A Research Programme in Post-Quantum Symmetric-Key Cryptography and the Foundations of Unclonable Primitives

Suprita Talnikar
Visiting Scientist, R. C. Bose Centre for Cryptology and Security

31 October 2025

Abstract

This document outlines a proposed programme of research in theoretical post-quantum cryptography. The work is organised into three primary areas. The first area involves a foundational study of advanced proof techniques in the Quantum Random Oracle Model, with the goal of developing a general framework for analysing cryptographic schemes that employ iterative optimisations. The second area consists of a systematic survey of the post-quantum security landscape for the Ascon family of algorithms, the new NIST standard for lightweight cryptography. The final area will explore the theoretical boundaries of unclonable cryptography, focusing on the efficiency of constructions that achieve strong multi-copy security guarantees.

1 Introduction: The Quantum Transition in Cryptography

The field of post-quantum cryptography is undergoing a period of rapid evolution, driven by two complementary goals: securing classical cryptographic paradigms against quantum adversaries and exploring entirely new functionalities enabled by quantum mechanics. The anticipated arrival of large-scale, fault-tolerant quantum computers necessitates a fundamental re-evaluation of our cryptographic infrastructure.

This research proposal addresses open problems at the intersection of these domains. It builds upon recent foundational work by Talnikar et al. in post-quantum proof techniques and the theory of unclonable primitives [25, 26, 27]. The proposed work is organised into three interconnected research thrusts. First, we will investigate a novel proof technique in the QROM, aiming to generalise its applicability to a wider class of practical cryptographic schemes. Second, we will conduct a systematic survey of the post-quantum security of the Ascon family of algorithms, a major new cryptographic standard for lightweight applications. Finally, we will push the theoretical boundaries of unclonable cryptography by analysing the efficiency of constructions that achieve strong multi-copy security guarantees.

1.1 The Post-Quantum Threat Landscape: A Dichotomy of Vulnerabilities

For decades, the security of global digital communication has rested on the computational difficulty of a few mathematical problems, primarily integer factorisation and the discrete logarithm problem [20]. Public-key cryptosystems like RSA and Elliptic Curve Cryptography (ECC), which are built upon these foundations, secure everything from financial transactions to national security communications. However, the theoretical work of Peter Shor in 1994 demonstrated that a sufficiently powerful quantum computer could solve these problems in polynomial time, rendering these

widely deployed cryptosystems insecure [24]. This creates an urgent and unusual threat scenario known as 'harvest now, decrypt later' (HNDL), or alternatively 'store now, decrypt later' [8]. An adversary can intercept and store vast quantities of encrypted data today, with the expectation of decrypting it years or decades from now once a cryptographically relevant quantum computer is available. The attack proceeds in three phases: harvesting, where encrypted data is collected through passive eavesdropping; storage, where the data is archived for the long term; and decryption, which occurs once quantum capabilities mature. For data with long-term confidentiality requirements—such as government secrets, intellectual property, and personal health information—this threat is immediate, even if the hardware to execute the final attack phase does not yet exist. The threat to symmetric-key cryptography, such as the Advanced Encryption Standard (AES), is more nuanced. While not vulnerable to Shor's algorithm, these primitives are susceptible to Grover's search algorithm, which provides a quadratic speed-up for unstructured search problems like brute-force key recovery [14]. This effectively halves the bit-security of a symmetric key; for instance, the 128-bit security of AES-128 is reduced to approximately 64 bits of security against a quantum adversary [16]. While this threat can be mitigated by doubling key sizes (e.g., migrating to AES-256), it underscores that no part of the cryptographic landscape is entirely unaffected by the quantum transition. This dichotomy of vulnerabilities—a catastrophic failure for public-key systems versus a quantifiable degradation for symmetric-key systems—creates a strategic 'urgency mismatch' that dictates the global cryptographic response. The threat to public-key cryptography necessitates an immediate and fundamental replacement of core infrastructure. In contrast, the threat to symmetric-key cryptography allows for a more measured, forward-looking migration to larger key sizes within existing, well-understood algorithmic families.

1.2 The Global Cryptographic Response: The NIST PQC Standardization Process

In response to this impending threat, the global cryptographic community, led by institutions like the U.S. National Institute of Standards and Technology (NIST), has been engaged in a multi-year effort to standardise a new suite of quantum-resistant algorithms. This public competition, initiated in 2016, solicited and evaluated algorithms from several distinct families, each based on mathematical problems believed to be hard for both classical and quantum computers. The main approaches included:

- **Lattice-based cryptography**, which relies on the hardness of problems such as Learning With Errors (LWE) in high-dimensional lattices. This family was favoured for its strong balance of performance and security.
- **Code-based cryptography**, based on the difficulty of decoding random error-correcting codes, a problem with a long history of resisting cryptanalysis.
- **Hash-based cryptography**, which derives its security from the properties of cryptographic hash functions. These schemes are noted for their conservative and well-understood security assumptions.
- **Multivariate cryptography**, based on the hardness of solving systems of multivariate polynomial equations over a finite field.

After several rounds of public scrutiny and analysis, this process culminated in the selection of a primary suite of post-quantum cryptographic (PQC) schemes, which were published as Federal Information Processing Standards (FIPS) in 2024. The initial standards are:

- **FIPS 203:** Module-Lattice-based Key-Encapsulation Mechanism (ML-KEM), based on the CRYSTALS-Kyber algorithm, for key establishment [20].
- **FIPS 204:** Module-Lattice-based Digital Signature Algorithm (ML-DSA), based on the CRYSTALS-Dilithium algorithm, for digital signatures [21].
- **FIPS 205:** Stateless Hash-based Digital Signature Algorithm (SLH-DSA), based on the SPHINCS+ algorithm, serving as an alternative signature standard with different underlying security assumptions [22].

While this standardisation effort provides a critical path forward for public-key infrastructure, it also illuminates new and urgent areas for fundamental research. The security of these new schemes, particularly in real-world implementations, must be established with the highest degree of rigour. This requires the development of sophisticated proof techniques capable of modelling the full power of a quantum adversary.

1.3 The Quantum Random Oracle Model: A New Proof Paradigm

Classical security proofs are often conducted in the Random Oracle Model (ROM), a theoretical framework where a cryptographic hash function is idealised as a perfect, publicly accessible random function. In the ROM, any party, including an adversary, can only learn the output $H(x)$ by making an explicit, classical query for the input x. This idealisation enables powerful proof techniques like lazy sampling, where a reduction can invent oracle outputs on the fly, and reprogramming, where a reduction can change an oracle's output for a specific input to facilitate a proof. However, the ROM is insufficient for post-quantum security analysis. A quantum adversary can evaluate a hash function on a superposition of inputs, gaining information about many input-output pairs simultaneously. To model this capability, Boneh et al. introduced the Quantum Random Oracle Model (QROM) in 2011 [5]. In the QROM, an adversary can make superposition queries of the form $\sum_x \alpha_x|x\rangle|0\rangle \rightarrow \sum_x \alpha_x|x\rangle|H(x)\rangle$. This transition to the QROM invalidates many classical proof techniques. A reduction cannot simply "record" an adversary's superposition query to see which inputs were queried, as this would violate the no-cloning theorem. Measuring the query to learn its contents would collapse the superposition, a disturbance the adversary could detect, thus invalidating the proof's simulation. This makes techniques like adaptive reprogramming exceptionally difficult and has led to formal separation results: cryptographic schemes that are provably secure in the ROM but demonstrably insecure in the QROM. Consequently, establishing security in the QROM has become a central and formidable challenge in modern cryptography.

1.4 Quantum Mechanics as a Cryptographic Enabler: The No-Cloning Principle

Simultaneously, the principles of quantum mechanics offer opportunities that go beyond merely defending against quantum attacks. The no-cloning theorem, a fundamental tenet of quantum physics, states that it is impossible to create an identical, independent copy of an arbitrary, unknown quantum state [30]. This principle arises from the linearity of quantum mechanics. If a universal cloning operator U existed such that for any two states $|\psi\rangle$ and $|\phi\rangle$, it held that $U(|\psi\rangle|s\rangle) = |\psi\rangle|\psi\rangle$ and $U(|\phi\rangle|s\rangle) = |\phi\rangle|\phi\rangle$ (where $|s\rangle$ is a standard blank state), the unitarity of U would require that the inner product be preserved: $\langle\psi|\phi\rangle = (\langle\psi|\phi\rangle)^2$. This equation holds only if $\langle\psi|\phi\rangle$ is 0 or 1, meaning the cloner could only copy states that are orthogonal to each other, not arbitrary states. This physical limitation provides a powerful cryptographic tool, giving rise to the field of unclonable

cryptography. The concept was first proposed by Stephen Wiesner circa 1970 in his work on "conjugate coding," which introduced the idea of quantum money [29]. By encoding information as a quantum state, that information can be made inherently copy-proof—a functionality unattainable in the classical world, where any bitstring can be perfectly duplicated [1]. This paradigm seeks to create primitives, such as quantum money or copy-protected software, that are physically impossible to duplicate [7].

1.5 Research Contributions and Proposal Structure

This research proposal addresses open problems at the confluence of these developments. It builds upon recent foundational work in post-quantum proof techniques and the theory of unclonable primitives. The proposed work is organised into three interconnected research thrusts, each aiming to make a distinct contribution to the field:

1. **Advancing QROM Proof Techniques:** We will conduct an independent theoretical study of advanced QROM proof techniques, taking the recently introduced 'adaptive multi-point reprogramming' technique as a starting point [17]. The goal is to generalise this method to analyse a broader class of practical cryptographic schemes that employ iterative optimisations, which are currently beyond the reach of formal analysis.

2. **Systematic Security Survey of Ascon:** We will perform a comprehensive survey and comparative analysis of the post-quantum security of the Ascon family of algorithms, the new NIST standard for lightweight cryptography [9]. By synthesising results from theoretical proofs [18] and concrete resource estimations [23], this work will provide a unified and actionable security assessment for implementers and standards bodies.

3. **Exploring the Frontiers of Unclonable Cryptography:** We will investigate the theoretical efficiency of primitives that achieve strong, multi-copy unclonability guarantees. Building on a recent breakthrough generic compiler [7], this research will involve a rigorous overhead analysis of such generic transformations and will aim to design more efficient, direct constructions.

The remainder of this proposal is structured as follows. Section 2 details the research objectives related to QROM proof techniques. Section 3 outlines the scope and methodology for the security survey of the Ascon family. Section 4 presents the research plan for unclonable cryptography. Finally, Section 5 discusses the expected methodologies and dissemination plan for the research findings.

2 Advanced Proof Techniques in the Quantum Random Oracle Model

2.1 Background: MPCitH Signatures and the Challenge of Iterative Optimisations

A significant number of modern post-quantum signature schemes are built using the Multi-Party-Computation-in-the-Head (MPCitH) paradigm [15]. Secure Multi-Party Computation (MPC) is a subfield of cryptography that enables a set of parties to jointly compute a function over their private inputs without revealing those inputs to one another. The MPCitH paradigm, introduced by Ishai, Kushilevitz, Ostrovsky, and Sahai, provides a generic transformation from an MPC protocol into

a zero-knowledge proof of knowledge. In this transformation, a prover demonstrates knowledge of a secret witness (e.g., a signing key) by simulating an MPC protocol 'in their head', where the witness is split into shares among virtual parties. The prover commits to the execution views of all parties and then, upon receiving a challenge from the verifier, opens a subset of these views. The verifier accepts if the revealed views are consistent and indicate a valid computation. To create a digital signature, this interactive proof is made non-interactive using the Fiat-Shamir transform [12]. Instead of receiving a random challenge from the verifier, the prover generates the challenge by hashing the commitments and other public information. The resulting non-interactive proof, which includes the message being signed in the hash input, constitutes the signature [10]. The MPCitH paradigm and its variants, such as Threshold Computation in the Head (TCitH) [11] and VOLE-in-the-Head (VOLEitH) [2], have proven to be a particularly fruitful approach for constructing signatures based on a wide range of post-quantum hardness assumptions. As shown in Table 1, this is reflected in the large number of MPCitH-based schemes submitted to the NIST Post-Quantum Cryptography (PQC) standardisation process for additional signature schemes [17].

Table 1: A Selection of NIST PQC Round 2 Signature Candidates based on MPCitH Variants.

Scheme	**Paradigm**	**Optimisations**	**Underlying Problem**
Mirath	TCitH	Grinding, Rejection Sampling	MinRank
RYDE	TCitH	Grinding, Rejection Sampling	Rank Syndrome Decoding
MQOM	TCitH	Grinding	Multivariate Quadratic (MQ)
FAEST	VOLEitH	-	AES Preimage
PERK	MPCitH	-	Permuted Kernel Problem
SDitH	VOLEitH	-	Syndrome Decoding

To improve efficiency and security, practical implementations of these schemes often employ optimisations such as grinding and rejection sampling.

- **Grinding** is a proof-of-work mechanism where the signer repeatedly computes a challenge hash with an incrementing counter until the output satisfies a certain public property (e.g., having a number of leading zero bits). This increases the cost for an adversary to find a challenge that is useful for forgery [17, 2].
- **Rejection sampling** is a statistical technique used to reduce signature size or ensure that the signature distribution does not leak secret information. The signer generates signature candidates and rejects them until one is found that satisfies a public condition (e.g., being shorter than a certain threshold) [17, 19].

While these optimisations are crucial for practical performance, they introduce significant challenges for security proofs in the QROM. Standard proof techniques, which often rely on reprogramming a random oracle at a single point, are unable to model these iterative, probabilistic procedures. This has created a disconnect between provable security and the practical implementations of some of the most promising post-quantum candidates [17].

2.2 A Survey of Modern QROM Proof Techniques

The difficulty of proving security in the QROM stems from the challenge of simulating a quantum-accessible oracle for an adversary while still being able to extract information or embed challenges.

The evolution of QROM proof techniques reflects a continuous effort to reclaim the power of classical ROM techniques in this more complex setting.

- **History-Free Proofs:** This early approach sidesteps the reprogramming problem by defining the random oracle's responses "a-priori" in a way that is already consistent with the proof's requirements. The reduction can answer queries without knowing the adversary's query history. However, this method is not generic; proofs are often highly complex, tailor-made for specific schemes, and may not exist for many important constructions, including standard Fiat-Shamir signatures.

- **The One-Way to Hiding (O2H) Lemma:** Introduced by Ambainis et al., the O2H lemma provides a more general tool for handling reprogramming-like scenarios [3]. It bounds an adversary's ability to distinguish between an oracle output derived from a one-way function and a truly random value. While powerful, early versions of the O2H lemma and related techniques often led to non-tight security reductions, introducing a square-root loss in security parameters, which is undesirable for practical schemes. Subsequent work has focused on developing tighter variants, such as "double-sided" O2H and the "Measure-Rewind-Measure" technique.

- **Tightly Secure Adaptive Reprogramming:** A more recent line of work directly tackles the reprogramming problem to achieve tight security bounds. The technique of Grilo et al. provides a method for tight adaptive reprogramming, provided the reprogramming is triggered by a classical query (as in a signature forgery attempt) and the location to be reprogrammed has sufficient entropy [13]. This approach successfully recovers much of the power and conceptual simplicity of classical ROM proofs.

This evolutionary path—from avoidance (history-free), to approximation (non-tight O2H), to direct and tight solutions (adaptive reprogramming)—shows a maturation of the field. The work of Kosuge and Xagawa, which extends this to handle multiple reprogramming points simultaneously, sits at the cutting edge of this trend, enabling the analysis of iterative optimisations like grinding [17].

2.3 Research Objective

This research proposes an independent theoretical study of advanced QROM proof techniques capable of analysing constructions that use grinding and rejection sampling [17]. The investigation will take the recently introduced 'adaptive multi-point reprogramming' technique of Kosuge and Xagawa as a starting point [17]. This technique, an extension of the adaptive reprogramming lemma of Grilo et al. [13], provides a method for reprogramming a random oracle on multiple points in a single logical step, as illustrated in Figure 1. This was instrumental in providing the first formal Existential Unforgeability under Chosen-Message Attack (EUF-CMA) security proofs for Mirath and RYDE [17].

The primary goal of this research objective is to generalise this technique into a broader theoretical framework for analysing iterative cryptographic processes. A key challenge to be addressed is extending the analysis to schemes with 'correlated' internal structures. An example is the signature scheme MQOM, which is also based on the TCitH framework but uses correlated Goldreich-Goldwasser-Micali (GGM) trees to reduce signature size [4, 17]. This correlation, which depends on the secret key, prevents a direct application of the existing reprogramming technique. Successfully developing a framework to handle such cases would significantly broaden the class of practical, optimised schemes that can be formally analysed for post-quantum security.

Figure 1: Illustrating the difference between single-point and batched adaptive reprogramming.

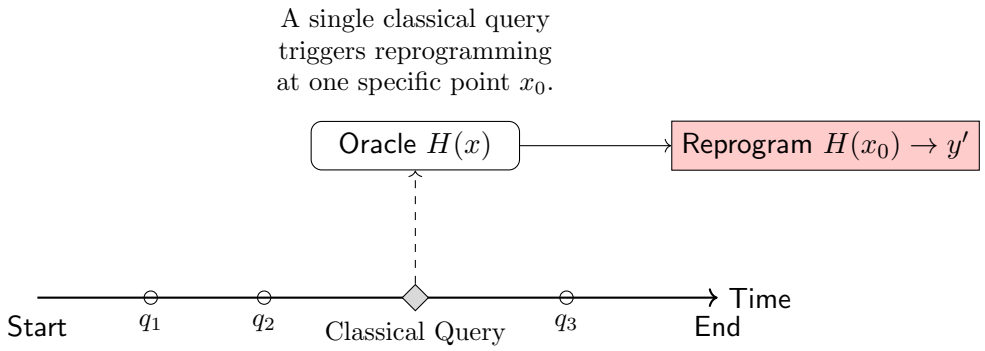

a) Single-Point Reprogramming

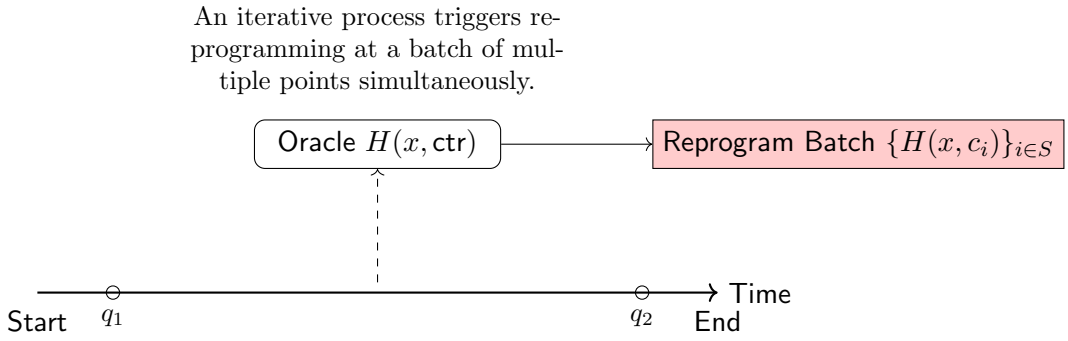

b) Batched Adaptive Reprogramming

3 A Comprehensive Security Survey of the Ascon Family

3.1 Background: Ascon and the Sponge Construction

The Ascon family of algorithms was selected as the winner of the National Institute of Standards and Technology (NIST) Lightweight Cryptography (LWC) standardisation process in 2023 [9]. The goal of the LWC process was to standardise algorithms suitable for resource-constrained environments such as Internet of Things (IoT) devices, embedded systems, and RFID tags. Ascon was chosen for its excellent all-around performance, efficiency, and strong security record, including its prior selection in the CAESAR competition. Given its expected long-term deployment horizon, a consolidated and robust understanding of its security posture against quantum adversaries is of high importance to the cryptographic community. Ascon is based on the cryptographic sponge construction, a versatile mode of operation for a fixed-length permutation, f, that can process variable-length inputs to produce variable-length outputs. The construction operates on a b-bit internal state, which is conceptually divided into two parts:

- **Rate (r):** The outer part of the state, comprising r bits. Input message blocks are XORed into the rate, and output blocks are extracted from it. A larger rate corresponds to higher throughput.
- **Capacity (c):** The inner part of the state, comprising $c = b - r$ bits. The capacity is shielded from direct interaction with input and output blocks, and its size is the primary determinant of the construction's security against generic attacks.

The operation proceeds in two phases. In the *absorbing phase*, the input message is padded and divided into r-bit blocks, which are sequentially XORed into the rate portion of the state, with an application of the permutation f after each block. In the *squeezing phase*, r-bit output blocks are extracted from the rate portion, again interleaved with applications of f. Ascon uses a 320-bit permutation with varying rate and capacity parameters for its different modes.

3.2 Generic Quantum Attacks on Sponge Constructions

The evaluation of post-quantum security for symmetric-key schemes like Ascon rests on two complementary pillars: theoretical proofs against generic attacks and concrete resource estimations for specific attacks.

1. **Theoretical Provable Security:** This 'top-down' approach aims to establish formal mathematical proofs of security against generic classes of quantum adversaries. The analysis is typically performed in a model where the adversary has classical access to the overall construction but quantum access to the underlying primitive (e.g., the unkeyed permutation) [18, 3].
2. **Concrete Resource Estimation:** This 'bottom-up' approach involves designing optimised quantum circuits for specific attacks and calculating the tangible quantum resources (e.g., qubits, gate counts, circuit depth) required to execute them [23, 16]. This provides a practical upper bound on the scheme's security.

When the upper bound from the best-known concrete attack is close to the lower bound from the formal proof, it creates a narrow 'security bracket' that provides a high degree of confidence in the true security level of the algorithm [27]. For sponge constructions, two primary generic quantum attacks define the theoretical security bounds:

- **Grover's Key Search:** For any keyed primitive with a k-bit key, Grover's algorithm provides a generic key-recovery attack with complexity $O(2^{k/2})$ [14]. This attack treats the entire construction as a black box and is the source of the $k/2$ term in the security bound.
- **Quantum State Recovery Attacks:** The internal state of the sponge construction makes it vulnerable to structural attacks. A quantum adversary can attempt to find "internal collisions"—two distinct input sequences that lead to the same c-bit state in the capacity. Quantum collision-finding algorithms, such as the Brassard-Høyer-Tapp algorithm, can solve this problem with complexity $O(2^{c/3})$ [6]. Such an attack can be used to distinguish the sponge from a random oracle and can be leveraged to break higher-level security properties. This threat is the source of the $c/3$ term in the security bound [18]. To be secure, a sponge construction must possess the "collapsing" property, a post-quantum strengthening of collision resistance that ensures it is hard to find a superposition of distinct inputs that map to the same output [28].

The overall security is thus determined by the weaker of these two attacks, highlighting a fundamental design duality: security depends on both the secrecy of the external key (k) and the integrity of the internal state, governed by the capacity (c).

3.3 Research Objective

This research objective is to conduct a comprehensive survey and comparative analysis of the post-quantum security of the Ascon family, synthesising the results from both theoretical proofs and concrete costing analyses. The survey will provide a unified view of the security levels for the primary variants, as summarised in Table 2.

Table 2: Parameters and Post-Quantum Security Bounds for Ascon Variants [18, 27].

Variant	Key Size (k)	Capacity (c)	Rate (r)	PQ Security Bound (bits)
Ascon-128	128	192	64	$\min\{128/2, 192/3\} = 64$
Ascon-128a	128	192	128	$\min\{128/2, 192/3\} = 64$
Ascon-80pq	160	256	64	$\min\{160/2, 256/3\} \approx 80$

The survey will focus on clarifying the security bounds offered by each variant and the crucial role that both key size (k) and sponge capacity (c) play in determining overall post-quantum security. For sponge-based authenticated encryption, the security against a generic quantum adversary is fundamentally limited by $\min\{k/2, c/3\}$ [18]. The $k/2$ term arises from Grover's key search, while the $c/3$ term arises from quantum state recovery attacks that exploit the sponge structure. This survey will explicitly compare the theoretical bounds with concrete resource estimations, as illustrated in Figure 2, to provide a clear and actionable assessment for implementers. The analysis will confirm that while Ascon-128 and Ascon-128a offer approximately 64 bits of post-quantum security, falling short of NIST's Level 1 security strength, the Ascon-80pq variant provides a robust 80 bits of security, making it a suitable choice for new post-quantum systems [27, 23].

4 Frontiers in Unclonable Cryptography

4.1 Background: From Single-Copy to Multi-Copy Security

Unclonable cryptography uses principles of quantum mechanics, most notably the no-cloning theorem, to achieve functionalities that are impossible in the classical world [1]. The field's origins

Figure 2: Comparing quantum attack costs for Ascon variants against the NIST Level 1 benchmark [23, 27].

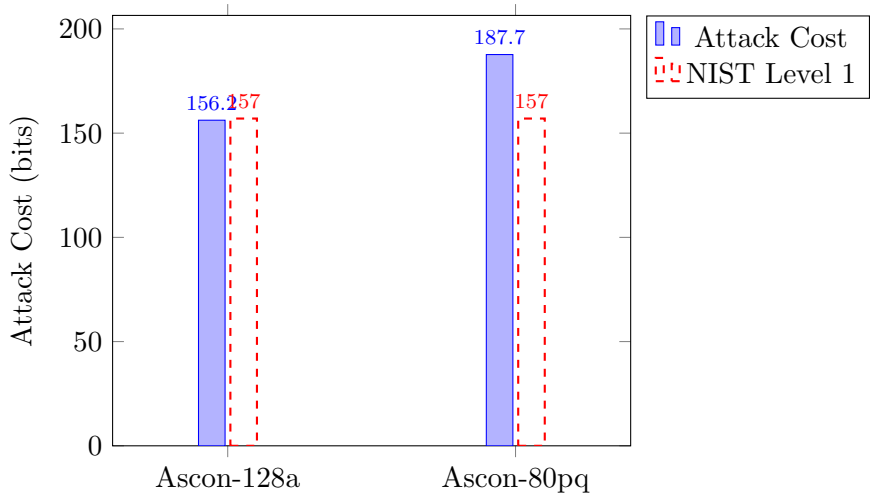

trace back to Stephen Wiesner's work on "conjugate coding" circa 1970, which introduced the concept of quantum money [29]. Wiesner's scheme proposed encoding information in quantum states (e.g., polarized photons) using one of two randomly chosen conjugate bases. A bank, knowing the sequence of bases, could verify the banknote, but a counterfeiter, ignorant of the bases, could not measure and reproduce the states without introducing detectable errors. A central challenge in this field is achieving a strong and realistic security guarantee. Most early research focused on **single-copy security**, often denoted as '$1 \to 2$' security. This notion guarantees that an adversary given one legitimate copy of a quantum object cannot produce two valid copies [7]. While this is a necessary first step, it is insufficient against a more powerful adversary who obtains multiple copies of the same object. The primary threat in the multi-copy setting is **quantum state tomography**. This is a general experimental procedure for reconstructing the full classical description (the density matrix) of an unknown quantum state. The process requires access to many identical copies of the state, upon which a comprehensive set of measurements are performed. By collecting statistics on the measurement outcomes, an adversary can solve for the density matrix. Once this classical description is learned, the adversary can use it to produce an arbitrary number of new copies, completely breaking the unclonability property. This threat motivates a significantly stronger guarantee: **multi-copy security**, or '$q \to q + 1$' security. This notion requires that an adversary given q legitimate copies of an object is prevented from creating $q + 1$ copies, for any polynomial q [7]. The distinction between single-copy and multi-copy security represents a fundamental "learnability threshold." Schemes that are vulnerable to tomography are based on "learnable" quantum states. In contrast, a multi-copy secure primitive must be constructed from quantum states that are computationally "unlearnable"—that is, it must be computationally intractable to derive a classical description of the state even when given polynomially many copies [7]. A recent foundational breakthrough by Çakan et al. provided the first generic compiler that transforms certain single-copy secure primitives into ones that achieve full, unbounded multi-copy security [7]. This result proves the feasibility of multi-copy security for a broad class of primitives by providing a generic "unlearnability wrapper." However, as is common with generic transformations, the compiler may introduce significant overhead, motivating the search for more efficient, direct constructions.

Table 3: Evolution of Unclonable Primitives and Security Models.

Primitive/ Concept	**Verification Model**	**Core Principle**	**Security Model**	**Known Limitations/ Open Problems**
Wiesner's Money	Private-Key (Bank only)	Conjugate Coding	Information-Theoretic (Single-Copy)	Not publicly verifiable; Vulnerable to adaptive verification attacks.
Aaronson-Christiano (Oracle)	Public-Key	Hidden Subspaces	Unconditionally Secure (Single-Copy)	Idealised black-box model; requires a secure instantiation.
Aaronson-Christiano (Explicit)	Public-Key	Polynomial Zero-Sets	Computationally Secure (Conjectured)	Broken by algebraic and quantum cryptanalysis.
Multi-Copy Secure Primitives	Public/Private	Varies	Computational ($q \to q+1$)	Constructing efficient instantiations from standard assumptions.

4.2 Research Objectives

This research will focus on the theoretical efficiency of multi-copy secure schemes, building directly upon the work of Çakan et al. The objectives are twofold:

- **Theoretical Overhead Analysis of Generic Compilers:** To conduct a rigorous theoretical analysis of the efficiency and complexity overhead introduced by the generic compiler of Çakan et al. [7]. This analysis will focus on the algorithmic complexity, query complexity, and security parameter trade-offs from a theoretical computer science perspective. The goal is to establish formal bounds on the overhead incurred by such generic transformations in terms of qubit requirements, gate complexity, and classical computation.
- **Design and Analysis of Direct Constructions:** To design and formally analyse new, direct constructions of multi-copy secure primitives. The objective is to create schemes that are purpose-built for multi-copy security, with the aim of achieving superior theoretical efficiency compared to the bounds achievable via the generic compiler. This will involve exploring novel quantum state encodings and verification procedures that are inherently resistant to tomographic attacks by embedding computational hardness directly into the structure of the states themselves.

5 Methodologies

This research programme will employ a range of methodologies from theoretical and quantum cryptography:

- **Formal Proofs in the Quantum Random Oracle Model:** The investigation into reprogramming techniques will involve the construction of rigorous, game-based security proofs to formally establish the security of new applications or extensions of the multi-point reprogramming lemma. This will require a deep understanding of existing QROM proof techniques, including the One-Way to Hiding (O2H) Lemma and prior work on adaptive reprogramming [3, 13]. The proofs will follow a reductionist argument, proceeding via a sequence of game hops

to relate the adversary's success probability to the hardness of an underlying computational problem.

- **Systematic Literature Review and Comparative Analysis:** The survey of Ascon's security will involve a comprehensive review of the existing academic literature, including theoretical proofs, concrete attack implementations, and performance benchmarks. The findings will be synthesised and presented in a comparative framework that juxtaposes theoretical lower bounds on security [18] with concrete upper bounds from resource estimations [23] to establish a clear "security bracket" for each variant.

- **Construction and Efficiency Analysis:** The research on multi-copy security will involve both the design of new cryptographic protocols and a detailed analysis of their concrete efficiency. This will include asymptotic complexity analysis of qubit and gate counts (with a focus on T-gate counts for fault-tolerant implementations) for the quantum components, as well as the computational cost of the classical components, to compare against the bounds achievable via existing generic transformations [7].

6 Dissemination

The findings from this research programme are expected to be suitable for publication in reputable academic journals and for presentation at leading international conferences in the field of cryptology.

References

[1] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, pages 41–60, 2012.

[2] Carlos Aguilar-Melchor, Andreas Hülsing, David Joseph, Christian Majenz, Eyal Ronen, and Dongze Yue. SDitH in the QROM. In *Advances in Cryptology – ASIACRYPT 2023*, pages 317–350. Springer, 2023.

[3] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In *Advances in Cryptology – CRYPTO 2019*, pages 269–295. Springer, 2019.

[4] Ryad Benadjila, Thibauld Feneuil, and Matthieu Rivain. MQ on my Mind: Post-Quantum Signatures from the Non-Structured Multivariate Quadratic Problem. In *2024 IEEE 9th European Symposium on Security and Privacy (EuroS&P)*, pages 468–485, 2024.

[5] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology – ASIACRYPT 2011*, pages 41–69. Springer, 2011.

[6] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum counting. In *Proceedings of the 25th International Colloquium on Automata, Languages, and Programming*, pages 820–831. Springer, 1998.

[7] Alper Çakan, Vipul Goyal, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Multi-copy security in unclonable cryptography. Cryptology ePrint Archive, Report 2025/1921, 2025. `https://eprint.iacr.org/2025/1921`.

[8] Cybersecurity and Infrastructure Security Agency. Prepare for post-quantum cryptography. CISA Insights, 2022. https://www.cisa.gov/resources-tools/resources/prepare-post-quantum-cryptography.

[9] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer. ASCON v1.2: Lightweight authenticated encryption and hashing. *Journal of Cryptology*, 34(3):33, 2021.

[10] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Efficient NIZKs and signatures from commit-and-open protocols in the QROM. In *Advances in Cryptology – CRYPTO 2022*, pages 729–757. Springer, 2022.

[11] Thibauld Feneuil and Matthieu Rivain. Threshold linear secret sharing to the rescue of MPC-in-the-Head. In *Advances in Cryptology – ASIACRYPT 2023*, pages 441–473. Springer, 2023.

[12] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology – CRYPTO '86*, pages 186–194. Springer, 1987.

[13] Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the QROM. In *Advances in Cryptology – ASIACRYPT 2021*, pages 637–667. Springer, 2021.

[14] Lov K. Grover. A fast quantum mechanical algorithm for database search. *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996.

[15] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multi-party computation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 21–30, 2007.

[16] Samuel Jaques, Michael Naehrig, Martin Roetteler, and Fernando Virdia. Implementing Grover oracles for quantum key search on AES and LowMC. In *Advances in Cryptology – EUROCRYPT 2020*, pages 280–310. Springer, 2020.

[17] Haruhisa Kosuge and Keita Xagawa. The new security proofs of MPC-in-the-Head signatures in the quantum random oracle model. To appear in Proceedings of Computer Security Symposium (CSS 2025). Preliminary version available as IACR ePrint Archive, Report 2025/1999, 2025. https://eprint.iacr.org/2025/1999.

[18] Nathalie Lang, Stefan Lucks, Bart Mennink, and Suprita Talnikar. Security of the Ascon authenticated encryption mode in the presence of quantum adversaries. Cryptology ePrint Archive, Report 2025/411, 2025. https://eprint.iacr.org/2025/411.

[19] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Advances in Cryptology – EUROCRYPT 2012*, pages 738–755. Springer, 2012.

[20] National Institute of Standards and Technology. Fips 203, module-lattice-based key-encapsulation mechanism standard. Technical report, U.S. Department of Commerce, 2024.

[21] National Institute of Standards and Technology. Fips 204, module-lattice-based digital signature standard. Technical report, U.S. Department of Commerce, 2024.

[22] National Institute of Standards and Technology. Fips 205, stateless hash-based digital signature standard. Technical report, U.S. Department of Commerce, 2024.

[23] Yujin Oh, Kyungbae Jang, Anubhab Baksi, and Hwajeong Seo. Depth-optimized quantum circuits for ASCON: AEAD and HASH. *Mathematics*, 12(9):1337, 2024.

[24] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 124–134. IEEE Computer Society, 1994.

[25] Suprita Talnikar, Anandarup Roy, and Kouichi Sakurai. The mirror breaks: Reformulating cryptographic proofs for quantum realities. In Proceedings of Computer Security Symposium (CSS 2025), 2025.

[26] Suprita Talnikar, Anandarup Roy, and Kouichi Sakurai. Post-quantum security: Bridging quantum mechanics and cryptographic proofs. In Proceedings of Computer Security Symposium (CSS 2025), 2025.

[27] Suprita Talnikar, Anandarup Roy, and Kouichi Sakurai. Provable and practical: Establishing post-quantum security bounds for ascon. In *Proceedings of Computer Security Symposium (CSS 2025)*, 2025.

[28] Dominique Unruh. Collapsing-secure encryption. In *Theory of Cryptography Conference*, pages 565–595. Springer, 2016.

[29] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983.

[30] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.