Mathematisches Forschungsinstitut Oberwolfach

Report No. 3/2025

DOI: 10.4171/OWR/2025/3

Cryptography

Organized by
Johannes Buchmann, Darmstadt
Shafi Goldwasser, Berkeley
Yael Kalai, Cambridge MA
Vinod Vaikuntanathan, Cambridge MA

19 January – 24 January 2025

Abstract. The science of cryptography lies at the very foundations of trust in current and future systems and devices for communication and computation. For example, secure communications over the internet today are possible thanks to mathematical work in cryptography in the 1970s and 80s. Security and efficiency of cryptographic schemes rely heavily on mathematics: mathematical models describe what security means, hard algorithmic problems are the basis of constructions that achieve the desired security which is established by mathematical proofs, and algorithmic optimizations make the schemes applicable in practice. Dramatic changes in computational technologies raise new challenges, and therefore new opportunities, for cryptography. These challenges include the near-ubiquitous use of remote storage and computation including cloud computing, which bring to fore new concerns of privacy, security and integrity; the advances in quantum computation which necessitate the development of more robust mathematical foundations of the field; and a revolution in machine learning and artificial intelligence that is poised to affect our lives in fundamental ways and which once again bring up an entirely new suite of problems for cryptography, ones related to trust, security, integrity and fairness of these systems. In the last half decade, there have been spectacular advances in cryptography, in the areas of program obfuscation, verifiable computation, elliptic curves and isogenies, lattice-based cryptography, quantum cryptography, cryptographic techniques in machine learning, and more. This workshop brought together experts from mathematics and computer science such as algorithmic number theory and algebra, quantum computation and complexity theory, in order to discuss recent advances and make progress in constructing the new generation of cryptographic systems that protect the future of information and computation.

Mathematics Subject Classification (2020): 94A60, 68P25, 81P94.

Introduction by the Organizers

The workshop on *cryptography*, organized by Jonannes Buchmann, Shafi Goldwasser, Yael Kalai and Vinod Vaikuntanathan, was a fantastic workshop with great talks and numerous fruitful collaborations and discussions. We had over 40 participants, with a healthy mixture of female and male participants, junior and senior participants, and a broad geographic representation.

This workshop addressed new and emerging questions of trust that arise in present and future information systems. Addressing these problems requires a joint effort of researchers with a wide range of expertise, ranging from algorithmic number theory and algebra, quantum computation, complexity theory, and theoretical machine learning. Indeed, as we elaborate on in this report, the workshop consisted of a diverse set of talks that covered all the areas mentioned above, and much more.

Description

Trust is paramount in the modern society driven by information and computation. Cryptography provides a mathematically rigorous and strong foundation of trust in information systems. For example, our internet communications are encrypted with protocols such as TLS (Transport Layer Security) which employ encryption and digital signature schemes. Cryptocurrencies employ cryptographic hash functions, digital signatures and zero-knowledge proofs of correctness. In fact, there will be no trust in current and future information systems without appropriate cryptography.

Mathematics is the lifeblood of modern cryptography. For starters, mathematical models are used to formally describe the security requirements for cryptographic algorithms (e.g. translating words such as security, privacy, integrity, fairness and knowledge into precise mathematical notions). Secondly, cryptographic constructions that meet these notions require complexity theory and the hardness of computational problems in mathematics. Classical examples are integer factorization and computing discrete logarithms in the group of points of an elliptic curve over a finite field, and more modern notions include finding short vectors in integer lattices, solving non-linear systems of equations over finite fields, and computing isogenies between elliptic curves. Cryptanalysis involves trying to design algorithms for the presumably hard problems, requiring deep mathematics. Finally, proving that the said constructions meet the said security notion is done through a complexity-theoretic reduction which in turn heavily uses mathematical tools.

Despite all these developments in mathematics-based cryptography, there are a growing number of important research problems that are seeing exciting growth and fresh new ideas in the recent years. The following topics (non-exclusively) were the focus of our workshop:

(1) **Program Obfuscation and Homomorphic Encryption:** Program obfuscation refers to the task of "scrambling" a software program in such

a way that it retains its functionality, namely its input-output behavior, yet hides all its inner workings. Back in the 1970s, Diffie and Hellman, in a work that is widely recognized as the genesis of modern cryptography, already recognized the importance of program obfuscation: they proposed creating a public-key encryption scheme by obfuscating the encryption algorithm of a private-key encryption scheme which contains the hardcoded private key (even though they then had no way of doing so). In modern times, the work of Barak et al. [BGI+12] proposed mathematical definitions of program obfuscation; Garg et al. [GGH+13, GGH+16] realized the first candidate construction; and following nearly a decade of research developing new tools and reductions [BV18, AJ15], the breakthrough works of Jain, Lin and Sahai [JLS21, JLS22] recently constructed ways of obfuscating programs under standard mathematical assumptions. The area is rife with exciting and difficult open questions including: (a) can we reduce the mathematical assumptions required for program obfuscation? (b) can we construct a program obfuscation scheme that is post-quantum secure, e.g. from hardness assumptions on integer lattices? and (c) can we construct obfuscators with a stronger security guarantee, referred to as virtual black-box obfuscation in the literature? A related notion is that of a fully homomorphic encryption scheme which allows us to compute on encrypted data: indeed, a program obfuscator immediately gives us a fully homomorphic encryption scheme. Whether one can construct a fully homomorphic encryption scheme from mathematical objects other than integer lattices remains a tantalizing open problem.

(2) **Verifiable Delegation of Computation:** Efficient verification of computation is one of the most fundamental problems in theoretical computer science, and is at the heart of the P versus NP problem. Recently, with the increasing popularity of blockchain technologies and cloud services, efficient verification schemes are increasingly deployed in practice. Such schemes provide a method for converting any proof into a "succinct computational proof," which is significantly shorter than the original proof though provides only *computational soundness.* Namely, computational proofs for false statements exist but finding them requires breaking a hardness assumptions (such as the computational hardness of factoring large numbers). A computationally sound proof is referred to as an argument, and a succinct (non-interactive) argument is referred to as a SNARG. Despite its growing use in practice, we only have heuristics for constructing SNARGs. Our holy grail is to construct a SNARG for any NP language under standard cryptographic assumptions, and this is one of the questions we propose to focus on in this workshop.

(3) **Cryptography and Quantum Computing.** The advancements in quantum computing raise many challenges and opportunities. For one, Shor's algorithm [Sho94], once implemented on a scalable quantum computer, will break existing public-key cryptosystems based on the hardness

of factoring or computing discrete logarithms. Thus, the first challenge is to construct new cryptographic schemes that are "post-quantum secure". Post-quantum security also includes ensuring that our security *proofs* extend to the post-quantum setting. Currently, our proof techniques include rewinding the adversary or running it many times. This cannot be done in general in the quantum setting since the "no cloning theorem" in quantum mechanics states that in general, quantum states cannot be replicated.

Beyond ensuring that our current schemes remain secure, quantum computing brings with it many new challenges. Envisioning a world where there will be some expensive quantum devices along with many cheaper classical devices, we need to build tools where a quantum device and a classical device will be able to interact securely and effectively. For example, how can a quantum device generate a classical commitment to its quantum state? How can a quantum device prove to a classical device that a quantum computation that it did was indeed correct [Mah22]?

Finally, quantum information enables us to achieve hitherto unimaginable tasks in cryptography such as quantum money, uncloneable programs and one-time programs [CLLZ21]. Quantum information also helps us construct cryptographic schemes that maintain security even when $\mathsf{P} = \mathsf{NP}$ [BB84]. One of goals in this workshop is to push the boundary of this interplay between classical and quantum cryptography.

(4) **New Mathematical Foundations: Lattices, Isogenies and Kolmogorov Complexity.** Cryptographers are constantly on the lookout for mathematical problems that are computationally hard yet possess enough structure to construct useful cryptographic objects such as public-key encryption, fully homomorphic encryption and program obfuscation schemes. At the same time, it is important to study algorithms that potentially break these hardness assumptions, a task that involves deep techniques from a range of areas in mathematics including algebraic geometry, (algebraic) number theory, combinatorics and complexity theory.

In addition to the above-mentioned areas, we included in the program emerging work at the intersection of machine learning and cryptography that, on the one hand, uses cryptography to insert undetectable backdoors in machine learning models [GKVZ22], but also shows how to watermark outputs from such models [CGZ23] and verify the training procedure of such models [GRSY21]. This is an exciting area with many important and consequential questions at the interface of the two fields.

Addressing all these problems required a joint effort of cryptographers with expertise in several areas of mathematics and computer science, such as algorithmic number theory and algebra, quantum computation, complexity theory, and theoretical machine learning. Our workshop participants included young and advanced researchers in these areas. Many fantastic talks were given (which are summarized below) and many new collaborations were fostered during this workshop, which we beleive will bring new insights and breakthroughs to the field of cryptography.

Acknowledgement: The MFO and the workshop organizers would like to thank the National Science Foundation for supporting the participation of junior researchers in the workshop by the grant DMS-2230648, "US Junior Oberwolfach Fellows".

Workshop: Cryptography

Table of Contents

Abstracts

Verifiable data science via interactive proofs and arguments

Guy Rothblum
(joint work with Noga Amir, Oded Goldreich, and Tal Herman)

Part I: Suppose we have access to a small number of samples from an unknown distribution, and would like to learn facts about the distribution. An untrusted data server claims to have studied the distribution and makes assertions about its properties. Can the untrusted data server prove that its assertions are approximately correct? Can a short efficiently verifiable proof be generated in polynomial time?

We study doubly-efficient interactive proof systems that can be used to verify properties of an unknown distribution over a domain $[N]$. On top of efficient verification, our focus is on proofs that the honest prover can generate in polynomial time. More generally, the complexity of generating the proof should be as close as possible to the complexity of simply running a standalone analysis to determine whether the distribution has the property.

Our main result is a new 2-message doubly-efficient interactive proof protocol for verifying any label-invariant distribution property (any property that is invariant to re-labeling of the elements in the domain of the distribution). The sample complexity, communication complexity and verifier runtime are all $\tilde{O}(\sqrt{N})$. The proof can be generated in quasi-linear $\tilde{O}(N)$ time and sample complexities (the runtimes of the verifier and the honest prover hold under a mild assumption about the property's computational complexity). This improves on prior work, where constructing the proof required super-polynomial time (Herman and Rothblum, STOC 2022).

Our new proof system is directly applicable to proving (and verifying) several natural and widely-studied properties, such as a distribution's support size, its Shannon entropy, and its distance from the uniform distribution. For these (and many other) properties, the runtime and sample complexities for generating the proof are within polylog(N) factors of the complexities for simply determining whether the property holds.

Part II: Suppose Alice has collected a small number of samples from an unknown distribution, and would like to learn about the distribution. Bob, an untrusted data analyst, claims that he ran a sophisticated data analysis on the distribution, and makes assertions about its properties. Can Alice efficiently verify Bob's claims using fewer resources (say in terms of samples and computation) than would be needed to run the analysis herself?

We construct [HR24] an interactive proof system for any distribution property that can be decided by uniform polynomial-size circuits of bounded depth: the circuit gets a complete description of the distribution and decides whether it has the property. Taking N to be an upper bound on the size of the distribution's support, the verifier's sample complexity, running time, and the communication

complexity are all sublinear in N: they are bounded by $\widetilde{O}(N^{1-\alpha} + D)$ for a constant $\alpha > 0$, where D is a bound on the depth of the circuits that decide the property. The honest prover runs in $\text{poly}(N)$ time and has quasi-linear sample complexity. Moreover, the proof system is tolerant: it can be used to approximate the distribution's distance from the property.

We show similar results for any distribution property that can be decided by a bounded-space Turing machine (that gets as input a complete description of the distribution). We remark that even for simple properties, deciding the property without a prover requires quasi-linear sample complexity and running time. Prior work [HR23] demonstrated sublinear interactive proof systems, but only for the much more restricted class of label-invariant distribution properties.

Part III: We initiate a study of doubly-efficient interactive proofs of proximity [AGR25], while focusing on properties that can be tested within query-complexity that is significantly sub-linear, and seeking interactive proofs of proximity in which:

(1) The query-complexity of verification is significantly smaller than the query-complexity of testing.
(2) The query-complexity of the honest prover strategy is not much larger than the query-complexity of testing.

We call such proof systems *doubly-sublinear IPPs* (dsIPPs).

We present a few doubly-sublinear IPPs. A salient feature of these IPPs is that the honest prover does not employ an optimal strategy. In particular, the honest prover in our IPP for sets recognizable by constant-width read-once oblivious branching programs uses a distance-approximator for such sets.

Additive Randomized Encodings: Part 1

Yuval Ishai

(joint work with Shai Halevi, Eyal Kushilevitz, and Tal Rabin)

A secure computation protocol for $f(x_1, \ldots, x_n)$ enables n parties to evaluate f on their local inputs x_i while hiding everything except the output. A useful special case, which is often easier to solve, is when f computes addition in a finite Abelian group G. Can we reduce the general case to this special case by first locally mapping each x_i to x'_i in G, and then securely computing the sum of all x'_i?

Such a reduction is captured by the abstract notion of *additive randomized encoding* (ARE) [HIKR23]. An ARE of $f(x_1, \ldots, x_n)$ is an n-tuple of randomized local mappings $g_i(x_i)$ whose sum reveals the output of f but hides (essentially) everything else about the inputs.

In this part, I will present positive results, negative results, and open questions about the existence of ARE with information-theoretic security.

I will also discuss several applications of ARE, including non-interactive secure computation protocol in the shuffle model, where parties can post messages on

an anonymous bulletin board. This implies a utility-preserving compiler from differential privacy in the central model to differential privacy in the shuffle model.

Additive Randomized Encodings: Part 2

Nir Bitansky
(joint work with Saroja Erabeli and Rachit Garg)

What are the minimal computational assumptions under which additive randomized encodings can be constructed?

In this part, I'll show a construction of (computationally-secure) ARE assuming public-key encryption [BEG25]. The key insight behind the construction is that one-sided ARE, which only guarantees privacy for one of the parties, are relatively easy to construct, and yet can be lifted to full-fledged ARE.

Better Garbling via Algebraic Assumptions

Rachel Lin
(joint work with Marshall Ball, Yuval Ishai, Hanjun Li, and Tianren Liu)

Part I: The beautiful work of Applebaum, Ishai, and Kushilevitz [FOCS'11] initiated the study of arithmetic variants of Yao's garbled circuits. An arithmetic garbling scheme is an efficient transformation that converts an arithmetic circuit $C : \mathcal{R}^n \to \mathcal{R}^m$ over a ring $\mathcal{R}$ into a garbled circuit $\widehat{C}$ and n affine functions L_i for $i \in [n]$, such that $\widehat{C}$ and $L_i(x_i)$ reveals only the output $C(x)$ and no other information of x. AIK presented the first arithmetic garbling scheme supporting computation over integers from a bounded (possibly exponentially large) range, based on Learning With Errors (LWE). In contrast, converting C into a Boolean circuit and applying Yao's garbled circuit treats the inputs as bit strings instead of ring elements, and hence is not "arithmetic".

In this work [BLLL23], we present new ways to garble arithmetic circuits, which improve the state-of-the-art on efficiency, modularity, and functionality. To measure efficiency, we define the *rate* of a garbling scheme as the maximal ratio between the bit-length of the garbled circuit $|\widehat{C}|$ and that of the computation tableau $|C|_\ell$ in the clear, where ℓ is the bit length of wire values (e.g., Yao's garbled circuit has rate $O(\lambda)$).

- We present the first *constant-rate* arithmetic garbled circuit for computation over large integers based on the Decisional Composite Residuosity (DCR) assumption, significantly improving the efficiency of the schemes of Applebaum, Ishai, and Kushilevitz.
- We construct an arithmetic garbling scheme for modular computation over $\mathcal{R} = \mathbb{Z}_p$ for any integer modulus p, based on either DCR or LWE. The DCR-based instantiation achieves rate $O(\lambda)$ for large p. Furthermore, our construction is modular and makes black-box use of the underlying ring and a simple key *extension gadget*.

- We describe a variant of the first scheme supporting arithmetic circuits over bounded integers that are augmented with Boolean computation (e.g., truncation of an integer value, and comparison between two values), while keeping the *constant rate* when garbling the arithmetic part.

To the best of our knowledge, constant-rate (Boolean or arithmetic) garbling was only achieved before using the powerful primitive of indistinguishability obfuscation, or for restricted circuits with small depth.

Part II: A *garbling scheme* transforms a program (e.g., circuit) C into a garbled program $\widehat{C}$, along with a pair of short keys $(k_{i,0}, k_{i,1})$ for each input bit x_i, such that $(C, \widehat{C}, \{k_{i,x_i}\})$ can be used to recover the output $z = C(x)$ while revealing nothing else about the input x. This can be naturally generalized to *partial garbling*, where part of the input is *public*, and a computation $z = C(x, y)$ is decomposed into a public part $C_{\text{pub}}(x)$, depending only on the public input x, and a private part $z = C_{\text{priv}}(C_{\text{pub}}(x), y)$ that also involves a private input y.

A key challenge in garbling is to achieve *succinctness*, where the size of the garbled program may grow only with the security parameter and (possibly) the output length, but not with the size of C. Prior work achieved this strong notion of succinctness using heavy tools such as indistinguishability obfuscation (iO) or a combination of fully homomorphic encryption and attribute-based encryption.

In this work [ILL24], we introduce new succinct garbling schemes based on variants of standard group-based assumptions. Our approach, being different from prior methods, offers a promising pathway towards practical succinct garbling. Specifically, we construct:

- *A succinct partial garbling scheme* for general circuits, where the garbled circuit size scales linearly with the private computation $|C_{\text{priv}}|$ and is independent of the public computation $|C_{\text{pub}}|$. This implies fully succinct *conditional disclosure of secrets* (CDS) protocols for circuits.
- *Succinct (fully hiding) garbling schemes* for simple types of programs, including truth tables, bounded-length branching programs (capturing decision trees and DFAs as special cases) and degree-2 polynomials, where the garbled program size is independent of the program size. This implies succinct *private simultaneous messages* (PSM) protocols for the same programs.

Our succinct partial garbling scheme can be based on a circular-security variant of the power-DDH assumption, which holds in the generic group model, or alternatively on the key-dependent message security of the Damgård-Jurik encryption. For bounded-depth circuits or the aforementioned simple programs, we avoid circular-security assumptions entirely.

At the heart of our technical approach is a new computational flavor of *algebraic homomorphic MAC* (aHMAC), for which we obtain group-based constructions building on techniques from the literature on homomorphic secret sharing. Beyond succinct garbling, we demonstrate the utility of aHMAC by constructing *constrained pseudorandom functions* (CPRFs) for general constraint circuits

from group-based assumptions. Previous CPRF constructions were limited to NC^1 circuits or alternatively relied on lattices or iO.

Part III: A major challenge in cryptography is the construction of *succinct garbling schemes* that have asymptotically smaller size than Yao's garbled circuit construction. We present a new framework for succinct garbling that replaces the heavy machinery of most previous constructions by lighter-weight *homomorphic secret sharing* techniques [ILL25].

Concretely, we achieve *1-bit-per-gate* (amortized) garbling size for Boolean circuits under circular variants of standard assumptions in composite-order or prime-order groups, as well as a lattice-based instantiation. We further extend these ideas to *layered* circuits, improving the per-gate cost below 1 bit, and to *arithmetic* circuits, eliminating the typical $\Omega(\lambda)$-factor overhead for garbling mod-p computations. Our constructions also feature "leveled" variants that remove circular-security requirements at the cost of adding a depth-dependent term to the garbling size.

Our framework significantly extends a recent technique of Liu, Wang, Yang, and Yu (Eurocrypt 2025) for lattice-based succinct garbling, and opens new avenues toward practical succinct garbling. For moderately large circuits with a few million gates, our garbled circuits can be *two orders of magnitude smaller* than Yao-style garbling. While our garbling and evaluation algorithms are much slower, they are still *practically feasible*, unlike previous fully succinct garbling schemes that rely on expensive tools such as iO or a non-black-box combination of FHE and ABE. This trade-off can make our framework appealing when a garbled circuit is used as a functional ciphertext that is broadcast or stored in multiple locations (e.g., on a blockchain), in which case communication and storage may dominate computational cost.

Shared Randomness: Friend or Foe?

Moni Naor
(joint work with Shahar Cohen and Adar Hadad)

What is the power of a shared random string known to several parties in a system who wish to perform some joint task? There are many settings where this is known to be helpful, especially for coordination between the parties. But what happens if the inputs to the system are chosen in a manner that is not independent of the string?

We will consider two settings, communication complexity [CN22] and distributed computing (LCL - locally checkable labelings) and discuss what happens there and the relationship to cryptographic and complexity assumptions, such as the existence of collision resistant hash functions and hardness on the average of TFNP (Total Function Nondeterministic Polynomial).

Somewhat Homomorphic Encryption from Linear Homomorphism and Sparse LPN

Alexandra Henzinger
(joint work with Henry Corrigan-Gibbs, Yael Tauman Kalai, and Vinod Vaikuntanathan)

We construct somewhat-homomorphic encryption schemes from the sparse learning-parities-with-noise (sparse LPN) problem, along with an assumption that implies linearly homomorphic encryption (e.g., the decisional Diffie-Hellman or decisional composite residuosity problems) [CGHKV24]. Our resulting schemes support an a-priori bounded number of homomorphic operations: $o(\log \lambda)$ multiplications followed by $poly(\lambda)$ additions, where λ is a security parameter. This gives the first somewhat-homomorphic encryption schemes that can evaluate the class of bounded-degree polynomials without relying on lattice assumptions or bilinear maps.

Tiny SNARKs in the Generic Group Model

Jesko Dujmovic
(joint work with Gal Arnon and Yuval Ishai)

In this talk we present new designated-verifier SNARKs. Succinct Non-interactive Arguments of Knowledge are very powerful tools in cryptography. They allow a prover to convince a verifier of a statement while communicating very little. Designated-verifier SNARKs are a special case where the proof can only be verified by a specific party. We present the smallest known SNARKs that do not require indistinguishability Obfuscation. Indeed, we only need cyclic groups and our proofs are one group element and $O(1)$ extra bits. We apply techniques from homomorphic secret sharing to compress proofs drastically.

Verifying Properties of Distributions: Constructions

Tal Herman

Suppose we have access to a small number of samples from an unknown distribution, and would like to learn facts about the distribution. An untrusted data server claims to have studied the distribution and makes assertions about its properties. Can the untrusted data server prove that its assertions are approximately correct? Can a short efficiently verifiable proof be generated in polynomial time?

In the talk we present a construction of a doubly-efficient proof system for any label invariant distribution property (any property that is invariant to re-labeling of the elements in the domain).

Obfuscation from Local Mixing

Ran Canetti

(joint work with Claudio Chamon, Eduardo Mucciolo, and Andrei Ruckenstein)

We explore the possibility of obtaining general-purpose obfuscation for all circuits by way of making only simple, local, functionality preserving random perturbations in the circuit structure [CCMR25]. Towards this goal, we use the additional structure provided by reversible circuits, but no additional algebraic structure.

We start by formulating a new (and relatively weak) obfuscation task regarding the ability to obfuscate *random* circuits of *bounded length*. We call such obfuscators *random input & output* (RIO) obfuscators. We then show how to construct indistinguishability obfuscators for all (unbounded length) circuits given only an RIO obfuscator – under a new assumption regarding the pseudorandomness of sufficiently long random reversible circuits with known functionality, which in turn builds on a conjecture made by Gowers (Comb. Prob. Comp. '96) regarding the pseudorandomness of bounded-size random reversible circuits. Furthermore, the constructed obfuscators satisfy a new measure of security – called *random output indistinguishability* (ROI) obfuscation – which is significantly stronger than IO and may be of independent interest.

We then investigate the possibility of constructing RIO obfuscators using local, functionality preserving perturbations. Our approach is rooted in statistical mechanics and can be thought of as locally "thermalizing" a circuit while preserving its functionality. We provide candidate constructions along with a pathway for analyzing the security of such strategies.

Given the power of program obfuscation, viability of the proposed approach would provide an alternative route to realizing almost all cryptographic tasks using the computational hardness of problems that are very different from standard ones. Furthermore, our specific candidate obfuscators are relatively efficient: the obfuscated version of an n-wire, m-gate (reversible) circuit with security parameter κ has n wires and $\mathrm{poly}(n, \kappa)m$ gates. We hope that our initial exploration will motivate further study of this alternative path to cryptography.

Exotic Lattice Assumptions and How to Tame Them

David Wu

A recent and exciting line of work has introduced new variants the classic short integer solutions (SIS) and learning with errors (LWE) problems in lattice-based cryptography. This talk surveys some of the recent developments in this area of research. I will start by describing the evasive LWE assumption and highlight some of its applications and cryptanalysis. In the second half of this talk, I will introduce the succinct LWE assumption, a new falsifiable lattice assumption introduced in the work of Wee. This assumption asserts that the LWE problem is hard even given a trapdoor for a related matrix. Finally, I will describe two types of applications

of succinct LWE. The first gives notions like functional commitments and succinct attribute-based encryption for circuits. The second gives trustless cryptographic schemes such as distributed broadcast encryption and registered attribute-based encryption. I conclude by describing some open problems.

Sparse LPN, LWE and Friends

Aayush Jain

In this talk, we survey recent and earlier progress on several related noisy linear algebraic assumptions with sparse coefficient vectors, including sparse Learning Parity with Noise (Alekhnovich, 2003) and sparse Learning with Errors (Jain, Lin, Saha, 2024). We will discuss the algorithms and complexity of these problems, their exciting cryptographic applications, and connections to statistical inference problems. The talk will also highlight important open problems in this area.

Recent advances in isogeny-based cryptography

Benjamin Wesolowski

A non-zero morphism between two elliptic curves is called an isogeny. Isogeny-based cryptography is based on the presumed hardness of the following problem: given two supersingular elliptic curves, find an isogeny between them. This talk reviews this problem, its properties, its variants (like the ℓ-Isogeny Path problem), and its "endomorphism" cousins (given a curve, find isogenies from the curve to itself: the One Endomorphism or the Endomorphism Ring problems). Remarkably, all these problems are equivalent, with some reductions contingent on the Generalized Riemann Hypothesis. I then present how the connections between these problems enable the design of a digital signature scheme: SQIsign.

Batching NISZK Proofs

Prashant Vasudevan
(joint work with Changrui Mu, Shafik Nassar, and Ron D. Rothblum)

In a zero-knowledge proof, a prover needs to convince a verifier that an input x is contained in a language Π without revealing any additional information. By repeating a zero-knowledge proof k times, it is possible to prove (still in zero-knowledge) that k separate inputs $x_1, \ldots, x_k$ all belong to Π. But this increases the communication by a factor of k. In this talk, I will show how to do better. In particular, we will see that any problem in NISZK has a non-interactive statistical zero-knowledge batch verification protocol with communication $\mathsf{poly}(n, \log(k))$ [MNRV24].

How to Prove False Statements: Practical Attacks on Fiat-Shamir

Ron Rothblum
(joint work with Dmitry Khovratovich, Lev Soukhanov)

We show an attack against the Fiat-Shamir security of a standard and popular interactive succinct argument, based on the GKR protocol [GKR08], for verifying the correctness of a non-determinstic bounded-depth computation.

For every choice of FS hash function, we show that a corresponding instantiation of this protocol, which was been widely studied in the literature and used also in practice, is not (adaptively) sound when compiled with the FS transform [KRS25].

Discussion group: Which direction(s) should evasive LWE research take?

Chris Brzuska

The evasive LWE assumption, proposed by Wee (Eurocrypt'22) for constructing lattice-based optimal broadcast encryption, has shown to be a powerful assumption, adopted by subsequent works to construct advanced primitives ranging from ABE variants to obfuscation for null circuits. However, a closer look reveals significant differences among the precise assumption statements.

Our current understanding is as follows: Vaikuntanathan, Wee and Wichs (Asiacrypt'22) give a heuristic obfuscation-based counterexample against all private-coin versions of evasive LWE. Assuming witness encryption and LWE, Brzuska, Ünal and Woo (Asiacrypt'24) give a provable counterexample against all those private-coin versions of evasive LWE where the sampler obtains the secret; and Branco, Döttling, Jain, Malavolta, Mathialagan, Peters, and Vaikuntanathan (ePrint '24) give a counterexample against binding evasive LWE via pseudorandom obfuscation. Finally, Brzuska, Ünal and Woo (Asiacrypt'24) also give simple examples against private-coin evasive LWE when P is partially known or B is known but P is not. There are no known counterexamples against public-coin evasive LWE (except when the sampler takes B as input).

Discussion group: Based on these counterexamples, how can we find meaningful subclasses of evasive LWE which are useful for strong applications and do not suffer from provable counterexamples? Should we move away from private-coin evasive LWE to public-coin evasive LWE and succinct LWE? Which other research directions would be useful to pursue?

How to Share an NP Statement or Non-Interactive Combiners and Amplifiers for Zero-Knowledge Proofs

Benny Applebaum
(joint work with Eliran Kachlon)

In Crypto 2019, Goyal, Jain, and Sahai introduced the elegant notion of *secret-sharing of an* NP *statement* (NPSS). Roughly speaking, a t-out-of-n secret sharing of an NP statement is a reduction that maps a circuit-SAT instance f into n circuit-SAT instances $(f_1, \ldots, f_n)$ over a common set of variables, such that: (1) Given a satisfying assignment x for f, it is possible to sample partial assignments $(y_1, \ldots, y_n)$ that consistently satisfy $(f_1, ..., f_n)$, and where the marginal distribution of every collection of $t-1$ partial assignments leaks no information about the witness x; (2) Conversely, any collection of t partial assignments that consistently satisfy t of the instances can be efficiently translated into an assignment x that satisfies f.

We present the first information-theoretic construction of NPSS for arbitrary values of t and n [AK25]. Previously, it was only known how to achieve computational privacy for the special case of $t = n$. Our constructions rely on a new notion of secure multiparty computation protocols that may be of independent interest. We use our NPSS to obtain several applications in the domain of zero-knowledge proofs and secure-multiparty computation.

Microcrypt

Dakshita Khurana
(joint work with Kabir Tomer)

Part I: One-way functions are central to classical cryptography. They are necessary for the existence of non-trivial classical cryptosystems, and also sufficient to realize meaningful primitives including commitments, pseudorandom generators and digital signatures. At the same time, a mounting body of evidence suggests that assumptions *even weaker* than one-way functions may suffice for many cryptographic tasks of interest in a quantum world, including bit commitments and secure multi-party computation.

This work [KT24a] studies one-way state generators [Morimae–Yamakawa, CRYPTO 2022], a natural quantum relaxation of one-way functions. Given a secret key, a one-way state generator outputs a hard to invert quantum state. A fundamental question is whether this type of *quantum* one-wayness suffices to realize quantum cryptography. We obtain an affirmative answer to this question, by proving that one-way state generators with pure state outputs imply quantum bit commitments and secure multiparty computation.

Along the way, we use efficient shadow tomography [Huang et al., Nature Physics 2020] to build an intermediate primitive with classical outputs, which we call a

(quantum) one-way puzzle. Our main technical contribution is a proof that one-way puzzles imply quantum bit commitments. This proof develops new techniques for pseudoentropy generation [Håstad et al., SICOMP 1999] from arbitrary distributions, which may be of independent interest.

Part II: Recent oracle separations [Kretschmer, TQC'21, Kretschmer et al., STOC'23] have raised the tantalizing possibility of building quantum cryptography from sources of hardness that persist even if the polynomial hierarchy collapses. We realize this possibility by building quantum bit commitments and secure computation from *unrelativized*, well-studied mathematical problems that are conjectured to be hard for $\mathsf{P}^{\#\mathsf{P}}$ – such as approximating the permanents of complex Gaussian matrices, or approximating the output probabilities of random quantum circuits. Indeed, we show [KT24b] that as long as *any one of the conjectures* underlying sampling-based quantum advantage (e.g., BosonSampling [Aaronson–Arkhipov, STOC'11], Random Circuit Sampling [Boixo et al., Nature Physics 2018], IQP [Bremner, Jozsa and Shepherd, Proc. Royal Society of London 2010]) is true, quantum cryptography can be based on the extremely mild assumption that $\mathsf{P}^{\#\mathsf{P}} \not\subseteq (\mathsf{io})\mathsf{BQP}/\mathsf{qpoly}$.

Our techniques uncover strong connections between the hardness of approximating the probabilities of outcomes of quantum processes, the existence of "one-way" state synthesis problems, and the existence of useful cryptographic primitives such as one-way puzzles and quantum bit commitments. Specifically, we prove that the following hardness assumptions are equivalent under BQP reductions.

- **The hardness of approximating the probabilities** of outcomes of certain efficiently sampleable distributions. That is, there exist quantumly efficiently sampleable distributions for which it is hard to approximate the probability assigned to a randomly chosen string in the support of the distribution (up to inverse polynomial relative error).
- **The existence of one-way puzzles**, where a quantum sampler outputs a pair of classical strings – a puzzle and its key – and where the hardness lies in finding the key corresponding to a random puzzle. These are known to imply quantum bit commitments.
- **The existence of state puzzles**, or one-way state synthesis, where it is hard to synthesize a secret quantum state given a public classical identifier. These capture the hardness of search problems with quantum secrets and classical challenges.

These are the first constructions of quantum cryptographic primitives (one-way puzzles, quantum bit commitments, state puzzles) from well-studied mathematical assumptions that do not imply the existence of classical cryptography.

Along the way, we also show that distributions that admit efficient quantum samplers but cannot be pseudo-deterministically efficiently sampled imply quantum commitments.

When the Universe Speaks in (Quantum) Crypto

Zvika Brakerski

We propose to study equivalence relations between phenomena in high-energy physics and the existence of standard cryptographic primitives, and show the first example where such an equivalence holds [Bra23]. A small number of prior works showed that high-energy phenomena *can be explained* by cryptographic hardness. Examples include using the existence of one-way functions to explain the hardness of decoding black-hole Hawking radiation (Harlow and Hayden 2013, Aaronson 2016), and using pseudorandom quantum states to explain the hardness of computing AdS/CFT dictionary (Bouland, Fefferman and Vazirani, 2020).

In this work we show, for the former example of black-hole radiation decoding, that it also *implies* the existence of secure quantum cryptography. In fact, we show an existential equivalence between the hardness of black-hole radiation decoding and a variety of cryptographic primitives, including bit-commitment schemes and oblivious transfer protocols (using quantum communication). This can be viewed (with proper disclaimers, as we discuss) as providing a physical justification for the existence of secure cryptography. We conjecture that such connections may be found in other high-energy physics phenomena.

How to Construct Random Unitaries

Fermi Ma
(joint work with Hsin-Yuan Huang)

The ability to efficiently implement a seemingly random function (i.e., a pseudorandom function) is a cornerstone of modern cryptography. Is an analogous statement true in the quantum world? Namely, can we efficiently implement seemingly random unitaries, i.e., pseudorandom unitaries (PRUs)? This would have broad implications for quantum computing and physics.

In this talk, I will present joint work with Hsin-Yuan Huang [MH24] in which we settle this question, proving that PRUs exist assuming one-way functions. Our proof analyzes a construction of Metger, Poremba, Sinha, and Yuen, using elementary arguments based on purification.

We achieve our results on PRUs by proving that any quantum oracle algorithm A^U that queries an n-qubit Haar-random unitary U can be *efficiently simulated* with a remarkably simple procedure:

(1) Initialize an external register E to the state $|\emptyset\rangle$, where $\emptyset$ denotes the empty set. (**Aside:** When we write a set inside a ket, e.g., $|S\rangle_{\mathsf{E}}$, we are simply using the set S as a label for a unit vector. The inner product $\langle R|S\rangle$ equals 1 if $R = S$ and 0 otherwise.)

(2) Run the oracle algorithm A, replacing each query to U with the following linear map:

$$V : |x\rangle|S\rangle_{\mathsf{E}} \mapsto \frac{1}{\sqrt{2^n - |S|}} \sum_{\substack{y\in\{0,1\}^n:\\ y\notin S_Y}} |y\rangle|S \cup \{(x,y)\}\rangle_{\mathsf{E}}, \tag{1}$$

where S_Y denotes the set of all y such that $(x,y) \in S$ for some x. In words, V maps x to a uniform superposition over $y \in \{0,1\}^n$, except those that already appear in S, and simultaneously "records" (x,y) by inserting it into S. We refer to V as the *path-recording oracle*.

We prove that the following mixed states have trace distance $O(t^2/2^n)$:

- $\mathbb{E}_U |\mathsf{A}^U\rangle\langle\mathsf{A}^U|$, the state of A after t queries to a Haar-random unitary U, where $|\mathsf{A}^U\rangle := U \cdot A_t \cdots U \cdot A_1|0\rangle$ denotes the state of A after t queries to U, and $|0\rangle$ denotes an arbitrary initial state.
- $\mathrm{Tr}_{\mathsf{E}}(|\mathsf{A}^V\rangle\langle\mathsf{A}^V|)$, where $|\mathsf{A}^V\rangle_{\mathsf{AE}} := V \cdot A_t \cdots V \cdot A_1|0\rangle|\emptyset\rangle_{\mathsf{E}}$ denotes the global state of the algorithm and the external register E after t queries to V.

Despite the extensive literature on Haar-random unitaries, to the best of our knowledge, this "path-recording" characterization was not known before.[1] As we now explain, this new path-recording perspective is the key to our PRU proof.

How to construct PRUs. The main technical step in our PRU proof is to show that a t-query oracle algorithm A can only distinguish between

- $P_\pi \cdot F_f \cdot C$, where $P_\pi = \sum_x |\pi(x)\rangle\langle x|$ for a random permutation $\pi \leftarrow S_{2^n}$, $F_f = \sum_x (-1)^{f(x)}|x\rangle\langle x|$ for a random function $f \leftarrow \{0,1\}^{2^n}$, and C is a random n-qubit Clifford.[2]
- a Haar-random n-qubit unitary U,

with probability $1/2 + t^2/2^n$.

Our proof works by *purifying* the randomness of the PRU. Ignoring C for now, suppose we initialize an external register to the uniform superposition $\propto \sum_{\pi\in S_{2^n}} |\pi\rangle \otimes \sum_{f\in\{0,1\}^{2^n}} |f\rangle$ over all permutations π and functions f. In this view, a query to a random $P_\pi \cdot F_f$ is equivalent to a query to a fixed unitary that applies $P_\pi \cdot F_f$ controlled on $|\pi\rangle|f\rangle$, i.e., the map

$$|x\rangle \otimes |\pi, f\rangle \mapsto (-1)^{f(x)} \cdot |\pi(x)\rangle \otimes |\pi, f\rangle. \tag{2}$$

Equivalently, we can view this map as sending x to a superposition over all y, while simultaneously multiplying the purifying register by the coefficient $\delta_{\pi(x)=y} \cdot (-1)^{f(x)}$:

$$|x\rangle \otimes |\pi, f\rangle \mapsto \sum_{y\in\{0,1\}^n} |y\rangle \otimes \Big(\delta_{\pi(x)=y} \cdot (-1)^{f(x)} \cdot |\pi, f\rangle\Big). \tag{3}$$

[1] This can also be viewed as an analog of Zhandry's compressed oracles for Haar-random unitaries [Zha19].

[2] This *PFC* construction was introduced by [MPSY24], who proved security against *non-adaptive* adversaries, i.e., adversaries that make all of their oracle queries at once, in parallel.

After t queries to the purified $P_\pi \cdot F_f$, the global state including the purifying registers is (proportional to) a sum of terms

(4)

$$|y_t\rangle\langle x_t|\cdot A_t \cdots |y_1\rangle\langle x_1|\cdot A_1|0^n\rangle \otimes \underbrace{\sum_{\pi\in S_{2^n}} |\pi, f\rangle \cdot \delta_{\pi(x_1)=y_1}\cdots\delta_{\pi(x_t)=y_t}\cdot(-1)^{f(x_1)+\cdots+f(x_t)}}_{\propto |\mathsf{pf}_{\{(x_1,y_1),\ldots,(x_t,y_t)\}}\rangle},$$

over all possible $x_1, y_1, \ldots, x_t, y_t \in \{0,1\}^n$, i.e., over all *Feynman paths*.

Crucially, when all the $x_1, \ldots, x_t$ are distinct, these $|\mathsf{pf}_{\{(x_1,y_1),\ldots,(x_t,y_t)\}}\rangle$ states are orthogonal and is isometric to $|\{(x_1,y_1),\ldots,(x_t,y_t)\}\rangle$. Since the algorithm is not given the purifying registers, a query to a random $P_\pi \cdot F_f$ is *identical* to a query to the path-recording oracle V described earlier—except on paths where there is a collision among the inputs $x_1, \ldots, x_t$.

This is where C comes in. We prove that V satisfies a key property: for any n-qubit unitary C,

$$(5)\quad (V\cdot C)\cdot A_t \cdots (V\cdot C)\cdot A_1|0^n\rangle|\emptyset\rangle_{\mathsf{E}} = ((C\otimes \mathsf{Id})^{\otimes t})_{\mathsf{E}}\cdot V\cdot A_t\cdots V\cdot A_1|0^n\rangle|\emptyset\rangle_{\mathsf{E}}.$$

This says that applying C to the **adversary's register** before each query to V is equivalent to applying C to each x_i in the **purifying register** $|\{(x_1,y_1),\ldots,(x_t,y_t)\}\rangle$. When C is sampled from any 2-design, the randomness of C ensures there are no collisions in the $x_1,\ldots,x_t$ with overwhelming probability. Consequently, we show that queries to V are indistinguishable from queries to $P_\pi\cdot F_f\cdot C$, as long as C is sampled from *any* 2-design. By instantiating the 2-design to be either (1) a random Clifford or (2) a Haar-random unitary, we show that both $P_\pi\cdot F_f\cdot C$ and Haar-random unitaries are indistinguishable from V, and thus, from each other.

Strong PRUs and a symmetrized path-recording oracle $\widetilde{V}$. To obtain strong PRUs, we use the construction: $D\cdot P_\pi\cdot F_f\cdot C$, where D, C are both random n-qubit Cliffords, P_π is the same as before, and F_f is a random q-ary phase (for any $q\geq 3$). By analyzing the purification of $P_\pi\cdot F_f$, we show that when A makes forward and inverse queries, the purifying registers, viewed in the right basis, "record" information from *two Feynman paths*: one set S^{for} consists of (x,y) tuples corresponding to the forward queries, and another set S^{inv} of tuples (x,y) corresponds to the inverse queries. Whereas each query in the standard PRU proof always inserts a tuple (x,y) into the set S, when both forward and inverse queries are allowed, the effect is more intricate:

- A forward query will sometimes add a tuple to S^{for}, but other times delete a tuple from S^{inv}.
- An inverse query will sometimes add a tuple to S^{inv}, but other times delete a tuple from S^{for}.

We prove that this behavior corresponds to a more general "symmetrized" path recording oracle $\widetilde{V}$. Moreover, as long as D, C are sampled from any 2-design, the adversary cannot distinguish between queries to $D\cdot P_\pi\cdot F_f\cdot C$ and queries to $\widetilde{V}$, and using similar reasoning as the standard PRU proof, conclude both of the

following (1) strong PRUs exist and (2) $\widetilde{V}$ is indistinguishable from Haar-random even under inverse queries.

Our proof requires handling several additional technical challenges and leverages the following property of 2-designs: if one samples C from a 2-design and applies $C \otimes \overline{C}$ to any state (where $\overline{C}$ denotes the complex conjugate), then with overwhelmingly high probability, the result is either (a) a pair of distinct elements, or (b) the maximally entangled state. The fact that there are two kinds of outcomes after twirling by $C \otimes \overline{C}$ is intimately related to the mechanism by which the purification "decides" whether it should add or delete a tuple (x, y).

A new way to analyze random unitaries. More broadly, the path-recording oracle unlocks a new way to proving theorems about random unitaries. Before this work, analyzing mixed states such as $\mathbb{E}_U |\mathrm{Adv}^U\rangle\langle \mathrm{Adv}^U|$ often necessitated the use of Weingarten calculus, involving intricate asymptotic bounds on Weingarten functions through sophisticated combinatorial and representation-theoretic calculations. Our approach circumvents this complexity entirely.[3]

To demonstrate the power of this approach, we give an elementary proof of the "gluing lemma" recently proven by [SHH24]. This lemma states that if two Haar-random unitaries U_1 and U_2 *overlap*, with U_1 acting on systems A, B and U_2 on B, C (where B has a super-logarithmic number of qubits), then queries to $U_2 \cdot U_1$ are indistinguishable from queries to a larger Haar-random unitary U acting on $\mathsf{A}, \mathsf{B}, \mathsf{C}$. Using this lemma (and our main theorem), [SHH24] constructed low-depth PRUs secure against forward queries. However, their proof of the gluing lemma is highly technical, relying on careful representation-theoretic analysis and tight bounds on Weingarten functions.

We demonstrate that the path-recording oracle yields an elementary proof of the gluing lemma. The key insight is to replace the Haar-random unitaries with path-recording oracles. This reduces to showing that the composition of two independent path-recording oracles $V_2 \cdot V_1$, where V_1 acts on $(\mathsf{A}, \mathsf{B}, \mathsf{E}_1)$ and V_2 acts on $(\mathsf{B}, \mathsf{C}, \mathsf{E}_2)$, approximates a single path-recording oracle V acting on $(\mathsf{A}, \mathsf{B}, \mathsf{C}, \mathsf{E})$.

Given the ubiquity of random unitaries in physics and quantum computing, we anticipate many future applications of the path-recording framework.

Error Detection and Correction in a Computationally Bounded World

Daniel Wichs
(joint work with Jad Silbak)

We study error detection/correction against PPT channels. We allow codes to rely on a public random seed, either only given to the encoder or to both encoder and decoder. In this model we construct codes that beat information theoretic counterparts, both in the large-alphabet [SW25b] and binary-alphabet [SW25a] settings.

[3] Alternatively, one can view our technique as deriving a simplified and approximate version of the Weingarten calculus from purely elementary arguments.

Simplified private information retrieval protocols

Amos Beimel
(joint work with Bar Alon and Or Lasri)

In a k-server private information retrieval (PIR) protocol, there are k servers holding an n-bit database and a user that wants to get one bit of the database while hiding this index from each server. The best constructions of PIR protocols use matching vectors and specific share conversions. The goal of this talk is to present simplified and generalized versions of the best known PIR protocols, the 2-server PIR protocol of Dvir and Gopi (J. ACM 2016) and the 3-server and multi-server protocol of Ghasemi, Kopparty, and Sudan (IACR Cryptol. ePrint 2024). The simplification is done by considering a new variant of matching vectors and by using a general share conversion [ABL24]. In addition to simplifying previous protocols, our 2-server protocol can use matching vectors over any m that is a product of two distinct primes. Our construction does not improve the communication complexity of PIR and CDS protocols; however, construction of better matching vectors over *any* m that is a product of two distinct primes will improve their communication complexity.

Lattices, Statistics, and Removing Backdoors From ML Models

Neekon Vafa
(joint work with Shafi Goldwasser, Jonathan Shafer, and Vinod Vaikuntanathan)

In this talk, I will discuss recent connections between cryptographic ideas, statistics, and security in machine learning.

The first part of the talk examines the average-case computational complexity of the number partitioning problem (NPP), which is a natural average-case discrepancy problem in dimension one. Specifically, for an input consisting of continuous uniform i.i.d. $a_1, \ldots, a_n \leftarrow [0,1]$, the goal of the problem is to output some vector $x \in \{-1,1\}^n$ so that $\left|\sum_{i\in[n]} a_i x_i\right|$ is below some threshold $\kappa(n)$. The statistical threshold for solutions to exist is $\kappa(n) \approx 2^{-n}$; below this threshold, solutions likely do not exist, and above this threshold, (many) solutions likely exist.

The beautiful differencing algorithm of Karmarkar and Karp (1982) gives a polynomial time algorithm for this problem that succeeds as long as $\kappa(n) \geq 2^{-\Theta(\log^2 n)}$. Other works have given some evidence of the hardness of achieving κ close to the statistical threshold, but none of them come close to the inverse quasi-polynomial bound of Karmarkar and Karp's algorithm. Specifically, Hoberg, Ramadas, Rothvoss, and Yang (2017) give evidence that a worst-case version of NPP is hard if certain lattice problems are hard, but only up to $\kappa(n) \leq 2^{-n/2}$. Gamarnik and Kizildag (2021) show that the (multi) overlap gap property sets in at $2^{-omega(\sqrt{n}\log n)}$, implying the failure of stable algorithms below this value of κ.

I will describe a reduction from approximate worst-case lattice problems to NPP [VV25]. This reduction shows that any (average-case) polynomial-time algorithm that achieves $\kappa(n) \leq 2^{-\log^{3+\epsilon} n}$ would beat state-of-the-art lattice algorithms. This brings the hardness threshold much closer to Karmarkar and Karp's differencing algorithm, and gives the first complexity-theoretic evidence of hardness from worst-case computational problems.

In the second part of the talk, I will discuss how to remove backdoors from machine learning (ML) models [GSVV24]. As society grows more reliant on ML, ensuring the security of ML systems against sophisticated attacks becomes a pressing concern. A recent result of Goldwasser, Kim, Vaikuntanathan, and Zamir (2022) shows that an adversary can plant undetectable backdoors in ML models under standard cryptographic assumptions, allowing the adversary to covertly control the model's behavior. Backdoors can be planted in such a way that the backdoored ML model is computationally indistinguishable from an honest model without backdoors. I will mention strategies for defending against backdoors in ML models, even if they are undetectable. The key observation is that it is sometimes possible to provably mitigate or even remove backdoors without needing to detect them, using techniques inspired by the notion of random self-reducibility. This depends on properties of the ground-truth labels (chosen by nature), and not of the proposed ML model (which may be chosen by an attacker).

First, we show an "offline mitigation" technique, which removes all backdoors from a ML model under the assumption that the ground-truth labels are close to a Fourier-heavy function. Second, we consider distributions where the ground-truth labels are close to a linear or polynomial function in $\mathbb{R}^n$. Here, we show "online mitigation" techniques, which remove backdoors with high probability for every input of interest, and are computationally cheaper than offline mitigation. All of our constructions are black-box, so our techniques work without needing access to the model's representation (i.e., its code or parameters).

The (Module) Lattice Isomorphism Problem

Alice Pellet-Mary

The lattice isomorphism problem (LIP) has been introduced in cryptography in 2022, as a hard problem to build public key encryption and signatures. Notably, the signature scheme Hawk [DPPvW22], which was submitted to the NIST standardization process (and is currently selected for the second round) is based on a variant of the lattice isomorphism problem: the module lattice isomorphism problem (module-LIP). This variant replaces plain lattices by structured lattices (named module lattices), which allows the signature scheme Hawk to be somewhat compact and efficient. On the other hand, introducing more structure raises the question of the security of the scheme. In this talk, we will review how the signature scheme Hawk is constructed, as well as recent results on the hardness of module-LIP.

Cloning Games, Black Holes and Cryptography

Alexander Poremba
(joint work with Seyoon Ragavan and Vinod Vaikuntanathan)

The no-cloning principle has played a foundational role in quantum information and cryptography. Following a long-standing tradition of studying quantum mechanical phenomena through the lens of interactive games, Broadbent and Lord (TQC 2020) formalized cloning games in order to quantitatively capture no-cloning in the context of unclonable encryption schemes. The conceptual contribution of this paper [PRV24] is the new, natural, notion of Haar cloning games together with two applications. In the area of black-hole physics, our game reveals that, in an idealized model of a black hole which features Haar random (or pseudorandom) scrambling dynamics, the information from infalling entangled qubits can only be recovered from either the interior or the exterior of the black hole – but never from both places at the same time. In the area of quantum cryptography, our game helps us construct succinct unclonable encryption schemes from the existence of pseudorandom unitaries, thereby, for the first time, bridging the gap between "MicroCrypt" and unclonable cryptography. The technical contribution of this work is a tight analysis of Haar cloning games which requires us to overcome many long-standing barriers in our understanding of cloning games. Answering these questions provably requires us to go beyond existing methods (Tomamichel, Fehr, Kaniewski and Wehner, New Journal of Physics 2013). In particular, we show a new technique for analyzing cloning games with respect to binary phase states through the lens of binary subtypes, and combine it with novel bounds on the operator norms of block-wise tensor products of matrices.

References

[ABL24] Bar Alon, Amos Beimel, and Or Lasri. Simplified pir and cds protocols and improved linear secret-sharing schemes. *Cryptology ePrint Archive*, 2024.

[AGR25] Noga Amir, Oded Goldreich, and Guy N Rothblum. Doubly sub-linear interactive proofs of proximity. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 6–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.

[AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 308–326. Springer, 2015.

[AK25] Benny Applebaum and Eliran Kachlon. How to share an np statement or combiners for zero-knowledge proofs. *Cryptology ePrint Archive*, 2025.

[BB84] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, 1984.

[BEG25] Nir Bitansky, Saroja Erabelli, and Rachit Garg. Additive randomized encodings from public key encryption. *Cryptology ePrint Archive*, 2025.

[BGI+12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.

[BLLL23] Marshall Ball, Hanjun Li, Huijia Lin, and Tianren Liu. New ways to garble arithmetic circuits. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part II*, volume 14005 of *Lecture Notes in Computer Science*, pages 3–34, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland.

[Bra23] Zvika Brakerski. Black-hole radiation decoding is quantum cryptography. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 37–65, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[BV18] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. *J. ACM*, 65(6):39:1–39:37, 2018.

[CCMR25] Ran Canetti, Claudio Chamon, Eduardo R Mucciolo, and Andrei E Ruckenstein. Towards general-purpose program obfuscation via local mixing. In *Theory of Cryptography Conference*, pages 37–70. Springer, 2025.

[CGHKV24] Henry Corrigan-Gibbs, Alexandra Henzinger, Yael Kalai, and Vinod Vaikuntanathan. Somewhat homomorphic encryption from linear homomorphism and sparse lpn. *Cryptology ePrint Archive*, 2024.

[CGZ23] Miranda Christ, Sam Gunn, and Or Zamir. Undetectable watermarks for language models. *CoRR*, abs/2306.09194, 2023.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584. Springer, 2021.

[CN22] Shahar P. Cohen and Moni Naor. Low communication complexity protocols, collision resistant hash functions and secret key-agreement protocols. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 252–281, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[DPPvW22] Léo Ducas, Eamonn W. Postlethwaite, Ludo N. Pulles, and Wessel P. J. van Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 65–94, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland.

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49. IEEE Computer Society, 2013.

[GGH+16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Hiding secrets in software: a cryptographic approach to program obfuscation. *Commun. ACM*, 59(5):113–120, 2016.

[GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 113–122, Victoria, BC, Canada, May 17–20, 2008. ACM Press.

[GKVZ22] Shafi Goldwasser, Michael P. Kim, Vinod Vaikuntanathan, and Or Zamir. Planting undetectable backdoors in machine learning models : [extended abstract]. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 931–942. IEEE, 2022.

[GRSY21] Shafi Goldwasser, Guy N. Rothblum, Jonathan Shafer, and Amir Yehudayoff. Interactive proofs for verifying machine learning. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPIcs*, pages 41:1–41:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[GSVV24] Shafi Goldwasser, Jonathan Shafer, Neekon Vafa, and Vinod Vaikuntanathan. Oblivious defense in ml models: Backdoor removal without detection. *arXiv preprint arXiv:2411.03279*, 2024.

[HIKR23] Shai Halevi, Yuval Ishai, Eyal Kushilevitz, and Tal Rabin. Additive randomized encodings and their applications. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 203–235, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[HR23] Tal Herman and Guy N. Rothblum. Doubley-efficient interactive proofs for distribution properties. In *64th Annual Symposium on Foundations of Computer Science*, pages 743–751, Santa Cruz, CA, USA, November 6–9, 2023. IEEE Computer Society Press.

[HR24] Tal Herman and Guy N. Rothblum. Interactive proofs for general distribution properties. In *65th Annual Symposium on Foundations of Computer Science*, pages 528–538, Chicago, IL, USA, October 27–30, 2024. IEEE Computer Society Press.

[ILL24] Yuval Ishai, Hanjun Li, and Huijia Lin. Succinct partial garbling from groups and applications. *Cryptology ePrint Archive*, 2024.

[ILL25] Yuval Ishai, Hanjun Li, and Huijia Lin. A unified framework for succinct garbling from homomorphic secret sharing. *Cryptology ePrint Archive*, 2025.

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 60–73. ACM, 2021.

[JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb {F}_p$, dlin, and prgs in nc^0. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part I*, volume 13275 of *Lecture Notes in Computer Science*, pages 670–699. Springer, 2022.

[KRS25] Dmitry Khovratovich, Ron D Rothblum, and Lev Soukhanov. How to prove false statements: Practical attacks on fiat-shamir. *Cryptology ePrint Archive*, 2025.

[KT24a] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th Annual ACM Symposium on Theory of Computing*, pages 968–978, Vancouver, BC, Canada, June 24–28, 2024. ACM Press.

[KT24b] Dakshita Khurana and Kabir Tomer. Founding quantum cryptography on quantum advantage, or, towards cryptography from #p-hardness. *arXiv preprint arXiv:2409.15248*, 2024.

[Mah22] Urmila Mahadev. Classical verification of quantum computations. *SIAM J. Comput.*, 51(4):1172–1229, 2022.

[MH24] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. *arXiv preprint arXiv:2410.10116*, 2024.

[MNRV24] Changrui Mu, Shafik Nassar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Strong batching for non-interactive statistical zero-knowledge. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 241–270, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.

[MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries. In *65th Annual Symposium on Foundations of Computer Science*, pages 485–492, Chicago, IL, USA, October 27–30, 2024. IEEE Computer Society Press.

[PRV24] Alexander Poremba, Seyoon Ragavan, and Vinod Vaikuntanathan. Cloning games, black holes and cryptography. *arXiv preprint arXiv:2411.04730*, 2024.

[SHH24] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. Random unitaries in extremely low depth. *arXiv preprint arXiv:2407.07754*, 2024.

[Sho94] Peter W. Shor. Polynominal time algorithms for discrete logarithms and factoring on a quantum computer. In Leonard M. Adleman and Ming-Deh A. Huang, editors, *Algorithmic Number Theory, First International Symposium, ANTS-I, Ithaca, NY, USA, May 6-9, 1994, Proceedings*, volume 877 of *Lecture Notes in Computer Science*, page 289. Springer, 1994.

[SW25a] Jad Silbak and Daniel Wichs. Binary codes for error detection and correction in a computationally bounded world. *Cryptology ePrint Archive*, 2025.

[SW25b] Jad Silbak and Daniel Wichs. Detecting and correcting computationally bounded errors: A simple construction under minimal assumptions. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 88–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.

[VV25] Neekon Vafa and Vinod Vaikuntanathan. Symmetric perceptrons, number partitioning and lattices. *arXiv preprint arXiv:2501.16517*, 2025.

[Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.

Reporter: Alex Lombardi

Participants

Prof. Dr. Benny Applebaum
Department of Electrical
Engineering Systems
Tel Aviv University
Ramat Aviv, Tel Aviv 69978
ISRAEL

Amos Beimel
Department of Computer Science
Ben Gurion University of the Negev
P.O. Box 653
Beer-Sheva 84105
ISRAEL

Nir Bitansky
New York University
and Tel Aviv University
New York NY 10012
UNITED STATES

Prof. Dr. Zvika Brakerski
Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O. Box 26
76100 Rehovot
ISRAEL

Prof. Dr. Chris Brzuska
Aalto University
P.O. Box 13000
00076 Espoo
FINLAND

Prof. Dr. Johannes Buchmann
Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Prof. Dr. Ran Canetti
Department of Computer Science
Boston University
Boston, MA 02215
UNITED STATES

Lalita Devadas
MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Dr. Nico Döttling
CISPA Helmholtz Center for
Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
GERMANY

Jesko Dujmovic
CISPA Helmholtz Center for
Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
GERMANY

Prof. Dr. Shafi Goldwasser
UC Berkeley
Melvin Calvin Lab
Berkeley, CA 94720
UNITED STATES

Rohan Goyal
Massachusetts Institute of Technology
70 Pacific St
Cambridge MA 02139
UNITED STATES

Alexandra Henzinger
MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Tal Herman
Dept. of Maths and Computer Science
Weizmann Institute of Sciences
Rehovot 76100
ISRAEL

Dr. Yuval Ishai
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Dr. Aayush Jain
Computer Science Department
Carnegie Mellon University
7215 Gates-Hillman Center
5000 Forbes Avenue
Pittsburgh, PA 15213
UNITED STATES

Prof. Dr. Abhishek Jain
NTT Research
Palo Alto 94301
UNITED STATES

Prof. Dr. Antoine Joux
CISPA Helmholtz Center for
Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
GERMANY

Dr. Yael Tauman Kalai
Microsoft Research Laboratory
One Memorial Drive
Cambridge MA 02142
UNITED STATES

Dr. Dakshita Khurana
Department of Computer Science
University of Illinois at
Urbana-Champaign
1304 W. Springfield Av.
Urbana, IL 61801-2987
UNITED STATES

Prof. Dr. Eike Kiltz
Lehrstuhl für Kryptographie
Ruhr-Universität Bochum
Universitätsstraße 150
44780 Bochum
GERMANY

Lisa Kohl
Centrum Wiskunde & Informatica
(CWI)
1090 GB Amsterdam
NETHERLANDS

Prof. Dr. Eyal Kushilevitz
Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Prof. Dr. Huijia (Rachel) Lin
Paul G. Allen School of Computer
Science & Engineering
University of Washington
Box 352350
Seattle WA 98195-2350
UNITED STATES

Alex Lombardi
Department of Computer Science
Princeton University
35 Olden Street
Princeton 08544-5233
UNITED STATES

Dr. Julian Loss
CISPA Helmholtz Center for
Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
GERMANY

Dr. Fermi Ma
Simons Institute and UC Berkeley
Berkeley 94720
UNITED STATES

Prof. Dr. Giulio Malavolta
Bocconi University
Via Röntgen 1
20136 Milano
ITALY

Prof. Dr. Tal Malkin
Department of Computer Science
Columbia University
560 Riverside Drive
New York, NY 10027
UNITED STATES

Surya Mathialagan
Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge MA 02139
UNITED STATES

Prof. Dr. Daniele Micciancio
Department of Computer Science &
Engineering
University of California, San Diego
9500 Gilman Drive, Mail Code 0404
La Jolla CA 92093-0404
UNITED STATES

Prof. Dr. Moni Naor
Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O. Box 26
Rehovot 76100
ISRAEL

Dr. Omer Paneth
Tel Aviv University
53631 Tel Aviv
ISRAEL

Dr. Alice Pellet-Mary
Institut de mathématiques de Bordeaux
(IMB)
351, cours de la Liberation
33405 Talence Cedex
FRANCE

Dr. Alexander Poremba
MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Willy Quach
CISPA Helmholtz Center for
Information Security
Stuhlsatzenhaus 5
66123 Saarbrücken
GERMANY

Prof. Dr. Alon Rosen
Universita Bocconi
Via Roentgen 1
20136 Milano
ITALY

Prof. Dr. Guy Rothblum
Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O. Box 26
Rehovot 76100
ISRAEL

Prof. Ron Rothblum
Technion
Haifa 32000
ISRAEL

Dr. Edlyn Teske
Sebastian-Bach-Str. 22
Leipzig 04109
GERMANY

Prof. Dr. Stefano Tessaro
Paul G. Allen School of Computer Science & Engineering
University of Washington
Box 352350
Seattle WA 98195-2350
UNITED STATES

Neekon Vafa
MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Dr. Vinod Vaikuntanathan
MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prashant Vasudevan
National University of Singapore
Computing 1
13 Computing Drive
Singapore 117417
SINGAPORE

Dr. Benjamin Wesolowski
École Normale Superieure de Lyon
46, Allee d'Italie
69364 Lyon Cedex 07
FRANCE

Prof. Dr. Daniel Wichs
College of Computer Science
Northeastern University
Boston MA 02115
UNITED STATES

Dr. David Wu
Department of Computer Science
University of Texas at Austin
2317 Speedway, D9500
Austin TX 78712
UNITED STATES

Or Zamir
School of Computer Science
Tel Aviv University
Ramat Aviv, Tel Aviv 69978
ISRAEL