of any nonlocal game converges to the quantum (commuting operator) value, but that they cannot quantitatively say how large the security parameter needs to be before the compiled value is within any given epsilon of the true value. Nonetheless, the result is the first to obtain a generic bound on the compiled quantum value.

2.1.4 John Wright: The Unitary Synthesis Problem

John Wright gave about progress on the unitary synthesis problem. Almost any computational task in quantum computation can be performed by implementing *some* unitary transformation, whether it is preparing a state, performing a measurement, or mapping between two states. The task of implementing a given n-qubit unitary is known as unitary synthesis. In classical computation, while every transformation can be understood as a mapping from $\{0,1\}^n \mapsto \{0,1\}^m$, they can be efficiently reduced to implementing some Boolean function, $f : \{0,1\}^n \mapsto \{0,1\}$. The *unitary synthesis problem* [2] asks whether all quantum transformations can also efficiently reduced to implementing a Boolean function. Formally, the unitary synthesis problems is the following: Does there exist a fixed query algorithm $\mathcal{A}^{(\cdot)}$ such that for every unitary U, there exists a function f such that $\mathcal{A}^f$ implements U (up to ϵ-diamond norm distance)? The talk was about a recent result [31] proving that any adversary solving the unitary synthesis problem must make at least 2 queries to a Boolean function.

While this bound may seem relatively weak, it has surprising implications for quantum cryptography. For example, it indicates that it might be possible to construct quantum cryptography in a world where all classical computation is easy (even undecidable languages like the halting problem). Formally, it implies that, relative to a random oracle, there is a construction of $1\mathrm{PRS}$ that is secure against adversaries that can make a single query to an arbitrarily powerful classical oracle, even one that itself can access the random oracle.

At a very high level, the proof reduces the problem of ruling out algorithms for unitary synthesis to a matrix concentration inequality. Here is a sketch of the proof. Consider a hypothetical algorithm, $\mathcal{A}$, for unitary synthesis that only makes a single query to some Boolean function, and implements any unitary U (where the algorithm is fixed and the function depends on U). Say that the unitary is on n qubits, and let $N = 2^n$. Then we can use this algorithm to distinguish between the following two distribution of states:

1. $|\psi\rangle$ sampled by sampling a random $k \in [\sqrt{N}]$ and outputting a binary phase state corresponding to a function r_k, $|\psi_{r_k}\rangle = \frac{1}{\sqrt{N}} \sum_x (-1)^{r_k(x)} |x\rangle$.

2. $|\psi\rangle$ is a random binary phase state, $|\psi_h\rangle = \frac{1}{\sqrt{N}} \sum_x (-1)^{h(x)} |x\rangle$.

If there is an algorithm for unitary synthesis, then the a distinguisher (with an appropriate choice of function) can synthesis the unitary that maps $\frac{1}{\sqrt{N}} \sum_x (-1)^{r_k(x)} |x\rangle$ to $|k\rangle$, and check whether the result is one of the first $\sqrt{N}$ basis vectors.

The next part of the proof involved expressing the distinguishing advantage of any algorithm (querying f) into an operator norm bound on a difference of two random matrices (one depending on r_k, the other depending on the random function h). This can most easily be seen when the adversary is modeled as just making a single query to the oracle without any ancilla and measuring a projector Π. In this case, we can see that the probability that the adversary accepts when run on input $|\psi_g\rangle$ (and querying an arbitrary function f) can be written as

$$\begin{aligned}
\langle \psi_g | \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f | \psi_g \rangle &= \langle +^n | \mathcal{O}_g \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot \mathcal{O}_g | +^n \rangle \\
&= \langle +^n | \mathcal{O}_f \cdot \mathcal{O}_g \cdot \Pi \cdot \mathcal{O}_g \cdot \mathcal{O}_f | +^n \rangle \\
&= \langle \psi_f | \mathcal{O}_g \cdot \Pi \cdot \mathcal{O}_g | \psi_f \rangle \,.
\end{aligned}$$

Then the difference between the accepting probability when the adversary is given a state from either of the two distributions can be written as the operator norm of the difference of two random matrices, where the randomness is over the choice of r_k or h. Applying a concentration inequality for Rademacher random matrices (i.e. those whose entries are random ± 1) completes the proof.

2.1.5 Aparna Gupte: One-time Programs

Aparna Gupte gave a talk on recent work on quantum one-time programs [25]. The notion of one-time programs, first proposed by Goldwasser, Kalai and Rothblum [23], allows us to compile a program into one that can be run on a single input of a user's choice, but only one. If realizable, one-time programs would have wide-ranging applications in software protection, digital rights management, electronic tokens and electronic cash. Unfortunately, one-time programs immediately run into a fundamental barrier: software can be copied multiple times at will, and therefore, if it can be run on a single input of a user's choice, it can also be run on as many inputs as desired.

Previous work has circuvented this barrier by assuming the existence of specialized stateful hardware devices called one-time memories, which can be boosted into general one-time programs [23, 24], but the security of these schemes rests on shaky grounds: security relies on how much one is willing to trust the impenetrability of these hardware devices against side-channel attacks.

One might hope that the quantum no-cloning theorem might give us a solution. The no-cloning theorem states that quantum information cannot be generically copied, so if one can encode the given program into an appropriate quantum state, one might expect to circumvent the barrier. However, there is a simple impossibility result by Broadbent, Gutoski and Stebila [18] that rules out quantum one-time versions of any *deterministic* program. Indeed, given a candidate quantum one-time program state $|\psi_f\rangle$, an adversary can evaluate f many times on different inputs as follows: it first evaluates the program on some input x, measures the output register to obtain $f(x)$. Since f is deterministic, the measurement does not disturb the state of the program at all (if the computation is perfectly correct). The adversary then uncomputes the first evaluation, restoring the initial program state. She can repeat this process on as many inputs as she wishes. Morever, [18] also rule out one-time programs for any quantum channel (not just classical deterministic functions) for a very strong notion of one-time program security.

While these impossibility result rules out one-time programs for a very strong notion of security and for deterministic functionalities, they raise the following natural questions:

Can we obtain one-time programs for randomized functionalities?
Are there weaker yet meaningful notions of one-time security that can be achieved without hardware assumptions?

This work [25] constructs one-time programs for randomized classical programs. More interestingly, they give a simulation-style security definition that bypasses the [18] impossibility in that it is both strong and meaningful, and achievable. We give an overview of the construction and security definition.

Construction with Classical Oracles. The construction in the oracle model is inspired by the use of the "hidden subspace states" in the literatures of quantum money [4], signature tokens and quantum copy protection [9, 20]. A subspace state $|A\rangle$ is a uniform superposition over all vectors in some randomly chosen, secret subspace $A \subset \mathbb{F}_2^\lambda$. Specifically, $|A\rangle \propto \sum_{v \in A} |v\rangle$, where dimension of A is $\lambda/2$ and λ is the security parameter. These parameters ensure that A has exponentially many elements but is still exponentially small compared to the entire space.

At a high-level, our one-time scheme requires an authorized user to query an oracle on subspace vectors of A or its dual subspace $A^\perp$. Let f be the function we want to one-time program. Consider the simple case where x is a single bit in $\{0,1\}$. Let H be a random oracle. The one-time program consists of a copy of the subspace state $|A\rangle$ along with access to the following classical oracle:

$$O(x,v) = \begin{cases} f(x, G(v)) & \text{if } x = 0, v \in A \\ f(x, G(v)) & \text{if } x = 1, v \in A^\perp \\ \perp & \text{otherwise} \end{cases}.$$

To evaluate on input x, an honest user will measure the state $|A\rangle$ to obtain a uniform random vector in subspace A, if $x = 0$; or apply a binary QFT to $|A\rangle$ and measure to obtain a uniform random vector in the dual subspace $A^\perp$, if $x = 1$. It then inputs (x, v) into the oracle O and will obtain the evaluation $O(x, G(v))$ where the randomness $G(v)$ is uniformly random after putting the subspace vector into the random oracle.

For security, we leverage an "unclonability" property of the state $|A\rangle$ ([9, 6]) called "direct-product hardness": an adversary, given one copy of $|A\rangle$, polynomially bounded in query to the above oracle should not be able to produce two vectors v, v' which satisfy either of the following: (1) $v \in A, v' \in A^\perp$; (2) $v, v' \in A$ or $v, v' \in A^\perp$ but $v \neq v'$.

Security Definitions. A minimal definition of one-time security, which [25] call "operational security" is that no adversary, is able to produce two correct input-output sample pairs (x_1, y_1), (x_2, y_2). However, this security definition might not be sufficient in some settings; even if the adversary is not able to output two full correct outputs, it might learn some other secrets about the program that might pose a breach in security. For example, even if a one-time program is secure in this "operational sense", an adversary could output one correct output, and half of a second output, which might be undesirable leakage.

When confronted with such a situations, cryptographers resort to the simulation security paradigm to capture the idea that the adversary should not learn anything other than one output corresponding to an input of its choice. Indeed, such a definition is formalized, and subsequently ruled out, by Broadbent, Gutoski, and Stebila [18].

All these impossibilities arise from the fact that we cannot prevent the adversary from making gentle queries and uncomputing as many times as it wants, and making at most one destructive query. The key insight is that if the adversary does not make a destructive measurement, it does not learn anything useful, if the program was "random enough". This motivates the definition of the *single-effective-query* model: consider an oracle O_{SEQ}, that allows at most one destructive query, and an unlimited number of gentle queries and uncomputations. Then, the SEQ security definition says that no adversary can learn more from a secure one-time program than it could given access to this oracle O_{SEQ}. The SEQ oracle is defined using the compressed oracle technique of [39] to keep track of the number of destructive queries made.

This definition is not only achievable (the construction outlined above achieves this this security notion) but also *meaningful*—for "unlearnable" randomized programs, SEQ simulation security implies the operational security notion alluded to above.

Using this framework, [25] list a few applications for one-time programs of randomized functionalities, but leave for future work more applications.

2.1.6 Fang Song: Pseudorandom States

Some central primitives in classical cryptography fall under the umbrella of *pseudorandomness*: how can we build classical objects (e.g. long strings) that "appear" more random than they truly are? Two such notions are those of pseudorandom generators (PRGs) [12] and pseudorandom functions (PRFs) [22]. Since their introduction, these have proven to be extremely central cryptographic primitives. Moreover, these notions naturally carry across to the world of post-quantum security [40].

However, in a quantum world this question can naturally be extended one step further: can we define and construct *quantum* pseudorandom objects? In other words, can we build quantum states or unitaries that "appear" more random than they truly are? This direction was initiated by Ji, Liu, and Song [27], and Fang Song gave a talk at this workshop surveying the developments in this direction. Quantum pseudorandom primitives can be classified along two axes: they can either be states (PRS) or unitaries (PRU), and their pseudorandomness can be either statistical or computational. In the computational case, a computationally bounded adversary should not be able to distinguish an a priori unbounded polynomial number of copies of the pseudorandom state/unitary from a truly Haar random state/unitary. In the statistical case, even an *all-powerful* adversary should not be able to distinguish a *fixed* (i.e. a priori bounded) number of copies of the pseudorandom state/unitary from a truly Haar random state/unitary. These notions have many applications in quantum information and cryptography, such as private-key quantum money, encryption, authentication, and randomized benchmarking of NISQ circuits.

Ji, Liu, and Song [27] put forth two techniques for constructing quantum pseudorandom states. The first of these is the *random phase* construction: starting from the superposition $\sum_{x\in\{0,1\}^n} |x\rangle$, apply a pseudo-

random phase depending on x in superposition to obtain a state

$$|\psi_k\rangle \propto \sum_{x\in\{0,1\}^n} (-1)^{F_k(x)} |x\rangle ,$$

where $\{F_k\}$ is a PRF. Ji, Liu, and Song initially proposed and proved security of a slightly more complicated variant of this proposal, but soon afterwards Brakerski and Shmueli [17] proved that the ensemble $\{\psi_k\}$ is a pseudorandom family of states.

The second candidate technique proposed by Ji, Liu, and Song [27] is a *random subset* construction. For a key $k \in \{0,1\}^n$ and pseudorandom permutation $P_k : \{0,1\}^{2n} \to \{0,1\}^{2n}$, they define

$$|\psi_k\rangle \propto \sum_{x\in\{0,1\}^n} |P_k(x||0^n)\rangle ,$$

which can be equivalently viewed as a uniform superposition over a uniformly random subset of $\{0,1\}^{2n}$ of size $\{0,1\}^n$.

The initial study by Ji, Liu, and Song also defined pseudorandom unitaries and proposed candidate constructions of pseudorandom unitaries. This constructions consist of interleaving diagonal phase unitaries with the tensored Hadamard gate $H^{\otimes n}$ (or a quantum Fourier transform gate). Intriguingly, the security of this original proposal remains open. However, there have been several exciting developments on the front of constructing pseudorandom unitaries and related objects in the past year, which Song surveyed at the end of his talk. Many of these constructions build in various ways on the ideas contained in the PRS proposals listed above:

- A construction with provable security [35, 34] (which was the subject of the subsequent talk by Fermi Ma, summarized in Section 2.1.7);
- A candidate proposal that may be secure even without the existence of secure classical cryptography [13];
- The definition and construction of pseudorandom quantum *isometries* [5];
- Somewhat pseudorandom unitaries with real entries (rather than complex) [16]; and
- Another construction of PRUs based on random walks [32, 33].

2.1.7 Fermi Ma: Pseudorandom Unitaries

Fermi Ma gave a talk on his recent work with Hsin-Yuan Huang on constructing pseudorandom unitaries [34], as well as a detailed description of the underlying path recording technique. A pseudorandom unitary (PRU), as described as well in the previous talk by Fang Song, is a notion of quantum pseudorandomness. Formally, a PRU is a family of unitary operations for which query access to a random member is indistinguishable from query access to a true Haar random unitary.

Before this work, a number of constructions have been proposed, the most notable of which is the PFC construction, which was proven secure in a weaker (non-adaptive) setting [35]. This construction works by first applying a random Clifford (C), then applying a random phase oracle (F), and finally a random permutation (P). In this work, the authors show that the PFC construction does in fact satisfy full, adaptive security. As a corollary, if one-way functions exist, then so do PRUs. In addition, the authors use one-way functions to construct what they call a "strong" PRU, a PRU secure against queries to the unitary's inverse.

The key idea behind these theorems is a new technique for analyzing the behavior of a Haar random oracle, which the authors call "path-recording". Path-recording observes that query access to the A register of the isometry

$$\mathsf{prO} : |x\rangle_A |D\rangle_S \to \frac{1}{\sqrt{N-|D|}} \sum_{y\notin D_Y} |y\rangle |D \cup \{(x,y)\}\rangle$$

(where register S is initialized with the empty set) is statistically close to query access to a Haar random unitary. This is essentially a version of "compressed oracles" [39] (a useful tool for analyzing random oracles in the quantum setting) that can be used to analyze random unitaries.

The remainder of the talk was spent describing the technical details necessary to show that the path-recording oracle is indistinguishable from a Haar unitary. The main idea is that queries to the "PF" part of the PFC construction are indistinguishable from the path-recording oracle as long as the queries are sufficiently random. And so concatenating a random unitary with the PF construction looks like concatenating a random unitary with the path-recording oracle, which is the same as just giving out the path-recording oracle. In fact, the same argument works if a Clifford is used instead of a random unitary, and so the PFC construction is indistinguishable from the path-recording oracle, and thus also a Haar unitary.

2.1.8 James Bartusek: Quantum Obfuscation

Major advancements have been made by studying the idea of obfuscating classical programs through the lens of a particular formalization, namely indistinguishability obfuscation (IO). Classically, IO is "crypto-complete", in that it (essentially) implies the existence of all the other cryptographic functionalities we care about. While we now have candidates for obfuscating classical circuits, the space of obfuscating quantum circuits is relatively less explored. A line of work has constructed obfuscation for a limited class of quantum circuits, assuming the existence of classical circuit obfuscation. James Bartusek gave a talk on two different approaches to build obfuscation for quantum circuits.

The first approach, taken by [8], is to build on quantum fully-homomorphic encryption (QFHE), which almost gets us what we want, except that the output is encrypted. To allow the user to obtain the output of the obfuscated program in the clear, [8] also provide an oracle that decrypts honestly generated ciphertexts.

The second approach, taken by [7], is a more first-principles-style approach, which breaks the circuit into many layers, and ensures that the only way to go from one layer to the next is to evaluate the circuit honestly at each step. This approach uses authentication schemes for quantum states, that at their heart, are error-correcting codes.

Both these approaches achieve security only for quantum circuits that are pseudodeterministic, that is, they approximately implement a deterministic map from classical strings to classical strings. An open problems is to generalize the construction and security to more general types of quantum circuits.

2.1.9 William Kretschmer: The Complexity of Breaking Quantum Cryptography

William Kretschmer gave a talk about the complexity of breaking quantum cryptography. In classical cryptography, it is known that the existence of one-way functions is straight forwardly related to the question of whether or not P equals NP. However, recent developments have provided evidence that forms of quantum cryptography can exist even if $\mathsf{P} = \mathsf{NP}$ [28]. Instead, the complexity of breaking various quantum cryptographic primitives can be broken into three distinct categories.

1. Primitives that are broken if $\mathsf{BQP} \supseteq \mathsf{QCMA}$.
2. Primitives that are broken if $\mathsf{BQP} \supseteq \mathsf{PP}$
3. Primitives that have no known upper bound on their complexity.

For example, quantum-implementable one-way functions can be broken in QCMA, whereas many "microcrypt" primitives like pseudo-random states and unitaries can only be broken using shadow tomography, which requires the power of PP. On the other hand, primitives like $1\mathsf{PRS}$ and EFI pairs have no known upper bound on their complexity, and the work of [31] showed that it is possible that even an oracle to the halting problem (or any arbitrarily powerful classical oracle) would not be enough to break the security of these primitives in general. William talked about work on exploiting the fact that quantum primitives require different kinds of complexity to break to find oracles relative to which $\mathsf{BQP} = \mathsf{QCMA}$ (i.e. one-way functions do not exist), but pseudo-random states exist [29], and for which $\mathsf{P} = \mathsf{NP}$ (i.e. classically implementable one-way functions do not exist) but quantumly implementable one-way functions exist [30].

2.1.10 Kabir Tomer: Quantum Complexity ∩ Quantum Cryptography

Outside of a very limited set of primitives, almost no cryptography can be secure against unbounded adversaries. Since we must therefore settle for security against computationally bounded adversaries, cryptography

fundamentally relies on computational assumptions. In the classical setting the minimum required assumption is the existence of one-way functions, which requires that $\mathsf{P} \neq \mathsf{NP}$.

In the quantum setting, however, the picture is somewhat different. Certain primitives like quantum key distribution can achieve security against unbounded adversaries. Additionally, it is possible to reduce the assumptions required for cryptographic primitives using quantum information. In particular, some recent works have shown that in the quantum setting, one-way functions suffice for building secure multi-party computation and (certain forms of) public-key encryption. Building these primitives from one-way functions in a black box manner is known to be impossible in the classical setting. Some recent exciting results have shown that it may be possible to weaken these assumptions even further. These results show that relative to certain oracles a primitive known as a pseudorandom state generator (PRS) can exist even if $\mathsf{BQP} = \mathsf{QMA}$ or $\mathsf{P} = \mathsf{NP}$. Additionally, PRS can be used to build quantum commitments, which in turn suffice to build powerful primitives like multi-party computation and zero-knowledge proofs. This world of cryptography below one-way functions is often referred to as Microcrypt.

This talk focused on some of the lenses through which the community has sought to understand Microcrypt. The first has been to study analogues of one-wayness as a natural form of computational hardness one may expect to encounter in the quantum setting. The first one-way primitive to be introduced was a natural relaxation of PRS called a One-Way State Generator (OWSG), which is similar to a one-way function except that the output of the function is a quantum state. Subsequently, other forms of one-wayness like quantum-computable one-way functions, one-way puzzles, and state-puzzles have been introduced in order to capture the one-wayness inherent in a large variety of primitives. All these forms of one-wayness have been shown to imply one-way puzzles (with the exception on mixed-state one-way state generators), which additionally can be used to build quantum commitments.

Another important approach has been to construct Microcrypt primitives from the hardness of concrete mathematical problems. There have been some promising candidates for such problems, but there is only limited evidence that breaking their security guarantees is harder than inverting one-way functions. This property is in general extremely challenging to argue, let alone prove. A recent work proposed a candidate construction of one-way puzzles (and hence quantum commitments) based on the average-case hardness of approximating the probabilities of outputs of certain quantum circuits (namely, the circuits used in BosonSampling, Instantaneous Quantum Polynomial-Time (IQP) circuits, or Random Circuit Sampling). This estimation problem has been conjectured in the literature on quantum advantage to be $\#\mathsf{P}$-hard, and there is a long line of works that aim to support this conjecture through complexity-theoretic evidence and worst-case to average-case reductions. The authors show that this estimation problem reduces to a distributional inversion task, giving rise to distributional one-way puzzles, which are a variant of one-way puzzles that have been shown to be equivalent to standard one-way puzzles. The construction is secure even if the estimation problem is not $\#\mathsf{P}$-hard but merely infeasible for QPT adversaries. However, if the problem is infeasible for $\mathsf{QPT}^{\mathsf{NP}}$ adversaries (which is implied by the $\#\mathsf{P}$-hardness conjecture and the mild assumption that $\mathsf{P}^{\#\mathsf{P}} \not\subseteq \mathsf{BQP}^{\mathsf{NP}}$) then the problem is hardness than inverting one-way functions (since one-way functions can be inverted by $\mathsf{QPT}^{\mathsf{NP}}$ adversaries). This gives a new class of cryptographic assumptions which under plausible conjectures are significantly milder than the existence of one-way functions. This provides compelling concrete evidence for cryptography from weaker assumptions than the existence of one-way functions.

Finally, the talk outlined several open questions in this area, including whether commitments and one-way puzzles are equivalent/separated, whether one-way puzzles can be used to build primitives beyond just commitments, and whether other concrete candidates can be found for other Microcrypt primitives.

2.2 Summary of Breakout Sessions

Participants not giving an plenary lecture were invited to give talks in a number of breakout sessions in the afternoons. Over these breakout sessions, the following talks were given:

1. Omri Shmueli: "Non-collapsing CRHs and one-shot signatures"
2. Tomoyuki Morimae: "Quantum advantage and quantum cryptography"
3. Eric Culf: "Unclonable encryption from decoupling"
4. John Bostanci: "Quantum money from non-Abelian group actions"

5. Kabir Tomer and Luowen Qian: "One-way puzzles = hardness of state synthesis/extrapolation"
6. Open: "Naming conventions in quantum cryptography"
7. Christian Majenz: "Quantum-accessible random permutation oracles"
8. Eli Goldin: "Metacomplexity and quantum cryptography"
9. Seyoon Ragavan: "Quantum factoring: the state of the art"
10. Harriet Apel: "Connections between cryptography and holography"
11. Or Sattath: "Quantum assumptions"

3 Scientific Progress Made

To give an indication of the progress made at the meeting, we summarize open problems that were discussed during the workshop. Progress was made in informal discussions, in the form of both new and continuing collaborations.

3.1 Quantum Lightning from Lattices

Constructing quantum money and lightning from "standard" post-quantum cryptographic assumptions is a major open problem in quantum cryptography. One such assumption is the learning with errors (LWE) problem, which involves distinguishing random vectors from noisy lattice elements. Possible constructions of quantum money, lightning, and one-shot signatures from lattices were discussed.

3.2 Private Unitary and State Synthesis

Private state (or unitary) synthesis is defined to be a (computationally unbounded) algorithm for mapping a state (or unitary) to a classical oracle such that there is an efficient synthesizer who can generate a copy of the state (or implement the unitary), but for which there exists a simulator that, given copies of the quantum state (or access to the unitary) can "spoof" queries to the function. Possible constructions of private state synthesis were discussed.

3.3 Universal constructions of cryptographic primitives

A universal construction of a primitive is one that satisfies security if *any* construction of the primitive is secure. One idea for constructing a universal pseudo-random unitary is to take a random quantum circuit and show that any other construction appears in the random quantum circuit with high probability (a similar idea appeared in [19]). During the workshop, attempts to extend the ideas in [19] to psuedo-random unitaries were discussed.

3.4 Idealized quantum primitives

Random functions are typically treated as an "ideal" hash function, and a question that was considered classically was whether other ideal functionality (for example, a truly random permutation) could be constructed from a random function. The idea of "indifferentiability" allows for these kinds of constructions to be composed with each other. In quantum cryptography, some ideal primitives include query access to a Haar random unitary, query access to 2^n many Haar random unitaries, and copies of Haar random quantum states. Possible constructions of ideal primitives from other ideal primitives were discussed.

3.5 Post-quantum parallel repetition

Parallel repetition is probably the most natural way to amplify the security of cryptographic primitives, but in the post-quantum and fully-quantum settings, it has only recently been shown to work [14, 26]. Today, it is not known whether tight parallel repetition is possible for public-coin post-quantum arguments, or whether modifying the challenger to randomly terminate helps with parallel repetition. Ideas for achieving better parallel repetition theorems for post-quantum arguments were discussed.

4 Outcome of the Meeting

The meeting was a resounding success, with all of the researchers agreeing that it was a positive and productive experience. Many of the participants came into the workshop knowing some, but not all, of the other researchers. Through the talks, breakout sessions, and social events (hikes, meals, board games), the participants were brought closer and many new connections and friendships were made.

The workshop was beneficial for both senior and junior researchers. Senior researchers appreciated the opportunity to meet new researchers and students, as well as to hear new ideas and directions in the field of quantum cryptography. Junior researchers took advantage of the opportunity to tap into the breadth of knowledge that the senior researchers brought to the workshop, as well as the chance to situate their research into the broader picture of quantum cryptography. The workshop was an amazing opportunity to mix ideas from different sub-branches of quantum cryptography, as well as discuss challenges, potential solutions, and ideas among participants, topics that are often omitted from finished research papers.

Fostering an environment of encouragement and empowerment for early-career researchers was an important part of the workshop design. Of the 11 plenary talk sessions, 7 of them were given by early career researchers (graduate students or postdocs). Nearly all of the breakout sessions were led by students or postdocs. Given that the branch of "fully quantum" cryptography is relatively new and young, it was fitting to have early career researchers be front and center in the workshop program.

Through breakout sessions and unstructured time, the workshop facilitated many discussions and led to progress in both new and old collaborations. We look forward to seeing the progress that will come from this event, hopefully at a future iteration of this workshop.

Acknowledgments

The organizers (Zvika Brakerski, Yael Kalai, Henry Yuen) thank BIRS for hosting this workshop. The bulk of this report was prepared by PhD students John Bostanci (Columbia), Eli Goldin (NYU), Aparna Gupte (MIT), Seyoon Raghavan (MIT), Kabir Tomer (UIUC), Tina Zhang (MIT). Everything praiseworthy about the report is due to them, and any errors or omissions are due to the organizers.

References

[1] Aaronson, S. Limitations of quantum advice and one-way communication. *Proceedings. 19th IEEE Annual Conference On Computational Complexity, 2004.*. pp. 320-332 (2004)

[2] S. Aaronson and G. Kuperberg, Quantum versus classical proofs and advice. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, IEEE (2007), 115—128

[3] Aaronson, S. (2009, July). Quantum copy-protection and quantum money. In 2009 24th Annual IEEE Conference on Computational Complexity (pp. 229-242). IEEE.

[4] Aaronson, S. & Christiano, P. Quantum money from hidden subspaces. *Proceedings Of The Forty-fourth Annual ACM Symposium On Theory Of Computing*. pp. 41-60 (2012)

[5] Ananth, P., Gulati, A., Kaleoglu, F. & Lin, Y. Pseudorandom Isometries. *Advances In Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference On The Theory And Applications Of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part IV*. **14654** pp. 226-254 (2024), https://doi.org/10.1007/978-3-031-58737-5

[6] Bartusek, J., Kitagawa, F., Nishimaki, R. & Yamakawa, T. Obfuscation of pseudo-deterministic quantum circuits. *Proceedings Of The 55th Annual ACM Symposium On Theory Of Computing*. pp. 1567-1578 (2023)

[7] Bartusek J, Brakerski Z, Vaikuntanathan V. Quantum state obfuscation from classical oracles. InProceedings of the 56th Annual ACM Symposium on Theory of Computing 2024 Jun 10 (pp. 1009-1017).

[8] Bartusek J, Kitagawa F, Nishimaki R, Yamakawa T. Obfuscation of pseudo-deterministic quantum circuits. InProceedings of the 55th Annual ACM Symposium on Theory of Computing 2023 Jun 2 (pp. 1567-1578).

[9] Ben-David, S. & Sattath, O. Quantum tokens for digital signatures. *Quantum*. **7** pp. 901 (2023)

[10] Bennett, Charles H., and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. Theoretical computer science 560 (2014): 7-11.

[11] Bhattacharyya, A., & Culf, E. (2025). Uncloneable Encryption from Decoupling. arXiv preprint arXiv:2503.19125.

[12] Blum, M. & Micali, S. How to Generate Cryptographically Strong Sequences of Pseudorandom Bits. *SIAM Journal On Computing*. **13**, 850-864 (1984)

[13] Bostanci, J., Haferkamp, J., Hangleiter, D. & Poremba, A. Efficient Quantum Pseudorandomness from Hamiltonian Phase States. *CoRR*. **abs/2410.08073** (2024), https://doi.org/10.48550/arXiv.2410.08073

[14] Bostanci, J., Qian, L., Spooner, N., & Yuen, H. (2024, June). An efficient quantum parallel repetition theorem and applications. In Proceedings of the 56th Annual ACM Symposium on Theory of Computing (pp. 1478-1487).

[15] Bostanci, J., Nehoran, B., & Zhandry, M. (2025, June). A general quantum duality for representations of groups with applications to quantum money, lightning, and fire. In Proceedings of the 57th Annual ACM Symposium on Theory of Computing (pp. 201-212).

[16] Brakerski, Z. & Magrafta, N. Real-Valued Somewhat-Pseudorandom Unitaries. *Theory Of Cryptography - 22nd International Conference, TCC 2024, Milan, Italy, December 2-6, 2024, Proceedings, Part II*. **15365** pp. 36-59 (2024), https://doi.org/10.1007/978-3-031-78017-2

[17] Brakerski, Z. & Shmueli, O. (Pseudo) Random Quantum States with Binary Phase. *Theory Of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*. **11891** pp. 229-250 (2019), https://doi.org/10.1007/978-3-030-36030-6

[18] Broadbent, A., Gutoski, G. & Stebila, D. Quantum one-time programs. *Annual Cryptology Conference*. pp. 344-360 (2013)

[19] Chen, C. F., Haah, J., Haferkamp, J., Liu, Y., Metger, T., & Tan, X. (2024). Incompressibility and spectral gaps of random circuits. arXiv preprint arXiv:2406.07478.

[20] Coladangelo, A., Liu, J., Liu, Q. & Zhandry, M. Hidden cosets and applications to unclonable cryptography. *Advances In Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*. pp. 556-584 (2021)

[21] Farhi, E., Gosset, D., Hassidim, A., Lutomirski, A., & Shor, P. (2012, January). Quantum money from knots. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (pp. 276-289).

[22] Goldreich, O., Goldwasser, S. & Micali, S. How to construct random functions. *J. ACM*. **33**, 792-807 (1986), https://doi.org/10.1145/6490.6503

[23] Goldwasser, S., Kalai, Y. & Rothblum, G. One-time programs. *Advances In Cryptology–CRYPTO 2008: 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings 28*. pp. 39-56 (2008)

[24] Goyal, V., Ishai, Y., Sahai, A., Venkatesan, R. & Wadia, A. Founding Cryptography on Tamper-Proof Hardware Tokens. *Theory Of Cryptography, 7th Theory Of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*. **5978** pp. 308-326 (2010), https://doi.org/10.1007/978-3-642-11799-2

[25] Gupte, A., Liu, J., Raizes, J., Roberts, B. & Vaikuntanathan, V. Quantum one-time programs, revisited. *ArXiv Preprint ArXiv:2411.01876*. (2024)

[26] Huang, A., & Kalai, Y. T. (2025). Parallel Repetition for Post-Quantum Arguments. arXiv preprint arXiv:2506.02277.

[27] Ji, Z., Liu, Y. & Song, F. Pseudorandom Quantum States. *Advances In Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*. **10993** pp. 126-152 (2018), https://doi.org/10.1007/978-3-319-96878-0

[28] Kretschmer, W. (2021). Quantum pseudorandomness and classical complexity. arXiv preprint arXiv:2103.09320.

[29] Kretschmer, W., Qian, L., Sinha, M., & Tal, A. (2023, June). Quantum cryptography in algorithmica. In Proceedings of the 55th Annual ACM Symposium on Theory of Computing (pp. 1589-1602).

[30] Kretschmer, W., Qian, L., & Tal, A. (2025, June). Quantum-computable one-way functions without one-way functions. In Proceedings of the 57th Annual ACM Symposium on Theory of Computing (pp. 189-200).

[31] A. Lombardi, F. Ma, and J. Wright, A one-query lower bound for unitary synthesis and breaking quantum cryptography, In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, Vanouver, BC, CA (2024), 979–990

[32] Lu, C., Qin, M., Song, F., Yao, P. & Zhao, M. Quantum Pseudorandom Scramblers. *Theory Of Cryptography - 22nd International Conference, TCC 2024, Milan, Italy, December 2-6, 2024, Proceedings, Part II*. **15365** pp. 3-35 (2024), https://doi.org/10.1007/978-3-031-78017-2

[33] Lu, C., Qin, M., Song, F., Yao, P. & Zhao, M. Parallel Kac's Walk Generates PRU. *CoRR*. **abs/2504.14957** (2025), https://doi.org/10.48550/arXiv.2504.14957

[34] F. Ma and H. Huang, How to Construct Random Unitaries. https://arxiv.org/abs/2410.10116.

[35] Metger, T., Poremba, A., Sinha, M. & Yuen, H. Simple Constructions of Linear-Depth t-Designs and Pseudorandom Unitaries. *65th IEEE Annual Symposium On Foundations Of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*. pp. 485-492 (2024), https://doi.org/10.1109/FOCS61266.2024.00038

[36] C. Shannon, The synthesis of two-terminal switching circuits, In *The Bell System Technical Journal* **28.1** (1949), 59-–98

[37] P. W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, Santa Fe, NM, USA (1994), 124–134

[38] Stephen Wiesner. Conjugate coding. SIGACT News 15, 1 (Winter-Spring 1983), 78–88. https://doi.org/10.1145/1008908.1008920

[39] Zhandry, M. How to Record Quantum Queries, and Applications to Quantum Indifferentiability. *Advances In Cryptology – CRYPTO 2019*. pp. 239-268 (2019)

[40] Zhandry, M. How to Construct Quantum Random Functions. *J. ACM*. **68**, 33:1-33:43 (2021), https://doi.org/10.1145/3450745

[41] Zhandry, M. (2021). Quantum lightning never strikes the same state twice. or: quantum money from cryptographic assumptions. Journal of Cryptology, 34, 1-56.

[42] Zhandry, M. (2023). Quantum money from abelian group actions. arXiv preprint arXiv:2307.12120.