Gap Analysis on the Limits of Post-Quantum Cryptography

Juha Partala, Pauli Pauna, Sini Ruohomaa, John Preuß Mattsson, Valtteri Lipiäinen, Aleksi Kalsta, Kimmo Järvinen, Petri Jehkonen, Sanish Gurung, Gizem Akman, Valtteri Niemi, Erik Hieta-aho

August 14, 2026

Contents

1 Introduction **3**
1.1 Organization 4

2 Security Models and Transformations **5**
2.1 Random oracle model 5
2.1.1 Reductions in QROM 6
2.2 Transformations 6
2.2.1 Universal Composability 6
2.2.2 Fujisaka-Okamoto 7
2.2.3 Fiat-Shamir 8
2.2.4 Fischlin 9
2.3 One-way functions 9
2.4 Indistinguishability 9
2.5 Unforgeability 11

3 Symmetric cryptography **11**
3.1 Pseudorandom Generators 11
3.2 Pseudorandom Functions and Permutations 12
3.3 Hash Functions 13
3.3.1 Indifferentiability 14
3.4 Message Authentication 15
3.5 Symmetric encryption 15
3.5.1 IND-CPA 16
3.5.2 IND-CCA 16

4 Asymmetric Cryptography **17**
4.1 Key exchange and encapsulation 18
4.1.1 Lattice-based schemes 19
4.1.2 Code-based schemes 20
4.1.3 Hybrid methods for key exchange and encapsulation . . . 20
4.2 Public-key Encryption 21
4.3 Digital signatures 22
4.3.1 Hash-and-Sign 22
4.3.2 Fiat-Shamir signatures 23
4.3.3 Lattice-based signatures 23
4.3.4 Hash-based signatures 24
4.3.5 Hybrid methods 25

5 Advanced Cryptography **27**
5.1 Oblivious Transfer 27
5.2 Homomorphic Encryption 28
5.3 Password Authenticated Key Exchange 28
5.4 Multi-Party Key Exchange 29
5.5 Secure Multi-Party Computation 29
5.6 Zero-knowledge proofs 30
5.7 Polynomial Commitment Schemes 31
5.8 Anonymous credentials 31

6 Cryptographic Protocols **32**
6.1 SSH 32
6.2 TLS 1.3 34
6.3 Signal 36
6.4 Tor 38
6.5 European Digital Identity Wallet 40

7 Limiting factors in hardware implementations of PQC **42**
7.1 Effects of larger keys, signatures and ciphertexts 42
7.2 Microarchitecture 42
7.3 Side channel analysis and PQC algorithms 43
7.3.1 Current Results in Securing PQC Against Side-Channel Attacks 43
7.3.2 Observed Vulnerabilities 43
7.3.3 Countermeasure Strategies 44
7.3.4 Hardware-Level Protections 44
7.3.5 Timing attack on HQC 44
7.4 User interface differences for applying the cryptographic primitives 44
7.4.1 Key establishment constraints 45
7.4.2 Signing constraints 45

8 Transition **46**
8.1 Limits to Cryptographic Asset Visibility 46
8.2 Causes of Limited Cryptographic Agility 47
8.3 Potential Execution Gap Between Policy Timelines and Organizational Practice 48
8.4 Organizational Limits to Cryptographic Migration 49

1 Introduction

The emergence of large-scale quantum computers poses an existential threat to many widely deployed cryptographic primitives that underpin today's digital infrastructure. Shor's algorithm can break the integer-factorisation and discrete-logarithm problems on which RSA, Diffie–Hellman, and elliptic-curve schemes rely, while Grover's search provides a generic quadratic speed-up against symmetric constructions. In response, the cryptographic community has embarked on standardization efforts, as well as on academic research into algorithms that are believed to resist quantum adversaries.

The transition to post-quantum cryptography (PQC) is far from a simple "drop-in replacement". Security proofs that hold in the classical random-oracle model (ROM) often collapse under the more powerful quantum random-oracle model (QROM). Widely used cryptographic transformations such as Fujisaki-Okamoto (FO) or Fiat-Shamir acquire new subtleties. Many higher-level protocols, such as TLS and SSH, rely on composability guarantees that have not yet been lifted to the quantum setting. Moreover, the practical deployment of PQC is constrained by a lack of PQC options, dramatically larger keys, ciphertexts and signatures, heterogeneous hardware requirements, and limited cryptographic agility in legacy systems.

The purpose of this report is therefore twofold:

1. **Systematic Gap Identification**. We survey the current state-of-the-art of quantum security of cryptographic building blocks (security models, symmetric primitives, asymmetric schemes, advanced constructions such as oblivious transfer or functional encryption, and end-to-end protocols). We identify the theoretical limits that have been proven (e.g., separations between ROM and QROM, tightness of reductions, impossibility results) and highlight open gaps where applications are lacking PQC primitives, security proofs are missing, assumptions remain unverified, or performance constraints render a construction impractical.

2. **Application-Level Needs**. By mapping these gaps onto real-world use cases we identify the applications that require robust PQC solutions: key-exchange and encapsulation mechanisms for TLS/SSH handshakes, digital signatures for public-key infrastructure (PKI) and code signing, authenticated encryption in low-power IoT and mobile devices, privacy-preserving protocols such as anonymous credential, secure multi-party computation and private group messaging.

By providing a unified view of the state of knowledge and the open problems across the entire cryptographic stack, this report aims to guide both academic investigations and engineering efforts toward the most critical gaps. Ultimately, closing these gaps will be essential for achieving a secure, performant, and future-proof digital ecosystem in the quantum era.

1.1 Organization

The report is divided into eight sections:

- **Section 2: Security Models & Transformations** reviews the random oracle landscape, recent advances in history-free reductions, compressed oracles, and quantum-aware rewinding techniques. It also analyses the applicability of universal composability, Fujisaki-Okamoto, Fiat-Shamir, and Fischlin transforms in a post-quantum world.

- **Section 3: Symmetric Cryptography** examines pseudorandom generators, functions and permutations, hash constructions, MACs, and symmetric encryption modes. Particular emphasis is placed on indistinguishability notions under different quantum adversary models (Q1 and Q2) and the security status of symmetric encryption schemes against quantum adversaries.

- **Section 4: Asymmetric Cryptography** surveys limitations related lattice-based, code-based and hybrid key-encapsulation mechanisms; public-key encryption schemes; and digital signatures (hash-and-sign, Fiat-Shamir constructions, NIST PQC candidates). We discuss decryption-failure handling, side-channel leakage in FO-transformed KEMs, and the trade-offs between computational efficiency and communication overhead.

- **Section 5: Advanced Cryptography** surveys primitives that sit above the basic encryption/signature layer: oblivious transfer, homomorphic encryption, PAKE, multi-party key exchange, MPC, zero-knowledge proofs, polynomial commitments, and anonymous credentials. The sections highlights which of these already enjoy quantum-secure constructions and where research is still needed.

- **Section 6: Cryptographic Protocols** analyses concrete internet-scale protocols (SSH, TLS 1.3, Signal, Tor, EUDIW), comparing the maturity of their post-quantum transition and identifying protocol-specific security properties that remain open research questions.

- **Section 7: Implementation & Hardware Constraints** details how larger key and ciphertext sizes affect memory footprints, microarchitectural design, side-channel resistance, and energy consumption, particularly in constrained environments such as IoT, smart cards and mobile devices.

- **Section 8: Transition** discusses organisational limits to cryptographic agility, the scarcity of comprehensive cryptographic inventories (CBOMs), policy-implementation gaps, and the risk that migration timelines prescribed by standards bodies outpace realistic deployment schedules for post-quantum security.

2 Security Models and Transformations

2.1 Random oracle model

In the random oracle model (ROM), hash functions are idealized as a theoretical construct where each input maps to an output uniformly at random. In post-quantum contexts, we extend this concept to the quantum random oracle model (QROM), which allows inputs in quantum superposition, thus enabling potentially more complex interactions than classical queries. The QROM is still not fully comprehended within the cryptographic community. The ability of quantum systems to query oracles using superpositions can reveal information that would remain hidden under classical querying methods. This extension was initially introduced by Boneh et al. who demonstrated that a cryptographic scheme could be secure in ROM yet vulnerable when analyzed under QROM [37].

Yamakawa and Zhandry conducted an in-depth analysis of the differences between security properties of cryptographic schemes in the ROM and QROM [241]. Their research introduced a novel concept known as Proof of Quantum Access to a Random Oracle (PoQRO), which allows protocols to demonstrate quantum access to a random oracle to a classical verifier. This breakthrough led them to construct digital signature and public key encryption schemes that exhibit security in the ROM but are compromised when subjected to QROM analysis. In subsequent work, Yamakawa and Zhandry demonstrated that there exists a fundamental separation between ROM and QROM security even without relying on computational assumptions, providing counterexamples to illustrate this disparity [242].

A critical property of ROM is extractability, which ensures that the simulator in the proof can recover preimages from quantum queries. Achieving this in QROM presents significant challenges. However, it becomes feasible when cryptographic schemes are designed with "oracle-masking" [215]. This technique allows security proofs for unmodified transformations with explicit rejection in the QROM, thus bridging some gaps between ROM and QROM analyses.

In the quantum setting, various models extend the classical random oracle model to incorporate quantum properties. One notable extension is the Quantum Haar ROM, where both prover and verifier have access to a Haar random quantum oracle and its inverse [59]. This framework establishes a hierarchy of quantum random oracle models, reflecting varying levels of theoretical strength and practical feasibility. The hierarchy includes: the Common Haar Random State as the strongest but least instantiable model; Unitary oracles offering strong but challenging implementations; Isometry oracles which are weaker but

easier to construct; Classical oracles with quantum access as the most standard yet weakest model.

2.1.1 Reductions in QROM

Reductions in the QROM are harder to construct than in the classical ROM. Boneh et al. proposed "history-free reductions" where oracle responses can be determined without considering previous queries, to facilitate the transition of ROM proofs into the QROM [37]. When converting a proof from ROM to QROM, a reduction loss in security typically occurs, which can range from square-root of the input length to more significant losses. Furthermore, many ROM proofs rely on query history, complicating their conversion to QROM.

A beneficial feature of the ROM is "lazy sampling," where outputs are generated as they are first accessed by the oracle queries. This property becomes intricate in a quantum context due to the potential for superpositioned queries. Zhandry demonstrated that lazy sampling can be adapted using compressed oracles, albeit in a more limited form [248]. This adaptation allows for efficient simulation of random oracles within the quantum computing model. Additionally, techniques like parallel query handling and accommodating ranges beyond $\{0,1\}^n$ via Fourier transforms are feasible [67]. Moreover, it is possible to implement non-uniformly distributed functions in this framework [76].

The programmability of the random oracle is a powerful tool in classical security proofs within ROM. Reprogramming – replacing the original random function with another that differs on specific elements – is more complex in the QROM but achievable through various methods. The one-way to hiding theorem allows reprogramming at the cost of some security [230]. Other approaches, such as Measure-Rewind-Measure and Measure-Rewind-Extract [151] [105], provide tighter security bounds. Reprogram-After-Measure addresses reduction loss by enabling OW-CPA secure PKEs to transform into KEMs with IND-1CCA security without security loss in QROM [59]. Memory-tight reductions for digital signatures are explored in [238], and it is demonstrated that classical output from quantum QROM implies extractable preimage, which facilitates "rewinding" in ROM proofs [89].

In non-uniform models, the utility of quantum advice versus classical advice depends on the specific security game being considered. For many natural security games, quantum and classical advice offer comparable advantages or disadvantages. However, it is theoretically possible to construct scenarios where quantum advice provides an exponential advantage over classical advice [162].

2.2 Transformations

2.2.1 Universal Composability

Universal composability (UC) offers a theoretical framework that ensures the secure composition of cryptographic protocols. The UC framework's robustness is highlighted by its quantum lifting theorem, which asserts that any classical

protocol achieving statistical classical-UC-security inherently satisfies statistical quantum-UC-security as well [229].

Despite these strengths, notable distinctions exist between classical and quantum settings within the UC framework. A significant point is that while bit commitment alone cannot achieve general statistically secure multi-party computation (MPC) in the classical realm, the introduction of a quantum channel allows for such protocols to be constructed securely within the quantum environment [229]. PQC protocols need to be analyzed through the lens of a quantum-aware UC framework.

2.2.2 Fujisaka-Okamoto

The Fujisaka-Okamoto (FO) transform is widely utilized in post-quantum Key Encapsulation Mechanisms (KEMs). Since traditional proofs of security for FO are history-dependent, it requires a different approach in the quantum model. Its adaptation to QROM demands additional assumptions on the basis IND-CPA secure scheme [118]. Research has shown that while both the Fujisaka-Okamoto and OAEP transforms maintain security in the QROM, the FO transform is advantageous due to its fewer necessary assumptions [222]. Nonetheless, these studies typically do not address the issue of decryption failures within the QROM context. A comprehensive post-quantum security proof for the textbook FO transform is available in [89].

Decryption failures present significant challenges, which can be addressed by employing either explicit or implicit rejection strategies. Explicit rejection involves returning a "reject" message when decryption fails, offering clear communication of the failure without leaking information. This method is preferred as it avoids susceptibility to timing side-channel attacks and eliminates the need for Grover search-based mitigations [123]. Implicit rejection, on the other hand, returns a random key upon failure, which can obscure the success or failure of decryption but introduces additional security considerations.

Shan et al. introduced "oracle-masked schemes" and a corresponding plaintext extraction procedure, providing a more precise IND-qCCA (see Section 2.4) security proof for the unmodified FO transformation by addressing quantum measurement-related disturbances [215]. This work marks the first successful establishment of IND-qCCA security proofs for both REACT and TCH transformations in the QROM.

In general, the security and universality of FO are preserved when transitioning from classical to quantum settings: the FO transform effectively converts any quantum-resistant IND-CPA public key encryption scheme into an IND-CCA2 key encapsulation mechanism. However, it's important to note that in the context of QROM, explicit rejection is notably superior to implicit rejection due to its ability to prevent information leakage without compromising security.

2.2.3 Fiat-Shamir

The Fiat-Shamir transform is a cornerstone for constructing post-quantum digital signatures. Research by Ambainis et al. highlights that, when considered relative to a random oracle, classically secure sigma protocols and the Fiat-Shamir heuristic may be insecure in quantum settings [12]. Specifically, this vulnerability is evident in computational sigma protocols with unique responses. In a quantum context, an adversary can create a superposition of all possible initial messages and measure the oracle's response. Subsequently, the adversary identifies valid responses corresponding to the measured challenge which is a strategy that classical rewinding techniques cannot prevent. This insecurity arises because the idealized behavior of random oracles underpinning such proofs may not accurately reflect real-world hash functions.

Don et al. have demonstrated that security properties like soundness and proof-of-knowledge are maintained in the Fiat-Shamir transformation within the QROM [88], provided that quantum computationally unique responses are required. This necessitates that the underlying sigma-protocol be a proof-of-knowledge against quantum adversaries, which is conjectured for some schemes. However, there is an $O(q^2)$ security loss for sigma-protocols, where q is the number of queries. This quadratic security reduction is optimal and stems from Grover's algorithm [87]. This illustrates a fundamental limit: quantum adversaries can achieve at most a quadratic advantage through superposition queries to the random oracle in sigma protocols.

Liu and Zhandry provided additional proofs supporting some assumptions made by Don et al. [163]. Interestingly, this security reduction does not seem to apply to commit-and-open identification schemes [54].

The Common Reference String (CRS) model extends naturally into the quantum domain through the introduction of quantum entanglement among parties, resulting in the Common Reference Quantum State (CRQS) model. In this enhanced model, it has been established that there is a fully-black-box impossibility result for applying the Fiat-Shamir transformation in scenarios where quantum resources are shared [91].

The Fiat-Shamir transform with aborts is an adaptation specifically relevant to PQC signatures such as ML-DSA. In this modified paradigm, the presence of an abort condition introduces additional security considerations. When operating in a quantum setting, the abort mechanism can lead to an extra layer of complexity. This results in an additional loss in security beyond the standard $O(q^2)$ incurred by Fiat-Shamir itself due to the potential for quantum adversaries to exploit the abort condition [20, 96]. The combination of these factors necessitates careful analysis and potentially new techniques to mitigate such vulnerabilities in post-quantum cryptographic protocols.

When implementing a construction, it is crucial that the concrete hash function used aligns with QROM assumptions. This alignment ensures that any theoretical security guarantees provided by the QROM model are applicable to real-world implementations [142].

2.2.4 Fischlin

The Fischlin transform provides a mechanism for generating non-interactive zero-knowledge proofs with extractability properties that hold in classical random oracle models. When adapted to the post-quantum setting, these properties can be preserved through techniques such as compressed oracles [169].

2.3 One-way functions

The concept of one-way functions (OWFs) is fundamental in cryptography, serving as a building block for various cryptographic protocols. Classically, an OWF is easy to compute but hard to invert without additional information. However, in the quantum realm, certain classical OWFs may not retain their hardness due to quantum properties. Yamakawa et al. have explored one-way functions that are secure against classical adversaries but vulnerable to quantum ones [242]. The quantum analogue of a one-way function is not yet completely understood which highlights a significant gap in our understanding and utilization of OWFs within post-quantum cryptography.

To address these challenges, researchers have proposed weaker cryptographic primitives specifically designed for the quantum setting:

- **One-way State Generators (OWSGs)**: These are functions that generate a sequence of states from an initial state. The difficulty lies not in reversing the function but in predicting future states without knowing the entire history, even with quantum capabilities.
- **EFI pairs**: These pairs generalize the concept of OWFs by providing mechanisms to extract information under certain conditions. They are a pair of efficiently samplable quantum states that are expansion-based, statistically far apart and computationally indistinguishable. Such pairs are designed to maintain security properties when transitioning from classical to quantum settings, particularly in scenarios involving zero-knowledge proofs and digital signatures.

OWSGs imply EFI pairs [143], suggesting a hierarchical relationship where OWSGs serve as a foundational element that supports the structure of EFI pairs. Conversely, while EFI pairs do not directly imply OWSGs [25], they are constructed using principles that may be inspired by or related to those found in OWSG frameworks. The impossibility of constructing black-box commitments from OWSGs alone highlights a fundamental limitation in extending classical cryptographic concepts into the quantum realm without additional constructs or assumptions [41].

2.4 Indistinguishability

In post-quantum cryptography (PQC), indistinguishability is analyzed under two distinct models: Q1, where adversaries make classical queries with quantum

computational capabilities, and Q2, where adversaries can also perform quantum queries. The security implications vary significantly between these models. The Q2 model introduces complexity as it lacks a clear hierarchy among different notions of indistinguishability, necessitating comprehensive analysis against several non-comparable definitions [132].

Indistinguishability under chosen plaintext attack (IND-CPA) is a widely used, weakest definition of security for encryption. In the quantum computing model, there are various generalizations of this model:

- **IND-1CPA**: This model applies to adversaries with quantum computational power but who can only make classical queries. Corresponds to Q1.
- **IND-qCPA**: Here, the encryption algorithm can be queried in superposition by a quantum adversary, adding a layer of complexity to potential attacks. Corresponds to a limited version of Q2.
- **qIND-qCPA**: This represents a further extension where both the challenge ciphertext and queries can exist in a quantum superposition, offering a framework for security against advanced quantum threats.

Indistinguishability under chosen ciphertext attack (IND-CCA) is the standard security notion for encryption. Similarly to CPA, there are different versions of CCA regarding the quantum capabilities of the adversary:

- **IND-CCA**: The standard security notion for encryption where the adversary can make both encryption and decryption queries, with restrictions on decrypting challenge ciphertexts.
- **IND-CCA2**: An extended version of the classical IND-CCA where adaptive queries can be performed.
- **IND-qCCA2**: An adaptation for quantum settings, allowing quantum queries during learning phases but restricting challenge queries to the classical domain. This model can give false security assurances as schemes meeting IND-qCCA2 might still be vulnerable if adversaries are able to use superpositions in challenge queries [38].
- **qIND-qCCA2**: A fully quantum framework allowing quantum challenge queries and ensuring indistinguishability even with superposed plaintexts. This model uses Zhandry's compressed oracle technique to address the quantum recording barrier [248].

A significant theoretical challenge in PQC is ensuring the indistinguishability of symmetric encryption under the Q2 model, where adversaries possess superposition access to message states. It has been demonstrated that "quasi-length-preserving" encryption schemes that minimally extend message length (such as standard block ciphers like AES and stream ciphers) are inherently

unable to achieve robust quantum indistinguishability (qIND-qCPA) or quantum semantic security [103]. This limitation arises because classical encryption methods cannot obscure the quantum phase of message states without extending the ciphertext length to incorporate additional randomness. While these schemes may maintain security under weaker, more restrictive models like IND-qCPA, where challenges remain classical, they fail in scenarios where adversaries can interact with an encryption oracle using quantum superpositions. Consequently, achieving high-level quantum semantic security necessitates moving beyond traditional length-preserving architectures towards constructions that involve message expansion, such as those utilizing quantum-secure pseudorandom permutations (qPRPs) (see Section 3.2).

2.5 Unforgeability

In PQC, existential unforgeability under chosen message attack (EUF-CMA) is adapted to account for the capabilities of quantum adversaries. In particular, the PO (plus-one unforgeability) model introduced by Boneh and Zhandry is a significant development in this area [38]. It defines security for Message Authentication Codes (MACs) where a quantum adversary with q queries to a MAC oracle cannot produce $q + 1$ valid, distinct classical message-tag pairs. EUF-qCMA model extends PO by allowing the adversary to make superposition queries to the oracle.

3 Symmetric cryptography

3.1 Pseudorandom Generators

Pseudorandom generators (PRGs) serve as essential components for constructing symmetric cryptographic primitives. PRGs can be constructed from one-way functions in the classical model [124]. Furthermore, this foundational result is shown in a blackbox manner and implies that if the underlying one-way function is secure against quantum adversaries, then the resulting pseudorandom generator will also be secure in a quantum context [249].

PRGs can also be constructed from random oracles. A classical PRG that remains secure under polynomially many queries in the ROM is assured of retaining its security properties even when faced with quantum adversaries. This principle, encapsulated by the Katz-Sela Lifting Theorem, guarantees that if a PRG is proven secure in the classical ROM and adheres to certain conditions, specifically, making only classical queries to the random oracle, achieving unconditional ROM security, and being pseudo-deterministic if it were to make quantum queries, it will maintain its security against quantum adversaries accessing the ROM via quantum queries [139].

Interestingly, standard black-box constructions of pseudorandom generators from one-way functions may remain secure in the QROM even when the underlying one-way function is quantumly insecure [242].

3.2 Pseudorandom Functions and Permutations

Pseudorandom functions (PRF) are designed to be computationally indistinguishable from a truly random function. The security of PRFs against quantum adversaries is defined in two models: Q1, where adversaries make classical queries but perform quantum computations, and the stronger Q2 model, where adversaries can also submit quantum queries. While the Q1 model reflects a realistic post-quantum scenario for current cryptographic deployments, the Q2 model provides a more rigorous framework by accounting for full quantum query capabilities.

Zhandry's work demonstrates that many classical PRF constructions retain their security in the face of quantum adversaries [249]. This includes well-known constructions such as the GGM construction derived from a length-doubling PRG, synthetic methods based on PRGs, and direct constructions grounded in one-way functions like those related to the learning with errors problem.

However, it is theoretically established that there exists a separation between classical security models and their quantum counterparts. Specifically, assuming the existence of classically secure PRFs, we can identify certain PRFs that maintain their security against adversaries restricted to classical queries but fail when such adversaries are allowed to make quantum queries [249]. Furthermore, within the Q1 model, permutation-based PRFs exhibit a security reduction: an adversary can recover keys using $O(2^{n/3})$ classical queries and $O(2^{n/3})$ quantum time, utilizing only $O(n^2)$ qubits [220].

In the Q2 model, the definition of a PRF is not straightforward. There exists different quantum analogues for PRFs such as pseudorandom function-like state generators (PRFSGs) [117], strong pseudorandom unitaries in the quantum Haar random oracle model and pseudorandom isometries [15]. These analogues appear to be non-equivalent [112]. Typically, security is considered for a "direct" generalization into the quantum model using quantum-secure pseudorandom functions (qPRF). Here, the qPRF can be queried in superposition and needs to be quantum indistinguishable from a random function.

Similar to PRFs, classical security proofs for pseudorandom permutations (PRPs) can collapse when adversaries make quantum superposition queries. Quantum secure PRPs can still be constructed from quantum-secure one-way functions, but require different techniques compared to their classical counterparts [250]. Specifically, format-preserving encryption inherently maintains full-domain security even in the presence of quantum attacks, allowing for rigorous reductions from quantum-resistant one-way functions to quantum-secure PRPs. Unlike PRFs, the compressed oracle technique does not apply straightforwardly to permutations, necessitating alternative approaches for constructing secure quantum PRPs [248].

The Even-Mansour construction, a well-known approach in classical cryptography, have been shown to maintain its security properties under quantum attacks within the Q1 model [2].

The Sum of Even-Mansour (SoEM21) construction faces a significant reduction in its security bounds when subjected to quantum analysis. A preprint

indicates that the classical birthday bound of $O(2^{n/2})$ deteriorates to a tight $\Theta(2^{n/3})$ under specific quantum attacks [221]. This reduction is due to offline Simon-type key-recovery attacks exploiting quantum queries to public permutations while maintaining classical interaction with the keyed oracle, effectively bypassing traditional Grover search constraints. In the Q2 model, SoEM21 lacks provable security due to polynomial-time period-finding attacks. In contrast, within the Q1 model, the Tweakable Even-Mansour cipher upholds a query complexity lower bound of $\Omega(2^{k/3})$, where k is the key length, and this bound remains independent of the number of keys [216]. For the FX construction, the bound decreases to $\Omega(2^{(k+n-u)/3})$ when considering 2^u keys.

The Keyed Sum of Permutations (KSoP) is a simplified cryptographic construction that uses two independent random permutations to build a pseudorandom function (PRF). It serves as an alternative to more complex constructions like SoEM21. Research shows that KSoP and its variants achieve an $O(2^{n/3})$ security bound in the Q1 model, indicating robustness against quantum attacks [78].

Two-round key-alternating ciphers (KAC) have demonstrated resilience to quantum attacks in both Q1 and Q2 models against non-adaptive adversaries. This confirms their resistance to polynomial-time attacks like Simon's algorithm, which compromises the 1-round Even-Mansour variant [18]. However, extending these security proofs to t-round KAC (where $t \geq 3$) in either Q1 or Q2 models remains unproven. Additionally, their security against adaptive quantum adversaries is yet to be established.

A three-round Key-Alternating Feistel network, even when employing the same random oracle for all rounds, functions as a PRP in the Q1 model [22]. The scenario where identical permutations are used across all rounds remains unresolved. Establishing security in the Q1 model involves addressing the complex interplay between adaptive classical queries and quantum queries to the primitives. This includes effectively managing collision events during random oracle reprogramming, a challenge not adequately addressed by existing proof frameworks. The Feistel construction with seven rounds has been proven to be a strong quantum PRP in the Q1 model [51].

3.3 Hash Functions

Hash functions are one the most widely used symmetric primitives and generally considered to be quantum secure. However, primitives proven secure in the classical ROM can be efficiently broken by quantum adversaries in the QROM using only polynomial queries, which is also true for some collision-resistant hash functions [242]. In fact, it has been shown that even the quantum version of collision-resistance is insufficient in the quantum model due to the difficulty in rewinding [12, 231]. Instead, an additional property called *collapsing* is also required [231].

The classic Merkle-Damgård extension has been shown to retains collision-resistance in the QROM [115]. However, since classical techniques such as lazy sampling and adaptive programming fail in the QROM, direct translation of

classical security proofs to the quantum setting is not possible. The preservation of all seven standard security properties (collision resistance, preimage and second-preimage resistance variants, and collapsing) against quantum adversaries need to be proven separately for each construction [115]. The ROX constructions is shown to retain all seven properties in the QROM.

Ambainis proved that the collision problem requires $\Omega(N^{1/3})$ quantum queries and element distinctness requires $\Omega(N^{2/3})$ even when the output range $M = N$ is small, not just for large range [11]. This means the collision resistance lower bounds hold for hash functions where the output length equals the input length.

The BHT algorithm can theoretically find collisions for any hash function using $O(2^{n/3})$ quantum queries, where n is the hash length [44]. However, implementing this algorithm in practice is currently hindered by the requirement to store exponentially many hashes in quantum memory, which is beyond current (and possibly even future) technological capabilities.

Carolan's work provides tight quantum lower bounds for preimage and collision resistance for both Sponge (SHA3) and Davies-Meyer (SHA2) constructions within the random permutation oracle model. These results underscore the resilience of these hash functions against quantum attacks [51].

The post-quantum security bounds in the Q1 model for keyed sponge constructions (including Ascon AEAD and KMAC) are theoretically capped at approximately $2^{k/3}$, where k is the key length, due to their reduction to the Even-Mansour and Tweakable Even-Mansour primitives [121]. These results do not address Q2 security or indifferentiability.

3.3.1 Indifferentiability

Standard proof techniques for indifferentiability are not adequate in the quantum model due to quantum no-cloning constraints and measurement-induced state disturbance. The compressed oracle technique is needed to show that Merkle-Damgård hash domain extender is quantum indifferentiable [248]. However, a recent preprint shows that when all inputs to a random oracle are public, which happens for example in Fiat-Shamir and full-domain-hash signatures, then Merkle-Damgård is provably not post-quantum public indifferentiable [120]. This shows a separation between indifferentiability in the classical setting and its post-quantum analogue.

The sponge construction, foundational to SHA-3 and all recent NIST post-quantum public-key standards is quantum indifferentiable from a random oracle [5]. The security bounds are better than for Merkle-Damgård which inherently forces looser security bounds.

Standard and reset indifferentiability frameworks are theoretically insufficient to guarantee security against adversaries equipped with pre-computation (classical or quantum advice) [52]. This limitation arises because such attacks can exploit specific structural weaknesses that are not addressed by the indifferentiability paradigm alone. Consequently, existing composable security paradigms cannot rigorously bound the efficiency of inverting these hash functions under quantum or classical advice and current frameworks leave a fun-

damental gap in verifying the post-quantum security of several widely used schemes. However, one-round sponge is indifferentiable (with pre-computation) from a random oracle [52].

3.4 Message Authentication

Plus-one unforgeability (PO) serves as a foundational security measure in the quantum setting, adapting classical unforgeability principles while addressing challenges posed by quantum mechanics [38]. Despite its utility, PO is considered relatively basic compared to stronger notions like blind unforgeability or quantum pseudorandom functions (qPRFs), which offer enhanced protection against sophisticated quantum attacks [157].

Message Authentication Codes (MACs) are often designed based on hash functions. One popular construction is Hash-based MAC (HMAC). Hosoyamada and Iwata have shown that both HMAC and NMAC have a distinguishing query bound of $\Theta(2^{n/3})$ in the QROM [122]. This means that HMAC-SHA-256 provides approximately 85 bits of quantum security. The security bound degrades for very large message lengths.

It has been shown in a preprint that the Learning parity with noise (LPN) based authentication protocols HB+ and HB# are fundamentally insecure in the Q2 setting. In this model, the independently sampled noise terms collapse into unobservable global phases, effectively neutralizing the noise and enabling adversaries to recover secret keys efficiently using the Bernstein-Vazirani algorithm [68].

3.5 Symmetric encryption

The primary quantum threat to symmetric encryption is Grover's algorithm, which can significantly speed up brute-force key search. A straightforward and effective countermeasure against this threat is to double the key size. For instance, while AES with 128-bit keys would require 2^{64} operations for a brute-force attack using Grover's algorithm, doubling the key length to 256 bits increases the required effort to 2^{128} operations. This effectively mitigates the quantum speed-up.

However, increasing the key size can have performance implications. Doubling the key size typically results in more computational resources being used, potentially slowing down encryption and decryption processes by a factor of four [14, 110, 188, 172]. Despite this, AES with 192-bit or 256-bit keys is widely regarded as secure against quantum attacks [172, 188].

AES remains a cornerstone in designing quantum-resistant systems due to its efficiency and ability to encrypt large volumes of data effectively. The NIST standard FIPS-203 [181], which outlines the Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM), specifies that even when using ML-KEM for key exchange, AES should be used for symmetric encryption. This highlights the continued reliance on AES for its robustness and efficiency in a post-quantum world.

3.5.1 IND-CPA

While classical IND-CPA security is well-understood, translating this notion directly to the quantum setting reveals significant challenges. In the quantum realm, traditional concepts do not map neatly due to the unique properties of quantum mechanics [38]. Specifically, there are 14 distinct equivalence classes of quantum qIND-qCPA notions, denoted by qIND-qCPA-P1 to qIND-qCPA-P14 [53]. These classifications reflect varying levels of security and assumptions in a quantum context.

Among these, qIND-qCPA-P14 represents the weakest form, aligning closely with IND-qCPA or the Q1 model. This indicates that while classical IND-CPA proofs provide some insight into quantum settings, they do not suffice to ensure security against adversaries with superposition-access capabilities [13]. Thus, a fundamental theoretical gap exists where standard PRF security does not extend to scenarios involving quantum queries.

Modes like CBC and CFB are vulnerable in the Q1 setting when based on standard PRFs. Quantum adversaries can exploit periodicity through Simon's algorithm to achieve full key recovery [13]. This vulnerability is addressed by enhancing the block cipher to a qPRF. In contrast, XTS mode remains susceptible: quantum superposition queries combined with partial plaintext control enable an adversary to recover uncontrolled plaintext portions, undermining confidentiality assumptions. However, using a qPRF secures modes such as CBC, CFB, OFB, and CTR against adversaries in the Q1 (qIND-qCPA-P14) model [13].

However, CBC, CFB, OFB, and CTR modes of operation are fundamentally insecure under the qIND-qCPA-P13 security notion regardless of the underlying block cipher [182]. Common cryptographic design practices actively compromise quantum security: schemes that preserve input length fail qIND-qCPA-P8 security, and those that use public functions (such as XOR) with adversary-accessible randomness for randomization fail qIND-qCPA-P5 security. There is a gap in the current quantum security framework, as existing qIND-qCPA notions lack direct equivalents to classical semantic security.

IND-CPA is also challenging for AEAD schemes due to the presence of AD. The difficulties arise because AEAD schemes must preserve the associated data register, which invalidates standard query model assumptions and enables trivial distinguishability attacks via quantum entanglement. The primary effect is that existing formal security frameworks cannot reliably evaluate or prove the post-quantum confidentiality of AEAD schemes without explicit circuit-level modeling, potentially leaving implementations vulnerable to quantum period-finding attacks (e.g., Simon's algorithm) and invalidating standard security reductions [251].

3.5.2 IND-CCA

Similar to IND-CPA, classical adaptive chosen-ciphertext security definitions are unachievable in the quantum setting due to the "recording barrier," which arises from quantum no-cloning and the destructiveness of measurements, prevent-

ing the storage and comparison of quantum challenge ciphertexts [61]. Widely deployed classical constructions including stream ciphers, block cipher modes (CTR, CFB, OFB), authenticated encryption modes like GCM, and traditional hybrid encryption schemes become insecure when quantum challenge queries are allowed, necessitating new approaches.

In the fully quantum setting (Q2), PO is theoretically insufficient to guarantee IND-qCCA security for encrypt-then-MAC generic composition, disproving prior assumptions that classical security thresholds extend to quantum adversaries [157]. This means that many widely deployed classical MACs (e.g., CBC-MAC variants, PMAC, and polynomial hashing-based constructions) are insecure in Q2, as they are vulnerable to quantum period-finding and linearization. Consequently, achieving post-quantum authenticated encryption via generic composition requires a qPRF for the underlying MAC. In addition, GCM and GCM-SIV become vulnerable to quantum period-finding and linearization attacks due to their reliance on universal hashing, breaking IND-qCCA security [100]. Security can only be restored by replacing universal hash functions with pairwise-independent variants and enforcing nonce-dependent keys, but this protection strictly applies to nonce-respecting adversaries. Finally, nonce reuse completely compromises IND-qCCA security, even for modified schemes.

Combined Feedback (COFB) mode succumbs to a distinguishing attack with a single quantum query in the fully quantum model [252]. It is IND-qCPA secure, but the theoretical bounds degrade significantly as message length increases. Without larger block sizes or stricter per-key limits on queries and message lengths, security guarantees rapidly weaken. Furthermore, nonce length critically impacts quantum resilience; shorter nonces substantially degrade bounds, indicating that full-block nonces are necessary to maintain robust post-quantum authenticated encryption.

In Q2, sponge-based authenticated encryption schemes, such as SLAE, fail to achieve ciphertext indistinguishability, as their underlying one-time pad/stream cipher components exhibit inherent quantum insecurities that yield high distinguishing advantages [132]. The landscape of quantum unforgeability for MACs and AEAD schemes remains theoretically fragmented. However, SLAE is secure in Q1.

4 Asymmetric Cryptography

Asymmetric cryptography faces significant challenges due to Shor's algorithm, which can efficiently break widely-used public-key cryptosystems like RSA and ECC. The NIST PQC competition has been instrumental in researching and standardizing post-quantum cryptographic algorithms designed to withstand quantum attacks. Despite the progress made, there are still gaps in both the security considerations and practical implementation of these new standards. One key issue is ensuring that the chosen algorithms provide adequate security margins against potential future quantum capabilities. Additionally, integrating PQC into existing infrastructure poses challenges, including performance

overheads and compatibility issues with current systems.

4.1 Key exchange and encapsulation

Modern key exchange protocols have historically relied on the Diffie-Hellman (DH) algorithm, valued for its security properties, efficiency, and manageable key sizes. However, with quantum computing posing a significant threat to DH-based systems, post-quantum alternatives are essential.

NIST has produced standards for post-quantum secure key encapsulation mechanisms (KEM) to replace Diffie-Hellman-based key exchange mechanisms. Functionally, KEMs can be used as a replacement to DH due to their similar output. However, theoretical limits arise from slightly different security assumptions. While finite-field and elliptic-curve Diffie-Hellman key exchange is secure even in the case where public key is reused with the cost being lost forward secrecy, this property has been challenging to implement in KEM alternatives like ML-KEM [46], and it is recommended to avoid static keys in those. Another positive characteristic for some applications in Diffie-Hellman is the mutual agreement on the secret, where both parties participate in the secret generation process, whereas in KEMs the randomness that is used to choose the secret is provided by the "encapsulator" if no additional mechanisms are added.

The primary practical limit in current PQC implementations is communication overhead driven by large public keys, ciphertexts, and signatures, which consistently dominates handshake latency. For PQC KEMs there is a trade-off between computational and communication overhead: PQC algorithms exhibit widely varying computational costs. Key and ciphertext expansion increases bandwidth demands and transmission energy, negatively affecting battery life and latency especially in constrained environments [185]. Certain PQC candidates (code-based, and multivariate schemes) impose prohibitive computational or memory overheads, rendering them unsuitable for lightweight IoT and mobile platforms, while fast structured-lattice schemes remain the only viable option for ephemeral key exchange in mobile TLS.

Standard post-quantum KEMs restrict key reuse exclusively to the decapsulator, lacking the bidirectional contributiveness of Diffie-Hellman. This limitation forces additional message flows in protocol designs, breaking asynchronicity and sometimes increasing handshake rounds [46]. Passively secure lattice-based KEMs naturally align with the required split-key structure but are vulnerable to known key-reuse attacks, whereas strongly secure variants typically employ the Fujisaki–Okamoto transform, which discloses the encapsulator's randomness and inherently prevents key reuse. Consequently, current PQC solutions do not support static-static key exchange, reversed message flows, or flexible ephemeral-static combinations [46].

Even though FO transformation provides theoretical guarantees for CCA security of KEMs in the quantum model, there are problems related to side channel resistance. The re-encryption and equality-checking phases introduce exploitable power/EM leakage during pseudorandom function/generator execution, which enables the construction of a plaintext-checking oracle without

implementation-specific knowledge or external profiling devices [204]. This oracle facilitates adaptive chosen-ciphertext attacks that achieve full key recovery across most KEMs (ML-KEM, FrodoKEM, NTRU, NTRU Prime, HQC). Standard countermeasures, including constant-time execution and software masking, prove insufficient against this generic attack, but it can be mitigated through hardware solutions.

4.1.1 Lattice-based schemes

The Module-Lattice-Based Key-Encapsulation Mechanism (originally known as Kyber and now as ML-KEM) was the first PQC KEM standardized by NIST. It is typically the prime candidate for replacing classical DH-based key exchange systems with robust quantum-resistant alternatives. ML-KEM benefits from a formally verified proof of correctness and IND-CCA security, grounded in the module-LWE assumption [8], and is also shown to provably offer anonymity guarantees [170].

Lattice-based schemes are among the most efficient of the PQC primitives. However, they are considerable less efficient than classical methods primarily due to key and ciphertext sizes compared to classical systems, which restricts applicability in resource-constrained environments and latency-sensitive applications. FrodoKEM demonstrates up to 20 times slower key generation and substantially higher encryption/decryption latency than optimized schemes like ML-KEM [232]. Scalability poses a challenge, as performance degradation occurs with increasing message and key sizes, challenging large-scale deployments. ML-KEM exhibits scalability limitations especially under increased client loads, manifesting as latency spikes and elevated energy consumption [119]. Furthermore, computationally intensive keypair generation in some lattice KEMs (e.g., NTRU) limits their feasibility for forward-secure protocols [185].

Lattice-based algorithms are vulnerable to implementation-specific side-channel and memory-hardness attacks on conventional hardware. ML-KEM, FrodoKEM, and SABER employ message decoding functions with ciphertext malleability, which enables targeted side-channel attacks that circumvent generic countermeasures such as masking and shuffling [202]. Some countermeasures exist, such as those suggested by Berthet [31].

PQC primitives are challenging to use on constrained devices. As noted in a recent preprint, communication bandwidth presents a significant challenge, as ML-KEM and ML-DSA keys, ciphertexts, and signatures exceed typical constrained IoT link payloads (e.g., LoRa and BLE), necessitating application-layer fragmentation [62]. There is also an excessive memory footprint and an inherent trade-off between execution speed and stack usage. Memory-optimized configurations can also introduce severe cycle-count penalties and execution latency. The mutual exclusivity of speed and memory optimizations restricts deployment. For FN-DSA, the lack of hardware acceleration or constant-time floating-point support exacerbates performance bottlenecks. Lattice-based schemes also exhibit pronounced efficiency degradation at higher security levels on ARM architectures compared to x86 systems [228].

4.1.2 Code-based schemes

Hamming Quaci-Cyclic (HQC) [102] is the second post-quantum KEM chosen to be standardized by NIST as a back-up for ML-KEM. HQC has a proof of IND-CCA2 security based on two coding theoretic assumption, both variants of syndrome decoding of quaci-cyclic codes.

Code-based schemes suffer from excessively large public keys (e.g., 1 MB for Classic McEliece) and high verification data overhead, creating substantial computational and communication burdens. HQC currently lack practical ciphertext compression techniques, resulting in ciphertext sizes up to 5.8 times larger than lattice-based counterparts, which critically impacts bandwidth-constrained cryptographic applications [33]. Theoretical analysis establishes a fundamental lower bound for lossy compression in the Hamming metric, demonstrating that even with optimal quantization and rejection mechanisms, ciphertext reduction is strictly limited to approximately 5.9% without parameter adjustments. Achieving greater compression necessitates a direct efficiency trade-off with public-key size, such as requiring a 15% key expansion to attain a 10% ciphertext reduction.

Memory constraints on resource-constrained hardware restrict the deployment of code-based KEMs. In fact, severe computational and transmission overhead has been observed especially for HQC rendering it unsuitable for time-sensitive applications on constrained devices [119]. It has been observed that handshake latency is highly sensitive to network degradation, with delays exceeding 150 ms, packet loss above 2–10%, and bandwidth below 2–5 Mbps causing disproportionate performance drops, particularly for data-intensive schemes like HQC [116].

4.1.3 Hybrid methods for key exchange and encapsulation

KEM combiners provide CCA security as long as one of the KEMs used in them provides it [106]. These combiners make it possible to use multiple cryptographic algorithms in parallel in a way that the security is depended on the security of all of the algorithms used. This makes it relatively safe to adapt new algorithms that are not yet tested in time. In practice this allows us to use new quantum safe algorithms together with classical mature ones to get security guarantees of both. As long as one of the input algorithms is considered FIPS-approved by NIST, the entire key exchange remains FIPS-approved.

However, combining the methods securely is not easy. Cryptographic combinators serve as the core security primitive, yet most lack rigorous practical evaluation. Simple XOR-based combiners are proven insecure against CCA, and more robust constructions like XOR-then-MAC, dual PRFs, and cascade KDFs are needed [107]. Furthermore, compression trade-offs for conveying large post-quantum public keys alongside classical counterparts arise and can cause packet fragmentation or excessive handshake latency in constrained network protocols.

4.2 Public-key Encryption

The algorithms currently standardized by NIST or chosen to be standardized do not provide capabilities to perform public key encryption (PKE) directly, in the model established by RSA in the past. Instead, KEM algorithms can be used to establish a recipient-specific symmetric key which can in turn be used to encrypt information. In case of multiple recipients and a large data object, an intermediate symmetric key is possible to apply such as used in Cryptographic Message Syntax for email-like scenarios: N symmetric keys, one per recipient, are established with for example ML-KEM, which in the case of email needs a static key, and those N symmetric keys in turn are used one by one to encrypt one common data encryption key, resulting in N small ciphertexts, and one larger ciphertext of the actual data blob.

Theoretically, fully quantum computation model introduces complications to PKE. Quantum PKE with classical public and secret keys alongside classical ciphertext cannot achieve perfect completeness (i.e. if all involved parties follow the protocol honestly, they are guaranteed to successfully compute the identical shared secret key) if key generation relies solely on classical queries [158]. For schemes utilizing quantum public keys, security necessitates that the keys be mixed and unclonable. Additionally, QPKE is unachievable in any oracle model if the decryption algorithm makes no oracle queries, and non-interactive quantum key agreement is strictly impossible. This means that standard classical key formats and non-interactive protocols are fundamentally incompatible with secure quantum PKE from one-way functions, while quantum public keys inherently sacrifice reusability, which is a core requirement for scalable cryptographic infrastructure. Furthermore, secure quantum PKE cannot be constructed using short classical secret keys in oracle models, necessitating either longer keys, quantum-key generation, or relaxed completeness guarantees.

The qIND-qCPA quantum indistinguishability notion for public-key encryption requires an in-place encryption operator that, under the standard assumption that the challenger only possesses the public key, cannot be efficiently constructed for non-recoverable schemes [148]. For example, OAEP construction is non-recoverable under the partial-domain one-wayness of its underlying trapdoor permutation. This limitation renders the qIND-qCPA notion inapplicable to OAEP in standard communication settings meaning that current quantum security definitions need to be revised to achieve universal applicability across real-world public-key encryption primitives.

Regarding CCA security, classical plaintext-awareness (PA2) combined with quantum IND-CPA security does not imply quantum indistinguishability under chosen-ciphertext attacks (IND-qCCA) [92]. This means that traditional classical security notions are theoretically insufficient for proving strong post-quantum chosen-ciphertext security, creating a breakdown in standard security reduction frameworks when adversaries possess superposition oracle access. In the standard model, constructing qCCA-secure public-key encryption (PKE) is challenging without generic transformations and the quantum random oracle model (QROM). However, Alamati and Maram have shown that qCCA can

be achieved through key-dependent message secure encryption or hash proof systems, and a single-bit qCCA-secure PKE is sufficient to realize a multi-bit qCCA-secure PKE [7].

4.3 Digital signatures

The quantum computation model necessitates a revision regarding unforgeability of signature schemes. Standard unforgeability notions fail to account for quantum adversaries leveraging superposition queries to signing oracles, leading to a potential security gap. In the Post-Quantum Existential Unforgeability (PQ-EUF) model, the adversary is capable of performing quantum computations but cannot query the challenger's circuit in superposition. This model is akin to the Q1 model used for symmetric encryption, providing a framework to evaluate signature schemes under post-quantum threats.

Similar to MACs, this model may not be sufficient, since quantum adversaries may try to attack the concerned protocol/primitive by executing it on a quantum computer, even if the protocol/primitive by design is only meant to be executed classically [38]. The plus-one unforgeability (PO) model can be applied instead, but can remain insufficient against attacks involving disjoint message subsets or superposition answers [6]. Alagic et al. propose blind-unforgeability (BU) as a more generic model, where quantum superposition attacks are defeated by modifying the signing oracle to always reject messages in a randomly chosen blinded subset, thereby forcing any successful adversary to forge a signature for a message they have never received a valid signature for. Lamport's and Winternitz one-time signatures provide blind-unforgeable security in the QROM.

4.3.1 Hash-and-Sign

Standard lifting theorems and existing QROM proof techniques cannot handle hash-and-sign schemes that sequentially query two dependent random oracles [58]. The challenge for signature schemes is that correlated inputs/outputs prevent the straightforward application of tight adaptive reprogramming and measure-and-reprogram techniques. The effect is that establishing QROM security for such schemes requires novel proof frameworks that support successive reprogramming of correlated oracles as is the case for the PSS-R scheme [58].

The hash-and-sign with retry (HSwR) paradigm does not have a proof of security for plus-one or blind unforgeability models. However, derandomized HSwR with bounded retries achieves both PO and BU security in the QROM, though the proofs necessitate statistical rather than computational preimage-simulatability due to the inability to perform adaptive reprogramming in quantum signing oracle simulations [145].

4.3.2 Fiat-Shamir signatures

Classical security proofs for Fiat-Shamir signature schemes do not account for quantum superposition access to signing oracles [247]. Under superposition attacks, standard honest-verifier zero-knowledge (HVZK) is theoretically insufficient to preserve blind unforgeability in the QROM, as quantum queries can produce entangled transcripts that leak information beyond classical bounds. This limitation necessitates either strengthening the underlying identification scheme to quantum special HVZK or modifying the construction to be deterministic or variants that require a qPRF.

4.3.3 Lattice-based signatures

As of writing this report, NIST has standardized two lattice-based digital signatures schemes: ML-DSA and FN-DSA. FN-DSA stands out as the fastest in terms of signature verification, while ML-KEM is faster in key generation and signing. Both incur an increase in both public key and signature sizes, which can be up to 50 times larger than in their classical counterparts. Compared to ECDSA, lattice-based schemes exhibit substantial key and signature bloat, creating storage and bandwidth constraints that often necessitate protocol or hardware redesign, and present an expanded physical attack surface: nearly all intermediate values are sensitive to side-channel and fault injection attacks, requiring complex countermeasure combinations [84]. FN-DSA's implementation is notably more difficult to secure against side channels, as due to optimizations it involves calculations on secret input that use floating point numbers instead of integers.

Both schemes face platform-specific bottlenecks; for instance, FN-DSA's verification performance degrades without double-precision floating-point units, making ML-DSA more practical for signing on constrained devices [185]. In addition, FN-DSA has been identified to suffer from excessively slow key-pair generation, with complexity comparable to SLH-DSA, and presents inherent trade-offs where higher security levels incur increased computational complexity and longer signature sizes [195]. The signing operation imposes severe latency and power penalties that cannot be easily mitigated in software [212].

There is also a significant certificate and signature overhead, where PQ X.509 objects are 1.5 to 6 times larger than classical equivalents, inflating handshake bytes and straining the limited flash and memory of constrained edge devices [214]. Finally, the absence of mature CA/PKI support for PQC certificates and the engineering overhead of hybrid/dual authentication schemes currently hinder seamless transition and complicate protocol design, with the greatest impact falling on mobile, cloud-connected, and IoT applications where energy efficiency, low latency, and limited bandwidth are critical. A recent preprint notes that, computationally, ML-DSA signing exhibits severe latency variance due to rejection sampling which poses challenges for time-critical applications [62]. Memory requirements are also highly restrictive; memory demands of ML-DSA signing functions limits deployment on Class-1 IoT devices without

explicit stack management.

One notable difference in the method of operation between the standardized post-quantum digital signature algorithms compared to RSA and ECDSA is in their inputs. While it has been possible to hash a message beforehand and sign that with RSA or ECDSA, with ML-DSA and SLH-DSA, it is no longer possible to feed in a message hash into a signing algorithm and produce the same output as signing the original message. This is because the original message must instead be hashed together with established context from the public part of the signing key to protect against certain attacks.

This has also led to pushback, as implementations of a signing infrastructure where very large messages must be moved to the secure elements used for signing are not always viable. With ML-DSA, an additional method of operation called external-mu has been introduced in the standard, which allows hashing the original message together with the public key outside the boundary of, for example, a hardware security module. Some have argued that making external-mu available in the signing interface may lead to attack vectors for finding out some information about the key by feeding in false hash values, but no exploitability has so far been discovered. Using a trusted environment for calculating the external-mu input should be considered good practice [98], as the source of the external-mu could also obviously request signatures for arbitrary messages in addition to providing arbitrary external-mu input that it does not have a preimage for.

ML-DSA was also standardized with a variant HashML-DSA which is specifically designed for signing hash inputs rather than original messages. However, $\mathsf{HashML\text{-}DSA.sign}(hash(M))$ cannot be used to produce a signature that would validate as $\mathsf{ML\text{-}DSA.sign}(M)$, and it is therefore yet another signature algorithm, and one that is not separately mentioned in many applying standards or government recommendations. It should also be noted that if the message to be signed is itself a hash of an original file, $M = hash(M')$, the resulting signature $\mathsf{ML\text{-}DSA.sign}(M)$ is still a valid signature, just comparable to what we would have expressed as $\mathsf{RSA.sign}(hash(hash(M')))$ for RSA. It has been argued that having separate hashing variants leads to unnecessary ambiguity that makes protocols more complicated to analyze [209].

4.3.4 Hash-based signatures

SLH-DSA is largely considered a conservative backup algorithm to the lattice-based ones. Cryptographically, SLH-DSA relies primarily on the security of the chosen hash function, either SHA2 or SHAKE. Since cryptographic hash function security is already generally required for all signing use cases in order to be able to represent arbitrary-length messages as smaller digests, SLH-DSA has the benefit of not having pressure to be hybridized in certain regions.

SLH-DSA is primarily characterized small keys, but computational inefficiency and excessive signature overhead compared to lattice-based alternatives. Hash-based signatures produce extremely large outputs, increasing bandwidth and storage demands, while stateful variants require strict state management

to prevent security degradation. Severe computational and transmission overhead has been observed for SLH-DSA rendering it unsuitable for time-sensitive applications on constrained devices [119]. These constraints severely limit SLH-DSA's suitability for high-throughput and resource-constrained cryptographic applications, where rapid signing and minimal bandwidth are essential cryptographic goals. Architecture-aware optimizations are required to fully utilize of the algorithm's inherent parallelism [255]. NIST is preparing to standardize new additional parameter sets for SLH-DSA that would be suited to produce smaller keys and signatures for use cases that do not need to sign millions of times with the same key. This would suit, for example, offline root keys that are only used to sign a handful of other keys.

The SHA2 variant of the originally submitted SPHINCS+ algorithm proposal had a flaw due to SHA2 not fulfilling the requirements set for the used hash function within the algorithm. This flaw was corrected during the NIST standardization competition process. However, as a result, the implementation of the SHA2 variant of SLH-DSA is clearly not as elegant as the SHAKE variant. While the performance of the SHAKE variant may be slower when measured on older hardware, SHAKE outperforms SHA2 on modern hardware that includes SHA3 acceleration.

There is also a predictability benefit in latency that SLH-DSA has over ML-DSA. The benefit becomes more relevant in use cases such as telecommunications where predictable latency is important. ML-DSA (FIPS 204) signing uses rejection sampling and therefore has variable execution time. While the average-case latency is low, the signing time follows a geometric-like distribution with high variance and a long tail, and the worst-case latency is infinite. While the SLH-DSA (FIPS 205) average-case latency is much higher than ML-DSA, the signing time follows a Irwin-Hall-like distribution with very low variance tightly concentrated around the mean, and with worst-case latency twice the average-case latency. When performance is evaluated using high-percentile tail latency, SLH-DSA's almost constant signing time compares much more favourably with ML-DSA's variable-time, rejection-sampling-based approach [206].

The study by Kempf et al. provides an in-depth analysis of SLH-DSA's impact on Time To First Byte (TTFB) during a TLS 1.3 handshake, compared with RSA and ML-DSA [141]. Conducted using software implementations on an x86 server CPU, the study's findings are broadly applicable to TLS 1.3 and DTLS 1.3 due to their shared handshake mechanism, as well as IKEv2, which is based on a similar design framework. In comparison with RSA-4096 (13.46 ms), ML-DSA-65 emerged significantly faster (2.14 ms). SLH-DSA exhibited slower performance than its counterparts with SLH-DSA-SHA2-128f (20.89 ms) and SLH-DSA-SHAKE-128f (41.89 ms). However, test was performed on hardware without native SHA3 support.

4.3.5 Hybrid methods

Similar to KEMs, hybrid signatures attempt to combine the security of both classical and PQC methods. Signatures can be hybridized at two levels: at

the algorithm level, where the hybrid is represented to the protocol as a single algorithm, and at the protocol level, where the protocol is aware of handling multiple signatures.

Unlike key exchange hybrids, signature hybrids benefit greatly from application or protocol level awareness, such as having two signatures that by verification policy must come from two algorithm groups and trace back to two independent trust anchors. In an algorithm-level hybrid, in contrast, multiple trust anchors are merged together inseparably and signature certificates cannot be revoked independently. For two algorithms that can independently fail to sign or validate on the same input, protocol level awareness also allows for improved error handling.

Signature hybridization causes multiple challenges. A naive hybrid combination of two signatures can produced a scheme that is less secure than its components. For example, failing to provide integrity guarantees for the signatures themselves have lead to serious signature malleability attacks in the past [79]. Such attacks can undermine system integrity, auditability, and provenance and, in the worst case, may even facilitate replay attacks [193]. To protect against such threats, strong existential unforgeability (sEUF-CMA) is needed to ensure that an adversary cannot produce a novel valid signature for a message when they have previously obtained a valid signature for it [63]. Several widely deployed cryptographic libraries enforce canonical (low-s) ECDSA signatures to eliminate signature malleability. This should be best practice; however, it is not compliant with ECDSA as specified in FIPS 186-5 [180].

The same is true for certificates. There is a significant gap between what people think a certificate fingerprint represents and what cryptography actually guarantees when a certificate is signed with an algorithm that does not provide signature integrity. With such signature algorithms, a CA does not issue a single certificate; instead, it issues a set of valid certificates, each with its own fingerprint. This mismatch has practical consequences. Logging, SIEM, and threat intelligence systems have been seen to record events such as "Observed certificate fingerprint X connecting to service Y," implicitly treating the fingerprint as a stable identifier. Similarly, some firewall blocklists operate on fingerprints (e.g., “Block fingerprint X”) [177], and incident response workflows often rely on fingerprints as unique identifiers when searching for the attacker across datasets.

In the absence of signature integrity guarantees, these assumptions break down, as the same underlying certificate can appear under many fingerprints. That is, given a certificate (tbsCertificate, Signature), an attacker can derive one or more additional valid certificates (tbsCertificate, Signature') that have different fingerprints $Hash$(tbsCertificate, Signature'). Nation-state attackers are known to deliberately confuse logging systems, evade signature-based detection, and introduce ambiguity into forensic analysis, and they can be expected to exploit signature malleability for these purposes as well.

If we consider that composite signatures are in fact new cryptographic algorithms, and should be held to the same standards as standalone algorithms and composite KEMs, we should follow the principles of the European definition [71],

where a hybrid is defined as "a combination of a post-quantum algorithm and a quantum-vulnerable algorithm for the same mechanism, such that the security is as high as the stronger of the two components." A composite signature should never provide weaker security properties than its components.

Naive composite signature constructions by concatenation with ECDSA as applied in IETF draft for X.509 not only significantly reduce the security properties of ML-DSA by inheriting the malleability of ECDSA [192], but also introduces additional malleability weaknesses. In particular, when hedged or randomized signing is used, an attacker who has observed n valid signatures of a message M can derive up to $O(n^2)$ distinct new valid signatures for the same message. In practice, repeated signing of the same message M with the same private key k often occur when a signing request is retried after failure or interruption, as well as in high-availability systems where the same message is submitted to multiple HSMs. The composites proposed also significantly weaken the security of ML-DSA by failing to preserve its beyond-unforgeability properties [74, 219].

As shown by Jackson et al. [129], existential unforgeability alone does not capture the guarantees required by real-world protocols, and the lack of beyond unforgeability properties in traditional signature schemes has enabled concrete attacks on deployed systems [74, 219].

The security vulnerabilities identified in naive composite hybrids are not surprising. The new cryptographic algorithms were designed without input from Crypto Forum Research Group (CFRG), and the primary objective was not to improve security. Rather, the intention was to enable vendors to market these hybrids as "quantum-resistant" and FIPS-validated, even when only the quantum-vulnerable components (RSA or ECDSA) had been validated.

The concatenation hybrid is not the only option for hybridization, however. Signature combiners have been proposed that preserve both strong unforgeability and beyond-unforgeability properties [131].

It is also worth noting that if for example European ECCG ACM is interpreted strictly, a cryptographic mechanism is approved only if all of its underlying cryptographic primitives are approved. This makes all hybrid trust anchors under threat of being invalidated at a regulatory level when their non-PQC component is deprecated, which is particularly problematic if they were primarily introduced to satisfy governmental recommendations to begin with. Because roots of trust often remain deployed for decades, selecting mechanisms with uncertain long-term approval status introduces operational risk and may necessitate costly and disruptive replacement of trust anchors.

5 Advanced Cryptography

5.1 Oblivious Transfer

An oblivious transfer (OT) protocol defines a two-party exchange between a sender and a receiver, where the sender holds a set messages, and the receiver

obtains a subset of these messages according to its choice, without revealing the selection to the sender and without learning any information about the remaining messages [240]. OT protocols are commonly categorized into 1-out-of-2 OT, 1-out-of-n OT, and k-out-of-n OT variants.

As quantum computing advances, the design of post-quantum OT protocols has become an active research area. Li et al. propose an efficient 1-out-n OT protocol based on the ring learning with errors (RLWE) assumption [160]. Libert et al. present an adaptive OT protocol based on short integer solution (SIS) and learning with errors (LWE) [161]. Mi et al. construct a lattice-based 1-out-of-n OT protocol using the NTRU cryptographic primitive [174].

There is also a quantum version of OT, Quantum Oblivious Transfer (QOT). In both classical and quantum models, one-sided two-party computation is impossible with information-theoretic security [173]. Perfect QOT is only possible with restrictions on dishonest parties [243]. Santos et al. propose a hybrid quantum oblivious transfer (HQOT) protocol [208], which relies on oblivious keys generated via quantum processes.

5.2 Homomorphic Encryption

The idea of homomorphic encryption is to enable calculations on ciphertexts without first decrypting them. In partially homomorphic encryption, only one operation, typically either addition or multiplication, is doable in encrypted form. In fully homomorphic encryption, it is possible to carry out combinations of two different operations, typically both addition and multiplication, in encrypted form.

Bhoi, Arakala, Corman, and Rao have reviewed current lattice-based Post-Quantum Homomorphic Encryption schemes and provided justification for diversifying the field with code-based alternatives [32].

5.3 Password Authenticated Key Exchange

Password authenticated key exchange (PAKE) protocols enable participants sharing a low-entropy password to establish an authenticated and encrypted session without revealing the password itself. PAKE protocols are designed to prevent offline dictionary attacks by requiring an adversary to interact with a legitimate participant for each password guess. As a result, dictionary attacks are limited to online attacks, which must be mitigated by system-level defenses.

According to an overview from Alnahawi et. al. [9], PAKE standards have not yet been upgraded into the PQC settings, even though many PQC-solutions have been proposed. Among the surveyed post-quantum PAKE constructions, lattice-based solutions were highly represented, while code-based alternatives were not found. Since the code-based alternative HQC was selected by NIST as a standardized KEM, it would be natural to seek code-based alternative for PAKE as well. A promising methodology for constructing PAKE protocols is to use generic compilers that transform a KEM into a PAKE, but such constructions require additional security properties from the underlying KEM [24]. The

standardized lattice-based ML-KEM can be transformed to satisfy these properties (as described, for instance, in the IETF internet draft [234]), but whether KEMs based on other assumptions, such as HQC, satisfy or can be adapted to satisfy the same properties remains unclear.

As with KEMs and digital signatures, hybrid constructions are also desirable for PAKE protocols. Vos et. al. proposed such a construction based on classical CPace and their variant of post-quantum OQUAKE [235].

Jana and Dutta proposed two efficient PAKE protocols that, according to the authors, are the first post-quantum secure PAKE's that achieve identity privacy and resistance to pre-computation attacks [130].

5.4 Multi-Party Key Exchange

The problem of sharing a key with $n \geq 3$ parties is significantly more challenging than the two-party case. The concept of multi-party key exchange has been developed to securely establish a shared key among group members. Ingemarsson et al. proposed a solution using rings of integers [126]. Burmester et al. introduced a system with a cyclic function [48], assuming the authentication scheme is secure due to the discrete logarithm problem. Bresson et al. presented a group Diffie-Hellman key-exchange mechanism that achieves authentication key exchange with implicit authentication [47].

With the development of quantum computing, all of these systems become vulnerable. Dabra et al. propose a lattice-based password-authenticated key exchange (SL3PAKE) for two parties [77], based on existing 3PAKE, which involves a trusted third party.

Fan et al. provide multi-party key exchange protocols using an isogeny-based cryptographic scheme [97]. Li et al. propose a group key exchange protocol that uses symmetric key primitives and hash functions that are already quantum resistant [159].

5.5 Secure Multi-Party Computation

Secure multiparty computation (MPC) enables two or more parties, each holding private inputs, to jointly compute a function over their inputs without revealing them to each other [253]. The standard security requirements for MPC include privacy, correctness, independence of input, guarantee of output, and fairness.

Guizani et al. survey the impact of quantum computing on MPC security and assess emerging quantum-resistant MPC protocols. They also proposed a unified framework for modeling quantum-secure MPC [111].

Büscher et al. developed an oblivious transfer protocol based on the ring learning with errors assumption [49], which they use to implement a post-quantum variant of Yao's garbled circuits protocol with comparable efficient. While their construction improve resilience against quantum adversaries, the authors note that formal security proofs are not fully established.

5.6 Zero-knowledge proofs

Zero-knowledge proofs (ZKP) enable *provers* to convince *verifiers* of the validity of statements without disclosing any other information. Three properties need to be satisfied: completeness (if a statement is true, the prover is able to convince the verifier), soundness (malicious prover is not able to convince the verifier if the statement is false), and zero-knowledge (verifier learns nothing except that the statement is true). Originally, the zero-knowledge property needs to be satisfied unconditionally. However, there are variants, *argument systems*, that provide only computational security. In addition, classical ZKPs are interactive, but non-interactive variants have been designed for specific applications.

Quantum computing model places restrictions on zero-knowledge. There is a proven impossibility result that demonstrates that black-box expected quantum polynomial-time simulation is unachievable for constant-round protocols under measured-runtime models [64]. However, Lombardi et al. have shown the post-quantum security of several foundational systems such as the GMW graph non-isomorphism protocol, Feige-Shamir NP arguments, and Goldreich-Kahan constant-round proofs by redefining simulation complexity via coherent-runtime tracking to preserve quantum superpositions [166]. They argue that this model is the appropriate quantum analogue of classical expected polynomial-time simulation.

Chiesa, Manohar and Spooner lifted the original SNARG (succint non-interactive argument) construction of Micali from ROM to QROM [66]. The result applies to any SNARG obtained from public coin interactive oracle proof via BCS transformation. They also showed how QROM-secure SNARGs can be constructed based on certain properties.

The Merkle tree structure is used in various cryptographic commitment schemes that form a basis of zero-knowledge proof systems. Direct adaptations of classical Merkle constructions fail against quantum attacks, rendering them insecure in standard models and necessitating the quantum Haar ROM for viable commitments. However, the security of the quantum Merkle tree remains unproven against arbitrary malicious provers due to the lack of analytical tools tailored to this model [60].

Folding is a technique that can be used to build more efficient proof systems, as well as to construct incrementally verified computation. Pairing-based folding schemes exist [146], and are relatively straightforward. There is a literature on lattice-based folding schemes starting from Latticefold [36, 184, 152, 56] [35]. The main technical difficulty is efficiently proving the shortness of witnesses, which is required for the security of the underlying SIS-assumption. Techniques to help achieve this include random projections (also used in other lattice-based proof systems) and use of the sumcheck protocol. Lattice-based folding schemes focus largely on efficiency, and whether they are zero-knowledge has not been established.

5.7 Polynomial Commitment Schemes

Polynomial Commitment Schemes (PCSs) are a fundamental component of modern practical (zero-knowledge) Succinct Non-interactive Arguments of Knowledge (zk-SNARKs). They enable the compilation of information-theoretically secure Polynomial Interactive Oracle Proofs (PIOPs) into succinct arguments [65]. Consequently, the efficiency of the PCS largely determines the proof size, verification complexity, and computational overhead incurred by the prover. Many practically efficient systems, such as Spartan [213], HyperPlonk [57], and GKR [108], employ the SumCheck protocol as a subroutine and therefore require a PCS that supports multivariate polynomials [167].

The predominant paradigms for constructing PCSs are based on groups (particularly bilinear pairings) [137], variations of FRI [27], and lattices. Among these, pairing-based schemes typically provide the best practical efficiency and optimal asymptotic performance in the trusted-setup setting, with some constructions achieving constant-size proofs, linear prover complexity, and polylogarithmic verifier complexity [26]. A significant drawback of these pairing-based schemes is that they do not provide post-quantum security.

In contrast, hash-based and lattice-based schemes are plausibly post-quantum secure. However, despite being computationally efficient, hash-based schemes generally suffer from large proof sizes. For example, pairing-based schemes can achieve proof sizes of only a few hundred bytes [104], whereas hash-based schemes often require proofs larger than 100 KB [113], and sometimes on the order of megabytes. Lattice-based schemes can achieve smaller proof sizes, but then the verifier is typically not asymptotically or concretely efficient. Recent work by Klooß et al. [144] presents the first lattice-based polynomial commitment scheme, RoKoKo, that is both concretely and asymptotically efficient, achieving linear-time prover complexity and polylogarithmic-time verification, with a proof size of roughly 150 KB. RoKoKo has comparable concrete verifier complexity to many hash-based constructions, while offering significantly improved prover efficiency. Using lattices, it is also possible to obtain even smaller proof sizes (around 50 KB), at the cost of less efficient verification [183].

In conclusion, many alternatives for constructing post-quantum SNARKs still do not match the performance of pairing-based systems. From a security perspective, to the best of our knowledge, Fiat–Shamir-transformed lattice-based constructions have not yet been analyzed in the QROM and therefore lack a complete end-to-end post-quantum security analysis. In contrast, hash-based constructions are generally built using the BCS transformation [28], which has been shown to be secure in the QROM [66]. Consequently, hash-based constructions can be considered a more secure option, as they rely only on random oracles and their quantum security guarantees are better understood.

5.8 Anonymous credentials

Anonymous credentials (AC) allow a user to prove possession of a credential while revealing no additional information beyond the required attributes. Fur-

thermore, repeated presentations of the proof should be unlinkable. ACs are commonly constructed using zero-knowledge proofs, although blind-signature-based constructions are also possible, where an authority blindly signs a token representing the credential. Credential revocation is not as straightforward in AC systems as it is for regular digital certificates, since the verifier cannot simply check whether the credential appears in a credential revocation list. Instead, specialized revocation mechanisms have been proposed for some anonymous credential schemes.

As of yet, the efficient solution for post-quantum anonymous credentials remains an active research topic, and even the classical ACs lack standardization. BBS signatures are pairing-based signature schemes that support selective disclosure and zero-knowledge proofs of credential possession [225]. However, their security relies on assumptions vulnerable to quantum adversaries. Several post-quantum replacements have been proposed that replace the underlying pairing-based primitives with lattice-based constructions [237][135][153][39]. Although these constructions achieve post-quantum security, they significantly increase sizes of the credentials and the presentations.

Kemmoe et al. have proposed a method for credential revocation in their preprint paper [140], which is applicable to "state-of-the-art lattice-based anonymous credential systems" [80] [154] [134] [16].

Bouillaguet et. al. revisited Fischlin's framework for blind signatures and have presented a "efficient and compact blind signature schemes based on hash-and-sign post-quantum signatures" that outperform the most compact existing PQC blind signatures in terms of signature size [43].

6 Cryptographic Protocols

6.1 SSH

Secure Shell (SSH) is a protocol used to establish a secure client-server channels. It consists of three layers: the transport layer for encryption, server authentication, and message integrity [246], the authentication layer for user authentication [244], and the connection layer for multiplexing multiple logical channels over a single encrypted tunnel [245].

The connection layer operates within an encrypted and authenticated channel and does not introduce additional cryptographic mechanisms, and is therefore not relevant in post-quantum perspective. The authentication layer includes a public-key authentication method, where the server verifies the client's identity using a stored public key. Although this method requires a post-quantum solution, the transition is relatively straightforward due to the limited cryptographic complexity and is therefore not the main focus here. The transport layer relies on a more varied range of cryptographic primitives and is worth a closer examination, not only to understand the requirements for post-quantum transition of SSH, but also for comparison with similar protocols.

Public key distribution

The transport protocol in SSH, by default, uses a trust-on-first-use (TOFU) authentication model, where the client software locally stores the server's public key on the first connection after the user manually verifies the connection via an external source. Alternatively, SSH certificates can be used together with a certificate authority (CA), which is typically an local trust anchor within an organization. In the latter model, the server's public key is certified by the CA, which mitigates the risks associated with the TOFU model but introduces additional trust relationships.

Key establishment and authentication

The transport layer establishes a shared secret between the client and the server through a key exchange protocol. In the commonly used SSH key exchange methods, the server signs a hash of the exchange to bind its identity to the session. The SSH protocol also permits a key exchange methods with implicit authentication, where the key agreement alone suffices to establish the authenticity of the server. However, such methods are not implemented, for example, in OpenSSH.

Transmission

After deriving the shared keys, the transmission layer encrypts and authenticates all transmitted data using symmetric encryption and message authentication tags, or authenticated encryption with additional data (AEAD) schemes.

Post-quantum transition

Post-quantum migration in SSH primarily affects its asymmetric components used in key exchange and authentication, the latter covering both the server authentication in transport layer, and the client authentication in the authentication layer. OpenSSH already offers hybrid KEMs as a default for session key establishment, but the post-quantum digital signatures are not yet supported.

SSH's TOFU model simplifies small-scale migration to post-quantum authentication in one-to-one connections, compared to the larger-scale trust infrastructures associated with SSH certificate-based systems.

Results

Benčina et. al. performed a post-quantum security analysis of SSH in the quantum random oracle model (QROM), proving it to be a secure authenticated and confidential channel establishment (ACCE) protocol when the hybrid KEM used in the handshake phase is IND-CPA secure and the remaining primitives satisfy their respective security requirements[29]. This differs from the previous analysis, which relied on IND-CCA secure KEMs, such as those used in common post-quantum SSH implementations. The results suggest that

SSH does not require the stronger IND-CCA secure KEMs, and their experiments where they removed the Fujisaki-Okamoto transformation significantly reduced the number of hash function computations. The analysis focuses on the Harvest-Now-Decrypt-Later threat model and does not consider the active quantum threat.

The Open Quantum Safe (OQS) fork of OpenSSH enables a variety of post-quantum KEMs and digital signature algorithms for prototyping purposes. Using this fork, Crockett et al. successfully prototyped various NIST Round 2 submissions in OpenSSH, including both pure and hybrid Kyber and Dilithium [75]. Key establishment and authentication mechanisms were evaluated separately, and therefore a fully post-quantum OpenSSH was not deployed in the study. The experiments focused on a TOFU model with raw public keys and did not consider SSH certificate-based authentication model. The implementation effort was found to be relatively modest, which was partly attributed to one-to-one usage scenario.

6.2 TLS 1.3

Transport Layer Security (TLS) is a protocol to provide secure channel between two peers[205]. It consists of a handshake protocol which establishes the shared secret and authenticates the server, and optionally a client, and a record protocol that uses the established secret to protect subsequent traffic. The abstract framework of TLS is similar to the transport layer of SSH, with the most significant difference being the authentication model. Instead of the TOFU model, TLS relies on CAs to certify server public keys through a public key infrastructure (PKI).

Public key distribution

In TLS, the server will send its static public key to the client in a certificate signed by a CA, including any intermediate certificates. Similar to SSH, CAs and the public keys in their self-signed certificates act as trust anchors that bind identities to public keys. However, unlike the CAs typically used in SSH, TLS CAs typically act as global trust anchors that are also referred to as webPKI, public key infrastructure of the world wide web. Clients store trusted CA certificates locally to validate server certificates. Local CAs are possible for use cases that does not require global access.

Key establishment and authentication

A shared secret is negotiated in a handshake protocol, in which the parties first negotiate the used algorithms and exchange their ephemeral public keys to compute the key exchange. The server authenticates itself by signing the handshake transcript with the private key corresponding to the certified public key presented in the certificate. In case of mutual authentication, the client

similarly authenticates itself with its own certificate that is signed by a CA trusted by the server.

Transmission

The record layer protects the transmitted data with a symmetric AEAD scheme.

Post-quantum transition

The post-quantum TLS 1.3 is believed to be achieved by replacing classical asymmetric cryptographic methods in the handshake phase with quantum-secure counterparts similarly to SSH. In webPKI, authentication is based on a global public key infrastructure, which will take some time to include trust anchors for the new algorithms.

TLS does not rely on cryptographic properties beyond those required for ACCE protocols and forward secrecy, so in the classical setting the Diffie-Hellman with ephemeral keys has been quite straightforward solution for key agreement. Current post-quantum TLS deployments typically use hybrid key establishment combining Diffie-Hellman and ML-KEM. Properties of DH that are not present in ML-KEM (symmetrical key agreement, static-static key possibility) are not required or used in TLS, and therefore the changes in key agreement are not too significant in theory. Some practical issues arise from the large key sizes, however.

Transition to hybrid key exchange in TLS is already quite widely adopted, but quantum-secure authentication remains to be deployed. While key establishment mechanisms can be upgraded by the communicating endpoints, deployment of post-quantum authentication requires support throughout the PKI ecosystem. The tool chains necessary for the PKI transition, from secure CA key storage to software solutions to managing the certificates and updates to the relevant application standards are still being finalized [171], and the concern of whether ML-DSA should be hybridized, and if so, the maturity of hybrid options, is also impacting long-lived trust anchors in webPKI.

Results

Provable security has been demonstrated for the handshake protocol in TLS 1.3 against polynomial time adversaries [90]. In the chapter, the authors identify the need for the extension of their proofs to take into account the post-quantum extensions that were not implemented at the time, and the reliance on the Diffie-Hellman assumption.

Huguenin-Dumittan et al. have proven a CPA secure KEM to be sufficient in the TLS 1.3 handshake, but the security bounds were not tight, and the security was evaluated in ROM [125]. Biming et al. tightened the security bounds and lifted the proof into QROM [254].

Resource-constrained embedded devices face critical energy, latency, and memory limitations that severely restrict TLS deployment in battery-powered

IoT applications. Specific post-quantum combinations, particularly those employing SLH-DSA for digital signatures, incur execution delays and high energy consumption [223]. Resource-constrained end-nodes requiring mutual authentication or high quantum resistance face the most severe deployment bottlenecks, whereas devices operating in unilateral authentication configurations can sometimes achieve energy efficiency comparable to or exceeding traditional TLS.

TLS 1.3's limit in keyshare restriction poses challenges for large-key algorithms resulting in packet fragmentation. Increasing security levels or adopting hybrid implementations amplifies transmitted data, triggering TCP congestion window throttling and initial window delays that disproportionately extend handshake times [116]. However, ML-KEM and other structured lattice-based algorithms achieve very good overall performance and partially beat classical schemes.

6.3 Signal

Notation 6.1. Deniability is a security property that prevents Bob, the accuser, from using a transcript of communication with Alice, the denier, to prove to a third party, the judge, that Alice participated in communication. For 1-out-of-2 deniability, either party must be able to simulate a transcript indistinguishable from a real execution.

Signal is an end-to-end encrypted messaging protocol. Similar to SSH and TLS, Signal separates initialization and transmission into distinct protocols. However, communication is asynchronous, channels are longer-lived, and Signal provides additional security properties such as deniability. These properties significantly affect the cryptographic design and requirements for post-quantum transition.

Public key distribution

Users upload their public keys to Signal's server, which distributes them to other clients. For protection against a malicious server attempting a man-in-the-middle attack, clients can verify their public keys by comparing QR codes, either in person or through another secure channel.

Key establishment and authentication

Signal uses the PQXDH protocol [150], a post-quantum extension of X3DH, for key agreement and authentication. To send an initial message to Bob, Alice fetches a public prekey bundle containing his identity key, prekeys, and, when available, a one-time prekey, together with their respective signatures and identifiers. The initial message in a conversation consumes the recipient's one-time prekey to add additional entropy to initialization.

Alice verifies the signatures and generates an ephemeral key pair. To derive the shared key, she computes multiple keyshares from combinations of her

private keys and Bob's public keys. Each keyshare serves a distinctive purpose, binding the long-term identities for authentication and providing forward secrecy through short-term keys. The shared key is calculated using a key derivation function (KDF) from the concatenation of the keyshares.

The initial message Alice sends to Bob contains her identity key, ephemeral key, identifiers of Bob's used prekeys, and the first encrypted message, encrypted using the AEAD-scheme defined for the transmission layer. The initial message can be simulated by either of the participants, granting 1-out-of-2 deniability for both.

Transmission

Signal uses a Double Ratchet algorithm to exchange encrypted messages [196]. To achieve forward secrecy between messages, a symmetric-key ratchet is used to derive unique message keys from previous key material. To provide post-compromise security, an asymmetric ratchet introduces new ephemeral keys, effectively performing a new key agreement for each message chain.

Private Group System

Signal's private group system relies on pairwise one-to-one channels established between every pair of group members, secured with PQXDH and the Double Ratchet. To manage membership lists on the servers while preserving privacy, the lists must remain encrypted while still allowing authorized members to update the group state without revealing its contents to the server. The system for maintaining the membership lists is based on the construction proposed by Chase et. al [55]. To support these operations, the construction employs a homomorphic encryption scheme, allowing certain computations to be performed directly on encrypted data.

Post-quantum transition

Compared to TLS, the post-quantum transition for Signal requires more delicate solutions to preserve its security properties. Whereas in TLS, key agreement and the authentication are provided by separate cryptographic mechanisms during the handshake, in Signal these are combined into the same computations to support asynchronous operation, mutual authentication, and deniability, although the authenticity of the Bob's keys can be verified separately. The transition also differs in the message exchange phase, where the TLS transmission layer does not rely on asymmetric cryptography, whereas Signal's Double Ratchet does to provide post-compromise security. Additional requirements arise from Signal's private group system.

PQXDH is an extension of X3DH that introduces an ML-KEM-encapsulated shared secret, which is concatenated with keyshares for key derivation. Although this provides post-quantum confidentiality, the authentication still relies on classical assumptions. Even though signatures could be replaced with the

standardized post-quantum digital signatures, deniable mutual authentication in PQXDH relies on the symmetry of Diffie-Hellman key exchange, where Alice and Bob compute the same keyshare using their own private key and the other party's public key. As stated in the PQXDH documentation, "Post-quantum secure deniable mutual authentication is an open research problem" [150].

Since Diffie-Hellman is used in the Double Ratchet to provide post-compromise security, a ML-KEM-based ratchet has been added to construct a so-called Triple Ratchet [86].

Signal's group system is based on ElGamal for its homomorphic nature, and its security proof relies on the Diffie-Hellman assumption in ROM [55]. These imply that the confidentiality of encrypted group data, including membership lists, does not currently provide post-quantum security and is therefore vulnerable to harvest-now-decrypt-later attacks.

Results

Deniability of the classical X3DH and post-quantum PQXDH Signal protocols have been studied by Katsumata et. al. [138], and it was concluded that both provide deniability against a malicious classical accuser and a classical judge. In the model where we take into account the harvest-now-judge-later scenario, the accuser is classical, but the judge may be quantum-enabled. In this scenario, only the standard (honest-but-curious accuser) deniability holds. A paper by Fiedler and Langrehr challenges this result by arguing that the underlying EKDH assumption does not hold for deniability against a malicious accuser [99].

Brendel et al. have shown how to construct an asynchronous deniable authenticated key exchange protocol from a designated verifier signature scheme based on ring signatures [45]. However, it does not provide deniability against a malicious accuser who does not follow the protocol.

Protocols migrating to post-quantum algorithms should consider preparing for sudden changes in security assumptions by supporting backup algorithms. Typically, the lattice-based scheme is implemented first, followed by a code-based alternative. Juaneda et. al. propose a code-based ratcheting mechanism as an alternative to the ML-KEM mechanism in Triple Ratchet [136].

6.4 Tor

Tor is a distributed overlay network [226], which uses a circuit of nodes to conceal the relationship between the client and the destination. Communication between neighboring nodes is protected by TLS, but additional cryptographic layer is required to hide the content and the whole circuit from the individual nodes. A tor circuit includes the client (or Onion Service), entry node, one or more middle nodes, and the exit node (alternatively a rendezvous point or a introductory point). Entry node is the only node that knows the client's network address, while the exit node is the only node that knows the endpoint.

Onion Service is a server running inside the tor network, effectively hiding its real location. OS maintains a long-term introductory point through which the clients can initiate connections. In the initiation, client will request a connection to OS via a rendezvous point, which for they both build their own tor circuits.

Clients build circuits to exit node for anonymous access to public services, and to rendezvous point for anonymous access to onion services.

Public key distribution

Directory Authorities are the servers that maintain a consensus of the signed list of nodes, including necessary metadata such as IP-addresses and public keys. When the tor circuit it established, the list of available nodes is requested from the DA and appropriate nodes are selected at random, but weighted by bandwidth.

OS will publish the necessary details of its current introductory points in the tor network, and they are accessible for clients that possess the OS's onion address, which is a concentration of a OS's public key, version number, and the checksum. Onion address can be published for a selected audience, often via public web sites or chat groups.

Key agreement, authentication, and transmission

When the client is establishing the circuit, he first initiates ntor-v3 handshake with the entry node to form a private channel between the two based on a symmetric secret they share. By using the entry node as a proxy, client will extend the circuit by establishing similar channel with middle node without revealing the new symmetric secret for the entry node, or clients address for the middle one. Client will do this once more with the exit node to gain secure channel where the transmission, initially a TLS handshake to final endpoint, can be encrypted with three symmetric secrets, each decrypted by one node before relaying the transmission forward. Similar operations are conducted when the connection is established through the rendezvous point to a OS, but instead of regular TLS handshake with the endpoint, the client and OS perform a ntor handshake.

Post-quantum transition

Cryptography vulnerable to quantum algorithms is used in many parts of Tor protocol. TLS layer (dismissing the TLS 1.2), ntor handshakes, node authentication, signatures from DA's and OS's all rely on either Ed25519 or x25519. All of these imply practical sacrifices in terms of performance and memory requirements when migrated to PQC, especially for the multi-layered circuit establishment, but the post-quantum alternatives for these are not shown to result in theoretical issues regarding any subtle security properties.

Usage of OS's public key as a part of its onion address defines somewhat tight boundaries for the size of the public key, which suggests the usage of SLH-DSA if other changes to protocol are not being made.

Results

TLS part of the tor can be expected to follow the progression of PQ-TLS in general, and migration is on going with increasing number of nodes supporting the current implementations of it. Implementation of local PQ-Tor network by Berger et al. evaluates the practical overhead in performance and memory between classical and PQ variants of asymmetric operations such as the ntor-v3 [30]. Practical limitations arise from the limited space in handshake cells which leads to a fragmentation of a PQ keys, but the successful implementation suggests no significant challenges. They admit that their estimated PQC overhead is highly theoretical, since it is concluded in a constrained environment and misses a multitude of parameters. It's also noted that the plug-in replacement is probably not the most efficient way to handle this migration, and that there will probably be some reconstruction of the underlying protocols for efficiency gains. It should be noted that they used ML-DSA over the SLH-DSA for its efficiency without disclosing the problem with the onion addresses getting too large in practice.

6.5 European Digital Identity Wallet

The European Digital Identity Wallet (EUDIW) is a digital identity ecosystem that enables users to store, present, and manage their personal identification data (PID) and attestations using their wallet unit. To reduce repetition, we use the term attestation to refer to both PIDs and attestations. At the time of writing, EUDIW is specified by its Architecture and Reference Framework [95], and its reference implementation is being developed [94]. Among the actors in the EUDIW ecosystem, the most relevant for post-quantum considerations are identity providers (Providers), relying parties (RP), trusted lists of providers and RPs, and the wallet unit.

The cryptographic requirements, as summarized in the Cryptographers' Feedback on the EU Digital Identity ARF [23], include selective disclosure, unlinkability, pseudonymous authentication, unforgeability, and non-transferability.

Issuing attestations

Attestations are issued by providers following the OID4VCI protocol [164]. In this model, the wallet unit first authenticates to the provider and proves the user's eligibility to receive the attestation. The provider binds the attestation to a key that is controlled exclusively by the wallet unit to ensure non-transferability.

Attestation revocation

Attestation revocation is discussed in EUDIW ARF ANNEX 2 Topic 7. The providers shall use revocation by means of short-lived attestations, attestation status lists, or attestation revocation lists. These methods are not crypto-

graphic, but the cryptographic properties of the attestation affect the usability of these methods.

Presenting attestations

Presenting attestations to relying parties follows the OID4VP protocol [165]. The wallet unit receives a presentation request from the RP, typically encoded as a QR code, and parses the requested attributes. The wallet unit then authenticates the RP using the appropriate trusted list, searches for an attestation containing attributes that satisfy the request, and asks the user for consent to present the requested attributes.

The EUDIW ARF requires that presentations from the same user remain unlinkable for colluding malicious RPs, and even for colluding groups of RPs and providers.

Design challenges

Following the issues mentioned in ARF and by Carsten et al. [23], the proposed format of attestations and their presentation will either provide unlinkability or non-transferability, but not both. If attestations are strictly bound to a wallet unit, or its single key pair, then repeated presentations may be linkable by the relying parties. On the other hand, if the one-time key pairs are used and attestations are only used once, they become independent, weakening non-transferability since users can share credentials without revealing a long-term secret key. Additionally, maintaining a large number of keys would make the possibility of storing keys in secure enclaves infeasible. Short-lived attestations provide a partial compromise, but do not provide unlinkability against colluding groups of providers and relying parties.

A zero-knowledge proof-based verification mechanism is a recognized alternative in the ARF which would provide unlinkability even against the colluding providers and relying parties. Research on zero-knowledge proofs is active, but the lack of mature standards and hardware support complicates their integration.

Post-quantum transition

Carsten et al. note that the currently proposed authentication mechanisms not only contradict some of the classical privacy requirements, but also rely on classical cryptographic assumptions that are not post-quantum secure [23]. Among the anonymous credential systems they describe is the family of BBS schemes, which provide unconditional unlinkability, but the signatures remain forgeable for a quantum adversary.

Signature-based authentication could be migrated to a post-quantum setting by updating the signature schemes and presentation methods within the formats used with OID4VCI and OID4VP, but privacy-preserving post-quantum zero-knowledge proofs or suitable blind signatures would better satisfy the privacy requirements.

Results

Desmoulins et al. proposed anonymous credential scheme BBS# [83]. It is efficiently implementable on modern mobile devices, it provides credential unforgeability under classical assumptions, and it unconditionally provides strong unlinkability. Although it could be a significant improvement over the current authentication methods, a fully post-quantum construction based on it remains an open research question.

7 Limiting factors in hardware implementations of PQC

7.1 Effects of larger keys, signatures and ciphertexts

The dominant practical penalty of PQC in hardware implementations is the growth in the size of keys, ciphertexts and signatures, which translates directly into larger RAM, flash and register-file requirements in an imlementation [42]. When ECC keys occupy only a few tens of bytes, currently standardized PQC schemes require at least some thousands of bytes, some (e.g., code-based schemes) substantially more. These increased storage requirements may render it impractical to use PQC algorithms in certain extremely constrained environments.

The larger keys, signatures, and ciphertexts also set a burden for the interfaces between the cryptographic core and the surrounding system simply because more data needs to be transferred before and after the cryptographic operation. Having said this, PQC has not changed the fact that costs of hardware accelerated asymmetric cryptography operations are still dominated by computation rather than interfacing [19, 239, 101], but the interfacing cost is a factor that plays a larger role especially in tightly-coupled-accelerators that require frequent data transfers between the host and a cryptographic accelerator (such as a Keccak core) [101].

More bytes must be moved across the communication channel as well. In systems where the radio or contactless interface dominates the energy budget, such as smart cards, RFID/NFC tokens and low-power wireless embedded nodes, this transmission overhead can outweigh the cost of the cryptographic computation itself and may create problems with power budgets and reduce battery life. Also transaction latencies may increase beyond the limits tolerated by contactless standards. These aspects can be significant problems for the PQC-transition in extremely constained systems.

7.2 Microarchitecture

Classical public-key hardware is comparatively homogeneous: an RSA or ECC core is essentially an engine for modular arithmetic of large integers. PQC schemes, by contrast, are heterogeneous, combining several distinct computa-

tional kernels that each warrant their own hardware support [19]. A typical lattice implementation requires:

1. modular arithmetic engined with relatively small modulus for performing an NTT/inverse-NTT and various polynomial arithmetic operations,
2. a Keccak/SHAKE core for hashing and pseudorandom expansion,
3. samplers that convert pseudorandom bytes into coefficients drawn from a specific distribution, and
4. encoders and decoders for mapping byte strings to polynomial vectors, and vice versa.

These blocks have differing datapath widths, control flows and area–latency trade-offs, complicating both the microarchitecture and the scheduling to keep them busy in an optimal manner [19, 101]. Although these blocks alone can be simple and compact, the plain number of different operations that must be supported makes designing a compact PQC hardware implementation a difficult task and the footprints of compact PQC implementations tend to be significantly larger than of compact ECC implementations; this is emphasized if memory footprints are also considered, as discussed above.

7.3 Side channel analysis and PQC algorithms

Like all cryptographic hardware, PQC implementations must resist side-channel and fault attacks that are in the threat model for that particular use case [203, 197]. However, several structural features of PQC algorithms make this harder and more costly than for mature classical designs (in particular ECC) [40]. The difficulties arise partly from the heterogenous nature of PQC algorithms, which was discussed above, and partly from the relative newness of the schemes [203], which means that few established countermeasures are available and concensus about the level of required protections is still partly missing.

7.3.1 Current Results in Securing PQC Against Side-Channel Attacks

PQC schemes, while mathematically secure against quantum adversaries, remain vulnerable to *Side-Channel Attacks* (SCA) that exploit physical leakages such as timing, power consumption, and electromagnetic emissions [81, 85]. These attacks target the implementation rather than the underlying hard problem, making practical security a critical concern.

7.3.2 Observed Vulnerabilities

Recent studies have demonstrated that lattice-based schemes like CRYSTALS-Kyber and Dilithium exhibit non-trivial leakage under power and electromagnetic analysis, even when implemented with basic countermeasures [189]. Ma-

chine Learning (ML)-assisted SCAs further reduce the number of traces required for key recovery, highlighting the insufficiency of traditional protections [189]. Timing-based attacks exploiting rejection sampling and secret-dependent branches have also been shown effective against certain PQC primitives [207].

7.3.3 Countermeasure Strategies

Current countermeasures include:

- **Masking and Shuffling:** Applied to lattice-based schemes to randomize intermediate values and execution order [85].
- **Constant-Time Implementations:** Avoiding secret-dependent branches and memory accesses, particularly in rejection sampling routines [207].
- **Randomized Clocking and Delay Insertion:** Introduced to mitigate timing and power analysis attacks, though at significant performance cost [73].
- **AI-Driven Mitigation:** Emerging approaches leverage anomaly detection and adaptive countermeasures using machine learning to dynamically respond to leakage patterns [1].

7.3.4 Hardware-Level Protections

FPGA-based implementations of NIST PQC finalists have incorporated universal countermeasure components, such as integrated masking units and noise injection, to reduce leakage without excessive resource overhead [133]. Automated evaluation platforms like PQC-SEP now enable pre- and post-silicon leakage assessment, improving design-time security validation [224].

7.3.5 Timing attack on HQC

In a recent paper [155], it has been shown that HQC and other code-based KEMs are vulnerable to timing attacks. It is specifically mentioned that an attacker can exploit the timing leakages that occur when processing sparse vectors, which are ciphertext-independent, to recover the secret key by measuring the leakages only once. The authors state that CT-grind can find timing leakages in various implementations of HQC.

7.4 User interface differences for applying the cryptographic primitives

The structure of how to request signatures and other operations is different in the NIST standardized ML-KEM, ML-DSA and SLH-DSA algorithms. From a programmer's point of view, the Abstract Programming Interface (API) of the cryptographic primitives has changed in a way that has notable practical

implications that bear summarizing from the point of view of implementation considerations.

7.4.1 Key establishment constraints

In case of KEM algorithms such as ML-KEM, the implementation allows joint establishment of a key between two recipients. In online protocols such as TLS, an ephremeal ML-KEM key can be created and a separate algorithm is used to authenticate this one-time KEM key and the key exchange result. If ML-KEM is used with a static key, it is possible to exchange the key with a receiving party without the receiving party being online, for example for the purposes of email encryption.

The way KEMs work enforces three differences in ways of working:

- A KEM key cannot sign. A single RSA key for email signing and encryption has not been recommended for a long time, but with ML-KEM it becomes compulsory to have a separate key.
- A KEM key cannot encrypt a pre-generated key. For establishing a single key between multiple recipients, KEM must be used to establish a symmetric key with each recipient, then those keys can be used to encrypt a shared key.
- A KEM key proof of possession also cannot involve a direct signature operation. This has implications on KEM certificate issuance by a Certificate Authority.

As mentioned, unless static KEM certificates are used, the missing direct link to proof of possesion is not a problem. However, in situations where a certificate authority (CA) tries to create a certificate for a static KEM key, for example for email encryption, it needs proof of possession to be convinced the recipient requesting a certificate is in fact in possession of the private key. The CA can in this situation either a) create the key itself in which case it knows exactly where it came from, it can b) ask a separate signing key that it has already issued to be used to sign the certificate signing request (CSR), or it can c) try to establish some kind of custom online protocol to be convinced the requestor is able to decrypt information for which a key is established using the public key in the request.

7.4.2 Signing constraints

In the past, it has been possible to use RSA and ECDSA algorithms for signing a hash of a document independently of the document. In order to address certain attack models, both ML-DSA and SLH-DSA have been standardized in a way that binds the hash of the document into the signing key. By default, neither of them can be used to sign a hash of a document in a way that would allow validation directly against the original document.

For ML-DSA, a special construct called External-μ which also needs a public key as input can be used to produce a smaller representation of the document and sign that. This must separately be supported by the signature generator.

For SLH-DSA, the entire message to be signed must be passed to the module doing the signing.

This has a few practical implications:

- To sign a large file with ML-DSA, a hardware security module would need to support External-μ to allow passing it a smaller representative string.
- There are likely to be surprising limits to how large files SLH-DSA can sign with keys in a hardware security module in practice, as it has no equivalent workaround.

Practical experiments have demonstrated that in terms of Certificate Revocation Lists, a file containing only a few thousand revocations was already too much for the worst-performing hardware security modules [114].

While pre-hashed variants of the algorithms have been proposed, HashML-DSA and HashSLH-DSA remain distinctly different algorithms from the regular algorithm, so a file hash signed with HashML-DSA cannot be validated as a ML-DSA signature, in addition to re-introducing the security issues that motivated the interface change to begin with [209].

If the implementation of the signature is not following a specific standard such as CRLs are, this algorithm-level limitation does not stop signing a file that contains the hash of another file, and performing first a signature validation of that file, and a separate validation that the original message matches the hash in the signed file. This is essentially a protocol or application layer change, and is therefore unaffected by algorithm constraints.

8 Transition

8.1 Limits to Cryptographic Asset Visibility

Visibility into cryptographic assets is widely identified in practitioner guidance and emerging academic work as a foundational prerequisite for PQC migration [175, 200, 210, 198]. Comprehensive inventories are consistently described in these sources as the necessary starting point for understanding cryptographic exposure, which in turn enables coordinated transitions across complex systems. Taken together, these perspectives suggest that cryptographic visibility is not merely a supporting capability, but a core enabling condition for any systematic migration effort.

Despite this widely acknowledged importance, multiple sources indicate that organizations frequently lack comprehensive visibility into their cryptographic assets and dependencies [69, 127, 218, 198]. Survey evidence further suggests that only a minority of organizations report full visibility into cryptographic assets such as certificates, keys, and secrets [198]. Prior work attributes this limitation to the fact that cryptographic mechanisms are deeply embedded across

applications, infrastructure, and third-party systems, often in ways that are not immediately observable, even to system owners [69, 127]. As a result, inventories are reported to remain incomplete, inconsistent, or unreliable for migration planning [127, 198]. More broadly, these challenges reflect not only limitations in tooling, but also structural and operational factors, such as complex system dependencies, legacy infrastructure, and organizational constraints, that hinder the establishment of comprehensive and continuously maintained cryptographic visibility [233, 128, 198].

Recent standardization and industry efforts have focused on Cryptographic Bills of Materials (CBOMs) as a mechanism for improving cryptographic visibility. This focus reflects the diversity of analysis domains involved in cryptographic discovery, where different techniques capture different aspects of cryptographic usage [210]. While standardization initiatives such as CycloneDX have only recently emerged [147, 194], early academic and practitioner evidence suggests that currently available tooling remains immature and evolving [34]. Current CBOM generation approaches may produce divergent or incomplete results, often requiring multiple techniques to compensate for individual limitations. As such, existing evidence suggests that CBOMs are best understood as a promising but still developing component of broader cryptographic visibility strategies, rather than a complete solution.

8.2 Causes of Limited Cryptographic Agility

Cryptographic agility is generally understood as the capability of systems to replace and adapt cryptographic algorithms over time while preserving security and maintaining uninterrupted operation [21]. However, the concept remains insufficiently defined. A dedicated literature survey by Alnahawi et al. identifies the absence of a unified definition and highlights that the requirements enabling crypto-agility are not commonly agreed upon, while also noting a lack of clarity regarding how the concept can be realized in practice [10]. Similarly, a systematic literature review by Näther et al. reports inconsistencies in terminology, definitions, and best practices across the field, indicating that no clear consensus has yet emerged and that the overall state of the literature remains fragmented [178].

A recurring observation in standards and industry reports is that many existing systems were not originally designed to accommodate cryptographic change. In legacy and mission-critical environments, cryptographic functionality is often tightly embedded within system components and lacks clear abstraction from core application logic, reflecting historical design practices that did not anticipate future transitions [21]. As a result, cryptographic components are often difficult to update, and many systems cannot support the rapid adoption of new algorithms without substantial system changes [21, 186]. Evidence from industry contexts suggests that these limitations can be amplified in large-scale deployments, where cryptographic mechanisms are deeply embedded across system layers and require coordinated updates across multiple stakeholders and components. Such migration efforts are described as complex and non-trivial,

requiring careful coordination and introducing operational challenges [93].

In addition to tight coupling, the lack of widely adopted abstraction mechanisms and interface standardization contributes to fragmentation at the implementation level. NIST guidance highlights the importance of modularity and abstraction via application programming interfaces as key approaches to enabling cryptographic change, and identifies the development of common cryptographic APIs as an area requiring further work [21]. Complementing this, a qualitative interview-based study of cryptographic libraries indicates that design decisions can span a broad range of abstraction levels and that programming interfaces vary substantially, ranging from low-level to high-level abstractions and influenced by factors such as standards, legacy code, and developer practices [211]. This variability suggests a degree of heterogeneity in current implementations, which may hinder interoperability and increase the effort required to evolve systems consistently over time.

Efforts to introduce crypto-agility can themselves increase system complexity. Supporting multiple algorithms, configurations, and transition pathways requires additional abstraction layers, coordination mechanisms, and testing effort, as noted in NIST guidance [21]. A community workshop report on post-quantum cryptography migration further highlights concerns that agility mechanisms introduce additional complexity and may have unforeseen consequences, including implications for system fragility and attack surfaces [191]. More generally, increased system complexity can, if not properly controlled, lead to additional attack surface and failure modes, underscoring the importance of rigorous design, and systematic security evaluation when introducing crypto-agile mechanisms.

Taken together, these observations suggest that limited practical crypto-agility in current systems arises from three interrelated factors: the absence of a shared conceptual foundation, leading to inconsistent definitions and requirements; fragmentation at the implementation level due to heterogeneous abstraction mechanisms and interfaces; and increased system complexity, which complicates the secure design, integration, and evolution of crypto-agile systems.

8.3 Potential Execution Gap Between Policy Timelines and Organizational Practice

Regulatory guidance on PQC consistently defines policy milestones extending to 2035, but these timelines are normative in nature and do not claim to represent empirically validated execution schedules at the organizational or system level. In this context, *execution gaps* refer to potential misalignments between externally defined policy milestones and the practical ability of organizations to complete end-to-end cryptographic migration across heterogeneous systems and dependencies. This notion is used here as an analytical construct rather than an empirically established condition.

NIST IR-8547 defines concrete transition milestones, including the deprecation of public-key mechanisms providing approximately 112-bit security strength by 2030 and the disallowance of quantum-vulnerable public-key algorithms, such

as RSA and elliptic-curve cryptography, in NIST standards after 2035. These dates are explicitly framed as inputs to standards planning and risk management rather than as guarantees of system-level migration completion [176]. The European Commission and the UK National Cyber Security Centre similarly articulate staged targets through 2035 as policy or risk-management objectives, rather than strict guarantees of uniform implementation across organizations [72, 179].

Execution of PQC transition is consistently characterized in practitioner and industry guidance as challenging, reflecting a long-running, dependency-heavy transformation spanning systems, vendors, and organizational boundaries rather than a discrete technical upgrade [199, 70].

Only a limited number of sources provide explicit numerical estimates of transition duration. NIST IR 8547 does not specify explicit timelines but notes that previous ecosystem-scale cryptographic migrations required more than a decade and that the more complex PQC transition is expected to require at least comparable effort, providing a conservative lower-bound expectation [176]. Practice-oriented guidance and analyses, including those by QRAMM, TNO, and Campbell, suggest that full transitions may span multiple years to well over a decade depending on organizational complexity and scale [201, 227, 50].

When considered alongside fixed policy timelines extending to 2035, these estimates indicate a tightening of the effective transition window, suggesting that organizations which have not yet initiated PQC transition activities may face increasing difficulty in achieving comprehensive migration within current policy horizons.

In addition to timeline considerations, multiple sources identify economic and resource constraints as significant barriers to PQC adoption, with cost, budget limitations, and competing priorities contributing to delayed or phased migration strategies [187, 82, 109, 17]. At the same time, ongoing standardization efforts introduce additional uncertainty, as organizations must plan migration paths while technical standards and implementation guidance continue to evolve [3, 4].

At the ecosystem level, insufficient international coordination during the PQC transition may contribute to diverging implementation approaches and standardization practices, increasing the risk of fragmentation [17, 156]. It can be reasonably inferred that such fragmentation may, in turn, complicate interoperability across PQC-compliant systems and protocols.

Collectively, these observations indicate that aligning organizational execution with externally defined policy timelines is likely to be challenging without additional empirical validation or structural intervention. They highlight structural factors that may constrain the ability of organizations to complete comprehensive migration within current policy horizons.

8.4 Organizational Limits to Cryptographic Migration

Organizational factors are increasingly identified in both academic and practitioner literature as a critical constraint on cryptographic migration. Prior work on crypto-agility emphasizes that technical mechanisms alone are insufficient to

enable systematic transitions, as organizational conditions fundamentally shape how and whether cryptographic changes are implemented [168, 217, 149]. Collectively, these findings indicate that migration is not purely a technical challenge, but an organizational process requiring coordination, prioritization, and sustained operational support.

Despite the recognized importance of maintaining up-to-date cryptographic mechanisms, evidence indicates that such efforts are often deprioritized under operational pressures. In particular, everyday operational needs can take precedence over cryptographic agility considerations, resulting in updates being postponed or inconsistently addressed [168]. Empirical findings further show that cryptographic updates are frequently perceived as challenging and are therefore delayed or carried out in an ad hoc manner, with contributing factors including insufficient expertise, legacy system constraints, and a lack of structured processes [149]. These observations indicate that migration timelines are not determined solely by technical feasibility, but are significantly shaped by organizational constraints affecting planning and execution.

From an enterprise-level perspective, cryptographic migration is complicated by limited control over configuration and reliance on software providers. Many applications place cryptographic configuration in the hands of providers, leaving operators with little control and requiring coordination across numerous providers to implement changes [217].

Further compounding these challenges, empirical evidence indicates that structured processes for managing cryptographic updates are often lacking in practice [149]. Developers report limited guidance and insufficient cryptographic knowledge, while also facing constraints such as legacy system dependencies that hinder systematic updates [149].

Overall, existing evidence indicates that organizational constraints, manifested in competing operational priorities, limited control over cryptographic configuration, and the absence of structured processes and expertise, pose a significant barrier to systematic cryptographic migration. Taken together, these factors suggest that effective migration depends not only on technical mechanisms, but on organizational capabilities that enable coordinated planning, decision-making, and execution across systems.

Lack of Systematic Cryptographic Transition Engineering

Existing literature is consistent with our view of the current cryptographic transition landscape, indicating that research has primarily focused on algorithmic and protocol-level aspects, while system-level transition strategies remain comparatively underexplored [236]. At the same time, prior work shows that available guidance for migration is often high-level and lacks concrete implementation detail, leaving substantial aspects of the transition to practitioners [178, 190]. Taken together, these findings suggest that systematic engineering approaches for cryptographic transition, including concrete methods and design practices for achieving crypto-agility, are still insufficiently developed.

References

[1] Raymond Ajax, John Owen, and Fadare Dare, *Ai-driven side-channel attack mitigation in post-quantum cryptography*, ResearchGate Preprint (2025).

[2] Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz, *Post-quantum security of the even-mansour cipher*, Advances in Cryptology – EUROCRYPT 2022 (Cham) (Orr Dunkelman and Stefan Dziembowski, eds.), Springer International Publishing, 2022, pp. 458–487.

[3] Gorjan Alagic, Maxime Bros, Pierre Ciadoux, David Cooper, Quynh Dang, Thinh Hung Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Hamilton Silberg, Daniel Smith-Tone, and Noah Waller, *Status report on the fourth round of the nist post-quantum cryptography standardization process*, NIST Internal Report NIST IR 8545, National Institute of Standards and Technology, March 2025.

[4] Gorjan Alagic, Maxime Bros, Pierre Ciadoux, Quynh Dang, Thinh Hung Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Hamilton Silberg, Daniel Smith-Tone, and Noah Waller, *Status report on the second round of the additional digital signature schemes for the nist post-quantum cryptography standardization process*, NIST Internal Report NIST IR 8610, National Institute of Standards and Technology, May 2026.

[5] Gorjan Alagic, Joseph Carolan, Christian Majenz, and Saliha Tokat, *The sponge is quantum indifferentiable*, 2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS), 2025, pp. 2591–2630.

[6] Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song, *Quantum-access-secure message authentication via blind-unforgeability*, Advances in Cryptology – EUROCRYPT 2020 (Cham) (Anne Canteaut and Yuval Ishai, eds.), Springer International Publishing, 2020, pp. 788–817.

[7] Navid Alamati and Varun Maram, *Quantum cca-secure pke, revisited*, Public-Key Cryptography – PKC 2024 (Cham) (Qiang Tang and Vanessa Teague, eds.), Springer Nature Switzerland, 2024, pp. 193–226.

[8] José Bacelar Almeida, Santiago Arranz Olmos, Manuel Barbosa, Gilles Barthe, François Dupressoir, Benjamin Grégoire, Vincent Laporte, Jean-Christophe Léchenet, Cameron Low, Tiago Oliveira, Hugo Pacheco, Miguel Quaresma, Peter Schwabe, and Pierre-Yves Strub, *Formally verifying kyber*, Advances in Cryptology – CRYPTO 2024 (Cham) (Leonid Reyzin and Douglas Stebila, eds.), Springer Nature Switzerland, 2024, pp. 384–421.

[9] Nouri Alnahawi, David Haas, Erik Mauß, and Alexander Wiesmaier, *SoK: PQC PAKEs - design, security and performance*, Cryptology ePrint Archive, Paper 2025/119, 2025.

[10] Nouri Alnahawi, Nicolai Schmitt, Alexander Wiesmaier, Andreas Heinemann, and Tobias Graßmeyer, *On the state of crypto agility*, Tagungsband zum **18** (2022), 103–126, Accessed: 2026-01-12.

[11] Andris Ambainis, *Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range*, Theory of Computing **1** (2005), no. 3, 37–46.

[12] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh, *Quantum attacks on classical proof systems: The hardness of quantum rewinding*, 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, 2014, pp. 474–483.

[13] Mayuresh Vivekanand Anand, Ehsan Ebrahimi Targhi, Gelo Noel Tabia, and Dominique Unruh, *Post-quantum security of the cbc, cfb, ofb, ctr, and xts modes of operation*, Post-Quantum Cryptography (Cham) (Tsuyoshi Takagi, ed.), Springer International Publishing, 2016, pp. 44–63.

[14] Ravi Anand, Arpita Maitra, and Sourav Mukhopadhyay, *Grover on SIMON*, Quantum Information Processing **19** (2020), no. 9, 340.

[15] Prabhanjan Ananth, John Bostanci, Aditya Gulati, and Yao-Ting Lin, *Pseudorandom unitaries in the haar random oracle model*, Advances in Cryptology – CRYPTO 2025 (Cham) (Yael Tauman Kalai and Seny F. Kamara, eds.), Springer Nature Switzerland, 2025, pp. 301–333.

[16] Sven Argo, Tim Güneysu, Corentin Jeudy, Georg Land, Adeline Roux-Langlois, and Olivier Sanders, *Practical post-quantum signatures for privacy*, Cryptology ePrint Archive, Paper 2024/131, 2024.

[17] Abdullah Aydeger, Engin Zeydan, Awaneesh Kumar Yadav, Kasun T Hemachandra, and Madhusanka Liyanage, *Towards a quantum-resilient future: Strategies for transitioning to post-quantum cryptography*, 2024 15th International Conference on Network of the Future (NoF), IEEE, 2024, pp. 195–203.

[18] Chen Bai, Mehdi Esmaili, and Atul Mantri, *Quantum security analysis of the key-alternating ciphers*, Cryptology ePrint Archive, Paper 2025/945, 2025.

[19] Utsav Banerjee, Tenzin S. Ukyab, and Anantha P. Chandrakasan, *Sapphire: A Configurable Crypto-Processor for Post-Quantum Lattice-based Protocols*, IACR Trans. Cryptographic Hardware and Embedded Systems (TCHES) **2019** (2019), no. 4, 17–61.

[20] Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu, *Fixing and mechanizing the security proof of fiat-shamir with aborts and dilithium*, Advances in Cryptology – CRYPTO 2023 (Cham) (Helena Handschuh and Anna Lysyanskaya, eds.), Springer Nature Switzerland, 2023, pp. 358–389.

[21] Elaine Barker, Lily Chen, David Cooper, Dustin Moody, Andrew Regenscheid, Murugiah Souppaya, Bill Newhouse, Russ Housley, Sean Turner, William Barker, and Karen Kent, *Considerations for achieving crypto agility: Strategies and practices*, NIST Cybersecurity White Paper NIST CSWP 39, National Institute of Standards and Technology, 2025.

[22] Jyotirmoy Basak, Ritam Bhaumik, Amit Kumar Chauhan, Ravindra Jejurikar, Ashwin Jha, Anandarup Roy, André Schrottenloher, and Suprita Talnikar, *Post-quantum security of key-alternating feistel ciphers*, Advances in Cryptology – ASIACRYPT 2025 (Singapore) (Goichiro Hanaoka and Bo-Yin Yang, eds.), Springer Nature Singapore, 2026, pp. 446–478.

[23] Carsten Baum, Olivier Blazy, Jan Camenisch, Jaap-Henk Hoepman, Eysa Lee, Anja Lehmann, Anna Lysyanskaya, René Mayrhofer, Hart Montgomery, Ngoc Khanh Nguyen, Bart Preneel, abhi shelat, Daniel Slamanig, Stefano Tessaro, Søren Eller Thomsen, and Carmela Troncoso, *Cryptographers' feedback on the eu digital identity's arf*, https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/discussions/211, June 2024, available online at https://github.com/user-attachments/files/15904122/cryptographers-feedback.pdf.

[24] Hugo Beguinet, Céline Chevalier, David Pointcheval, Thomas Ricosset, and Mélissa Rossi, *Get a cake: Ge neric t ransformations from k ey encaspulation mechanisms to password a uthenticated k ey e xchanges*, International Conference on Applied Cryptography and Network Security, Springer, 2023, pp. 516–538.

[25] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa, *A new world in the depths of microcrypt: Separating owsgs and quantum money from qefid*, Advances in Cryptology – EUROCRYPT 2025 (Cham) (Serge Fehr and Pierre-Alain Fouque, eds.), Springer Nature Switzerland, 2025, pp. 23–52.

[26] Juraj Belohorec, Pavel Hubáček, Aleksi Kalsta, and Kristýna Mašková, *CHOPIN: Optimal pairing-based multilinear polynomial commitments from bivariate KZG*, Advances in Cryptology – CRYPTO 2026, 2026, To appear in CRYPTO 2026.

[27] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev, *Fast Reed-Solomon Interactive Oracle Proofs of Proximity*, 45th International

Colloquium on Automata, Languages, and Programming (ICALP 2018) (Dagstuhl, Germany) (Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, eds.), Leibniz International Proceedings in Informatics (LIPIcs), vol. 107, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018, pp. 14:1–14:17.

[28] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner, *Interactive oracle proofs*, Theory of Cryptography (Berlin, Heidelberg) (Martin Hirt and Adam Smith, eds.), Springer Berlin Heidelberg, 2016, pp. 31–60.

[29] Benjamin Benčina, Benjamin Dowling, Varun Maram, and Keita Xagawa, *Post-quantum cryptographic analysis of ssh*, 2025 IEEE Symposium on Security and Privacy (SP), 2025, pp. 595–613.

[30] Denis Berger, Mouad Lemoudden, and William J Buchanan, *Post quantum migration of tor*, Cryptology ePrint Archive, Paper 2025/479, 2025.

[31] Pierre-Augustin Berthet, *Ciphertext malleability in lattice-based KEMs as a countermeasure to side channel analysis*, Fundamenta Informaticae **194** (2025).

[32] Siddhartha Siddhiprada Bhoi, Arathi Arakala, Amy Beth Corman, and Asha Rao, *Post-quantum homomorphic encryption: A case for code-based alternatives*, Cryptography **9** (2025), no. 2.

[33] Sebastian Bitzer, Jean-Christophe Deneuville, Emma Munisamy, Bharath Purtipli, Stefan Ritterhoff, and Antonia Wachter-Zeh, *Hqc beyond the standard: Ciphertext compression and refined dfr analysis*, Advances in Cryptology – EUROCRYPT 2026 (Cham) (Joan Daemen and Emmanuel Thomé, eds.), Springer Nature Switzerland, 2026, pp. 544–574.

[34] Roman Bögli, Jonas Spieler, and Timo Kehrer, *BF-CBOM: Uncovering cryptographic assets through comparative CBOM analysis at scale*, Proceedings of the 34th IEEE/ACM International Conference on Program Comprehension (ICPC 2026) (Rio de Janeiro, Brazil), ACM, 2026, Preprint - Accepted for publication, pp. 1–5.

[35] Dan Boneh and Binyi Chen, *Latticefold: A lattice-based folding scheme and its applications to succinct proof systems*, International Conference on the Theory and Application of Cryptology and Information Security, Springer, 2025, pp. 330–362.

[36] ———, *Latticefold+: faster, simpler, shorter lattice-based folding for succinct proof systems*, Annual International Cryptology Conference, Springer, 2025, pp. 327–361.

[37] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry, *Random oracles in a quantum world*, Advances in Cryptology – ASIACRYPT 2011 (Berlin, Heidelberg)

(Dong Hoon Lee and Xiaoyun Wang, eds.), Springer Berlin Heidelberg, 2011, pp. 41–69.

[38] Dan Boneh and Mark Zhandry, *Secure signatures and chosen ciphertext security in a quantum computing world*, Advances in Cryptology – CRYPTO 2013 (Berlin, Heidelberg) (Ran Canetti and Juan A. Garay, eds.), Springer Berlin Heidelberg, 2013, pp. 361–379.

[39] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti, *A framework for practical anonymous credentials from lattices*, Advances in Cryptology – CRYPTO 2023 (Cham) (Helena Handschuh and Anna Lysyanskaya, eds.), Springer Nature Switzerland, 2023, pp. 384–417.

[40] Joppe W. Bos, Marc Gourjon, Joost Renes, Tobias Schneider, and Christine van Vredendaal, *Masking Kyber: First- and Higher-Order Implementations*, IACR Trans. Cryptographic Hardware and Embedded Systems (TCHES) **2021** (2021), no. 4, 173–214.

[41] John Bostanci, Boyang Chen, and Barak Nehoran, *Oracle separation between quantum commitments and quantum one-wayness*, Advances in Cryptology – EUROCRYPT 2025 (Cham) (Serge Fehr and Pierre-Alain Fouque, eds.), Springer Nature Switzerland, 2025, pp. 3–22.

[42] Leon Botros, Matthias J. Kannwischer, and Peter Schwabe, *Memory-Efficient High-Speed Implementation of Kyber on Cortex-M4*, Progress in Cryptology – AFRICACRYPT, LNCS, vol. 11627, Springer, 2019, pp. 209–228.

[43] Charles Bouillaguet, Thibauld Feneuil, Jules Maire, Matthieu Rivain, Julia Sauvage, and Damien Vergnaud, *Blinding post-quantum hash-and-sign signatures*, Cryptology ePrint Archive, Paper 2025/895, 2025.

[44] Gilles Brassard, Peter HØyer, and Alain Tapp, *Quantum cryptanalysis of hash and claw-free functions: Invited paper*, p. 163–169, Springer Berlin Heidelberg, 1998.

[45] Jacqueline Brendel, Rune Fiedler, Felix Günther, Christian Janson, and Douglas Stebila, *Post-quantum asynchronous deniable key exchange and the signal handshake*, Public-Key Cryptography – PKC 2022 (Cham) (Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, eds.), Springer International Publishing, 2022, pp. 3–34.

[46] Jacqueline Brendel, Marc Fischlin, Felix Günther, Christian Janson, and Douglas Stebila, *Challenges in proving post-quantum key exchanges based on key encapsulation mechanisms*, IACR Cryptol. ePrint Arch. **2019** (2019), 1356.

[47] Emmanuel Bresson, Olivier Chevassut, David Pointcheval, and Jean-Jacques Quisquater, *Provably authenticated group diffie-hellman key exchange*, Proceedings of the 8th ACM Conference on Computer and Communications Security (New York, NY, USA), CCS '01, Association for Computing Machinery, 2001, p. 255–264.

[48] Mike Burmester and Yvo Desmedt, *A secure and efficient conference key distribution system: Extended abstract*, Advances in Cryptology — EUROCRYPT'94, vol. 950, Springer Berlin Heidelberg, Berlin, Heidelberg, 1995, Series Title: Lecture Notes in Computer Science, pp. 275–286.

[49] Niklas Büscher, Daniel Demmler, Nikolaos P. Karvelas, Stefan Katzenbeisser, Juliane Krämer, Deevashwer Rathee, Thomas Schneider, and Patrick Struck, *Secure Two-Party Computation in a Quantum World*, Applied Cryptography and Network Security (Cham), vol. 12146, Springer International Publishing, 2020, Series Title: Lecture Notes in Computer Science, pp. 461–480.

[50] Robert Campbell, *Enterprise migration to post-quantum cryptography: Timeline analysis and strategic frameworks*, Computers **15** (2025), no. 1, 9.

[51] Joseph Carolan, *Compressed permutation oracles*, 2025.

[52] Joseph Carolan, Alexander Poremba, and Mark Zhandry, *(quantum) indifferentiability and pre-computation*, 2024.

[53] Tore Vincent Carstens, Ehsan Ebrahimi, Gelo Noel Tabia, and Dominique Unruh, *Relationships between quantum ind-cpa notions*, Theory of Cryptography (Cham) (Kobbi Nissim and Brent Waters, eds.), Springer International Publishing, 2021, pp. 240–272.

[54] André Chailloux, *Tight quantum security of the fiat-shamir transform for commit-and-open identification schemes with applications to post-quantum signature schemes*, 2021.

[55] Melissa Chase, Trevor Perrin, and Greg Zaverucha, *The signal private group system and anonymous credentials supporting efficient verifiable encryption*, Cryptology ePrint Archive, Paper 2019/1416, 2019.

[56] Binyi Chen, *Symphony: Scalable snarks in the random oracle model from lattice-based high-arity folding*, Cryptology ePrint Archive (2025).

[57] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang, *Hyperplonk: Plonk with linear-time prover and high-degree custom gates*, Advances in Cryptology – EUROCRYPT 2023 (Cham) (Carmit Hazay and Martijn Stam, eds.), Springer Nature Switzerland, 2023, pp. 499–530.

[58] Bohang Chen, Shuai Han, and Shengli Liu, *On post-quantum signature with message recovery from hash-and-sign in qrom*, Public-Key Cryptography – PKC 2026 (Cham) (Shi Bai and Edoardo Persichetti, eds.), Springer Nature Switzerland, 2026, pp. 302–333.

[59] Jinrong Chen, Yi Wang, Rongmao Chen, Xinyi Huang, and Wei Peng, *Tighter proofs for pke-to-kem transformation in the quantum random oracle model*, Advances in Cryptology – ASIACRYPT 2024 (Singapore) (Kai-Min Chung and Yu Sasaki, eds.), Springer Nature Singapore, 2025, pp. 101–133.

[60] Lijie Chen and Ramis Movassagh, *Quantum Merkle trees*, Quantum **8** (2024), 1380.

[61] Céline Chevalier, Ehsan Ebrahimi, and Quoc-Huy Vu, *On security notions for encryption in a quantum world*, Progress in Cryptology – INDOCRYPT 2022 (Cham) (Takanori Isobe and Santanu Sarkar, eds.), Springer International Publishing, 2022, pp. 592–613.

[62] Rojin Chhetri, *Pqc benchmarks on arm cortex-m0+ (rp2040)*, 2026.

[63] Jason Chia, Ji-Jian Chin, and Sook-Chin Yip, *Digital signature schemes with strong existential unforgeability*, F1000Research **10** (2021), 931.

[64] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa, *A black-box approach to post-quantum zero-knowledge in constant rounds*, Advances in Cryptology – CRYPTO 2021 (Cham) (Tal Malkin and Chris Peikert, eds.), Springer International Publishing, 2021, pp. 315–345.

[65] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward, *Marlin: Preprocessing zksnarks with universal and updatable srs*, Advances in Cryptology – EUROCRYPT 2020 (Cham) (Anne Canteaut and Yuval Ishai, eds.), Springer International Publishing, 2020, pp. 738–768.

[66] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner, *Succinct arguments in the quantum random oracle model*, Theory of Cryptography (Cham) (Dennis Hofheinz and Alon Rosen, eds.), Springer International Publishing, 2019, pp. 1–29.

[67] Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao, *On the compressed-oracle technique, and post-quantum security of proofs of sequential work*, Advances in Cryptology – EUROCRYPT 2021 (Cham) (Anne Canteaut and François-Xavier Standaert, eds.), Springer International Publishing, 2021, pp. 598–629.

[68] Carlos Cid, David Elkouss, and Manuel Goulão, *Superposition attacks against LPN-based authentication protocols*, Cryptology ePrint Archive, Paper 2025/1506, 2025.

[69] NSA CISA, *Quantum-readiness: Migration to post-quantum cryptography*, Tech. report, CISA, Tech. Rep, 2023.

[70] Cloud Security Alliance, *Strategic post-quantum cryptography migration: Enterprise roadmap v1*, 2026, Accessed: 24 April 2026.

[71] Euroean Comission, *A coordinated implementation roadmap for the transition to post-quantum cryptography*, `https://digital-strategy.ec.europa.eu/en/library/coordinated-implementation-roadmap-transition-post-quantum-cryptography`, Accessed: 2026-06-26.

[72] European Commission, *A coordinated implementation roadmap for the transition to post-quantum cryptography*, Tech. report, NIS Cooperation Group Work Stream on Post-Quantum Cryptography, Brussels, June 2025, First deliverable following the European Commission Recommendation on a Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography.

[73] ———, *Security of implementations of post-quantum cryptography algorithms*, Horizon Europe Programme, 2025.

[74] Cas Cremers, Samed Düzlü, Rune Fiedler, Marc Fischlin, and Christian Janson, *Buffing signature schemes beyond unforgeability and the case of post-quantum signatures*, 2021 IEEE Symposium on Security and Privacy (SP), 2021, pp. 1696–1714.

[75] Eric Crockett, Christian Paquin, and Douglas Stebila, *Prototyping post-quantum and hybrid key exchange and authentication in TLS and SSH*, Cryptology ePrint Archive, Paper 2019/858, 2019.

[76] Jan Czajkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur, *Quantum lazy sampling and game-playing proofs for quantum indifferentiability*, 2021.

[77] Vivek Dabra, Saru Kumari, Anju Bala, and Sonam Yadav, *SL3PAKE: Simple Lattice-based Three-party Password Authenticated Key Exchange for post-quantum world*, Journal of Information Security and Applications **84** (2024), 103826.

[78] Nilanjan Datta, Avijit Dutta, Sougata Mandal, Hrithik Nandi, and Amlan Sinha, *Post-quantum security of keyed sum of permutations and its siblings*, Cryptology ePrint Archive, Paper 2026/426, 2026.

[79] Christian Decker and Roger Wattenhofer, *Bitcoin transaction malleability and mtgox*, Computer Security - ESORICS 2014 (Cham) (Mirosław Kutyłowski and Jaideep Vaidya, eds.), Springer International Publishing, 2014, pp. 313–326.

[80] Rafael del Pino and Shuichi Katsumata, *A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling*, Cryptology ePrint Archive, Paper 2022/834, 2022.

[81] K. Delgado-Vargas et al., *A look at side channel attacks on post-quantum cryptography*, Computación y Sistemas (2024).

[82] Department for Science, Innovation and Technology, *Regulator and Industry Perspectives on the Current Plan for PQC Transition*, Tech. report, UK Government, November 2025.

[83] Nicolas Desmoulins, Antoine Dumanois, Seyni Kane, and Jacques Traoré, *Making bbs anonymous credentials eidas 2.0 compliant*, Cryptology ePrint Archive, Paper 2025/619, 2025.

[84] Daniel Dinu, *Migration to post-quantum cryptography: From ECDSA to ML-DSA*, Cryptology ePrint Archive, Paper 2025/2025, 2025.

[85] P. Dobias et al., *Sok: Reassessing side-channel vulnerabilities and countermeasures in pqc implementations*, IACR ePrint Archive (2025).

[86] Yevgeniy Dodis, Daniel Jost, Shuichi Katsumata, Thomas Prest, and Rolfe Schmidt, *Triple ratchet: A bandwidth efficient hybrid-secure signal protocol*, Cryptology ePrint Archive, Paper 2025/078, 2025.

[87] Jelle Don, Serge Fehr, and Christian Majenz, *The measure-and-reprogram technique 2.0: Multi-round fiat-shamir and more*, Advances in Cryptology – CRYPTO 2020 (Cham) (Daniele Micciancio and Thomas Ristenpart, eds.), Springer International Publishing, 2020, pp. 602–631.

[88] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner, *Security of the fiat-shamir transformation in the quantum random-oracle model*, p. 356–383, Springer International Publishing, 2019.

[89] ———, *Online-extractability in the quantum random-oracle model*, Advances in Cryptology – EUROCRYPT 2022 (Cham) (Orr Dunkelman and Stefan Dziembowski, eds.), Springer International Publishing, 2022, pp. 677–706.

[90] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila, *A cryptographic analysis of the tls 1.3 handshake protocol*, vol. 34, ch. 37, Springer International Publishing, International Association for Cryptographic Research, 2021.

[91] Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail, *Fiat-Shamir for Proofs Lacks a Proof Even in the Presence of Shared Entanglement*, Quantum **8** (2024), 1568.

[92] Ehsan Ebrahimi and Jeroen van Wier, *Post-quantum plaintext-awareness*, Post-Quantum Cryptography (Cham) (Jung Hee Cheon and Thomas Johansson, eds.), Springer International Publishing, 2022, pp. 260–285.

[93] Entrust, *Building telecom resilience in the quantum era*, April 2026, Accessed: 2026-05-18.

[94] European Commission, *Eudi wallet reference implementation documentation*, `https://docs.eudi.dev/latest/`, 2025, Documentation for the EUDI Wallet Reference Implementation, accessed 28 June 2026.

[95] European Digital Identity Cooperation Group, *European digital identity wallet architecture and reference framework*, `https://eudi.dev/2.9.0/architecture-and-reference-framework-main/`, 2026.

[96] Pouria Fallahpour, Serge Fehr, and Yu-Hsuan Huang, *Tighter quantum security for fiat-shamir-with-aborts and hash-and-sign-with-retry signatures*, Cryptology ePrint Archive, Paper 2025/985, 2025.

[97] Xuejun Fan, Fei Zhao, and Xiu Xu, *Multi-party post-quantum key exchange schemes*, Journal of Information Security and Applications **95** (2025), 104288.

[98] Niels Ferguson, Samuel Lee, Mitch Lindgren, and Aaron Hadley, *Do not use pure ml-dsa to sign arbitrary length messages*, Tech. report, 2025.

[99] Rune Fiedler and Roman Langrehr, *On deniable authentication against malicious verifiers*, Cryptology ePrint Archive, Paper 2025/470, 2025.

[100] Marc Fischlin and Tobias Schmalz, *Quantum-resistant authenticated encryption variants of gcm*, IACR Transactions on Symmetric Cryptology **2026** (2026), no. 2.

[101] Tim Fritzmann, Georg Sigl, and Johanna Sepúlveda, *RISQ-V: Tightly Coupled RISC-V Accelerators for Post-Quantum Cryptography*, IACR Trans. Cryptographic Hardware and Embedded Systems (TCHES), vol. 2020, 2020, pp. 239–280.

[102] Philippe Gaborit, Carlos Aguilar-Melchor, Nicolas Aragon, Slim Bettaieb, Loïc Bidoux, Olivier Blazy, Jean-Christophe Deneuville, Edoardo Persichetti, Gilles Zémor, Jurjen Bos, et al., *Hamming quasi-cyclic (HQC)*, NIST Post-Quantum Standardization, 4th Round; National Institute of Standards and Technology: Gaithersburg, MD, USA (2025).

[103] Tommaso Gagliardoni, Andreas Hülsing, and Christian Schaffner, *Semantic security and indistinguishability in the quantum world*, Advances in Cryptology – CRYPTO 2016 (Berlin, Heidelberg) (Matthew Robshaw and Jonathan Katz, eds.), Springer Berlin Heidelberg, 2016, pp. 60–89.

[104] Chaya Ganesh, Sikhar Patranabis, and Nitin Singh, *Samaritan: Linear-time prover snark from new multilinear polynomial commitments*, Advances in Cryptology – ASIACRYPT 2025 (Singapore) (Goichiro Hanaoka and Bo-Yin Yang, eds.), Springer Nature Singapore, 2026, pp. 456–489.

[105] Jiangxia Ge, Heming Liao, and Rui Xue, *Measure-rewind-extract: Tighter proofs of one-way to hiding and cca security in the quantum random oracle model*, Advances in Cryptology – ASIACRYPT 2024 (Singapore) (Kai-Min Chung and Yu Sasaki, eds.), Springer Nature Singapore, 2025, pp. 3–34.

[106] Federico Giacon, Felix Heuer, and Bertram Poettering, *Kem combiners*, Public-Key Cryptography – PKC 2018 (Cham) (Michel Abdalla and Ricardo Dahab, eds.), Springer International Publishing, 2018, pp. 190–218.

[107] Alexandre Augusto Giron, Ricardo Custódio, and Francisco Rodríguez-Henríquez, *Post-quantum hybrid key exchange: a systematic mapping study*, Journal of Cryptographic Engineering **13** (2023), no. 1, 71–88.

[108] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum, *Delegating computation: Interactive proofs for muggles*, J. ACM **62** (2015), no. 4.

[109] Trusted Computing Group, *State of PQC Readiness: Are Businesses Prepared For Q-Day?*, Tech. report, Trusted Computing Group, November 2025.

[110] Lov K. Grover, *A fast quantum mechanical algorithm for database search*, Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96 (Philadelphia, Pennsylvania, United States), ACM Press, 1996, pp. 212–219.

[111] Sghaier Guizani, Tehseen Mazhar, and Habib Hamam, *Quantum Secure Multiparty Computation: Bridging Privacy, Security, and Scalability in the Post-Quantum Era*, Computers, Materials & Continua **87** (2026), no. 1, 1–10 (en).

[112] Aditya Gulati, Yao-Ting Lin, Tomoyuki Morimae, and Shogo Yamada, *Black-box separation between pseudorandom unitaries, pseudorandom isometries, and pseudorandom function-like states*, 2025.

[113] Yanpei Guo, Xuanming Liu, Kexi Huang, Wenjie Qu, Tianyang Tao, and Jiaheng Zhang, *Deepfold: efficient multilinear polynomial commitment from reed-solomon code and its application to zero-knowledge proofs*, Proceedings of the 34th USENIX Conference on Security Symposium (USA), SEC '25, USENIX Association, 2025.

[114] Tomas Gustavsson, *Interoperability in action: Post-quantum cryptography*, Presentation at the OpenSSL Conference 2025, Prague, Czech Republic, October 2025, Presentation slides.

[115] Ben Hamlin and Fang Song, *Quantum security of hash functions and property-preservation of iterated hashing*, Post-Quantum Cryptography (Cham) (Jintai Ding and Rainer Steinwandt, eds.), Springer International Publishing, 2019, pp. 329–349.

[116] Johanna Henrich, Andreas Heinemann, Alex Wiesmaier, and Nicolai Schmitt, *Performance impact of pqc kems on tls 1.3 under varying network characteristics*, Information Security (Cham) (Elias Athanasopoulos and Bart Mennink, eds.), Springer Nature Switzerland, 2023, pp. 267–287.

[117] Minki Hhan and Shogo Yamada, *Pseudorandom function-like states from common haar unitary*, Theory of Cryptography (Cham) (Benny Applebaum and Huijia (Rachel) Lin, eds.), Springer Nature Switzerland, 2026, pp. 134–165.

[118] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz, *A modular analysis of the fujisaki-okamoto transformation*, Theory of Cryptography (Cham) (Yael Kalai and Leonid Reyzin, eds.), Springer International Publishing, 2017, pp. 341–371.

[119] Sanzida Hoque, Abdullah Aydeger, Engin Zeydan, and Madhusanka Liyanage, *Analysis of post-quantum cryptography in user equipment in 5g and beyond*, 2025 IEEE 50th Conference on Local Computer Networks (LCN), 2025, pp. 1–9.

[120] Akinori Hosoyamada, *The impossibility of post-quantum public indifferentiability for merkle-damgard*, Cryptology ePrint Archive, Paper 2026/128, 2026.

[121] Akinori Hosoyamada, *Post-quantum security of keyed sponge-based constructions through a modular approach*, Advances in Cryptology – ASIACRYPT 2025 (Singapore) (Goichiro Hanaoka and Bo-Yin Yang, eds.), Springer Nature Singapore, 2026, pp. 411–445.

[122] Akinori Hosoyamada and Tetsu Iwata, *On tight quantum security of hmac and nmac in the quantum random oracle model*, Advances in Cryptology – CRYPTO 2021 (Cham) (Tal Malkin and Chris Peikert, eds.), Springer International Publishing, 2021, pp. 585–615.

[123] Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz, *Failing gracefully: Decryption failures and the fujisaki-okamoto transform*, Advances in Cryptology – ASIACRYPT 2022 (Cham) (Shweta Agrawal and Dongdai Lin, eds.), Springer Nature Switzerland, 2022, pp. 414–443.

[124] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby, *A pseudorandom generator from any one-way function*, SIAM Journal on Computing **28** (1999), no. 4, 1364–1396.

[125] Loïs Huguenin-Dumittan and Serge Vaudenay, *On ind-qcca security in the rom and its applications*, Advances in Cryptology – EUROCRYPT 2022 (Cham) (Orr Dunkelman and Stefan Dziembowski, eds.), Springer International Publishing, 2022, pp. 613–642.

[126] I. Ingemarsson, D. Tang, and C. Wong, *A conference key distribution system*, IEEE Transactions on Information Theory **28** (1982), no. 5, 714–720.

[127] Marin Ivezic, *Dos & don'ts of crypto inventories for quantum readiness*, 2024, Accessed: 2026-05-14.

[128] _____, *How to perform a comprehensive quantum readiness cryptographic inventory*, April 2024, Accessed: 2026-05-14.

[129] Dennis Jackson, Cas Cremers, Katriel Cohn-Gordon, and Ralf Sasse, *Seems legit: Automated analysis of subtle attacks on protocols that use signatures*, Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (New York, NY, USA), CCS '19, Association for Computing Machinery, 2019, p. 2165–2180.

[130] Pratima Jana and Ratna Dutta, *Quantum-safe identity-binding password authenticated key exchange protocols*, Cryptology ePrint Archive, Paper 2025/2107, 2025.

[131] Jonas Janneck, *Bird of prey: Practical signature combiners preserving strong unforgeability*, Advances in Cryptology – EUROCRYPT 2026 (Cham) (Joan Daemen and Emmanuel Thomé, eds.), Springer Nature Switzerland, 2026, pp. 222–250.

[132] Christian Janson and Patrick Struck, *Sponge-based authenticated encryption: Security against quantum attackers*, Post-Quantum Cryptography (Cham) (Jung Hee Cheon and Thomas Johansson, eds.), Springer International Publishing, 2022, pp. 230–259.

[133] P. Jedlicka et al., *On secure and side-channel resistant hardware implementations of pqc*, ARES Conference, 2022.

[134] Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders, *Lattice signature with efficient protocols, application to anonymous credentials*, Cryptology ePrint Archive, Paper 2022/509, 2022.

[135] Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders, *Lattice signature with efficient protocols, application to anonymous credentials*, Advances in Cryptology – CRYPTO 2023 (Cham) (Helena Handschuh and Anna Lysyanskaya, eds.), Springer Nature Switzerland, 2023, pp. 351–383.

[136] Julien Juaneda, Marina Dehez-Clementi, Jean-Christophe Deneuville, and Jérôme Lacan, *RHQC: post-quantum ratcheted key exchange from coding assumptions*, Cryptology ePrint Archive, Paper 2025/481, 2025.

[137] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg, *Constant-size commitments to polynomials and their applications*, Advances in Cryptology - ASIACRYPT 2010 (Berlin, Heidelberg) (Masayuki Abe, ed.), Springer Berlin Heidelberg, 2010, pp. 177–194.

[138] Shuichi Katsumata, Guilhem Niot, Ida Tucker, and Thom Wiggers, *Comprehensive deniability analysis of signal handshake protocols: X3DH, PQXDH to fully post-quantum with deniable ring signatures*, Cryptology ePrint Archive, Paper 2025/1090, 2025.

[139] Jonathan Katz and Ben Sela, *A quantum "lifting theorem" for constructions of pseudorandom generators from random oracles*, 2025.

[140] Victor Youdom Kemmoe, Anna Lysyanskaya, and Ngoc Khanh Nguyen, *Lattice-based accumulator and application to anonymous credential revocation*, Cryptology ePrint Archive, Paper 2025/1099, 2025.

[141] Marcel Kempf, Nikolas Gauder, Benedikt Jaeger, Johannes Zirngibl, and Georg Carle, *A quantum of quic: Dissecting cryptography with post-quantum insights*, 2024 IFIP Networking Conference (IFIP Networking), 2024, pp. 195–203.

[142] Dmitry Khovratovich, Ron D. Rothblum, and Lev Soukhanov, *How to prove false statements: Practical attacks on fiat-shamir*, Advances in Cryptology – CRYPTO 2025 (Cham) (Yael Tauman Kalai and Seny F. Kamara, eds.), Springer Nature Switzerland, 2025, pp. 3–26.

[143] Dakshita Khurana and Kabir Tomer, *Commitments from quantum one-wayness*, Proceedings of the 56th Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC 2024, Association for Computing Machinery, 2024, p. 968–978.

[144] Michael Klooss, Russell W. F. Lai, Ngoc Khanh Nguyen, Michał Osadnik, and Lorenzo Tucci, *RoKoko: Lattice-based succinct arguments, a committed refinement*, Cryptology ePrint Archive, Paper 2026/575, 2026.

[145] Haruhisa Kosuge and Keita Xagawa, *The security of hash-and-sign with retry against superposition attacks*, Public-Key Cryptography – PKC 2025 (Cham) (Tibor Jager and Jiaxin Pan, eds.), Springer Nature Switzerland, 2025, pp. 317–349.

[146] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla, *Nova: Recursive zero-knowledge arguments from folding schemes*, Annual International Cryptology Conference, Springer, 2022, pp. 359–388.

[147] Jan Kowalleck and Steve Springett, *Ecma-424: Cyclonedx bill of materials specification*, Tech. report, Ecma International, December 2025, Defines the CycloneDX v1.7 specification and includes the Cryptography Bill of Materials (CBOM) capability.

[148] Juliane Krämer and Patrick Struck, *On quantum ciphertext indistinguishability, recoverability, and oaep*, Post-Quantum Cryptography (Cham) (Jung Hee Cheon and Thomas Johansson, eds.), Springer International Publishing, 2022, pp. 286–306.

[149] Alexander Krause, Harjot Kaur, Jan H Klemmer, Oliver Wiese, and Sascha Fahl, *" that's my perspective from 30 years of doing this": An interview study on practices, experiences, and challenges of updating cryptographic code*, 34th USENIX Security Symposium (USENIX Security 25), 2025, pp. 2907–2926.

[150] Ehren Kret and Rolfe Schmidt, *The pqxdh key agreement protocol.*, url:https://signal. org/docs/specifications/pqxdh/pqxdh.pdf, 2023.

[151] Veronika Kuchta, Amin Sakzad, Damien Stehlé, Ron Steinfeld, and Shi-Feng Sun, *Measure-rewind-measure: Tighter quantum random oracle model proofs for one-way to hiding and cca security*, Advances in Cryptology – EUROCRYPT 2020 (Cham) (Anne Canteaut and Yuval Ishai, eds.), Springer International Publishing, 2020, pp. 703–728.

[152] Shuto Kuriyama, Russell WF Lai, Michał Osadnik, and Lorenzo Tucci, *Salsaa–sumcheck-aided lattice-based succinct arguments and applications*, Cryptology ePrint Archive (2025).

[153] Qiqi Lai, Chongshen Chen, Feng-Hao Liu, Anna Lysyanskaya, and Zhedong Wang, *Lattice-based commit-transferrable signatures and applications to anonymous credentials*, Cryptology ePrint Archive, Paper 2023/766, 2023.

[154] ———, *Lattice-based commit-transferrable signatures and applications to anonymous credentials*, Cryptology ePrint Archive, Paper 2023/766, 2023.

[155] Zhenzhi Lai, Ruiyi Zhang, Zhiyuan Zhang, Julius Hermelink, Michael Schwarz, Van-Thuan Pham, and Udaya Parampalli, *You only decapsulate once: Ciphertext-independent single-trace passive side-channel attacks on HQC*, Cryptology ePrint Archive, Paper 2025/2162, 2025.

[156] Brian LaMacchia, Matt Campagna, and William Gropp, *The post-quantum cryptography transition: Making progress, but still a long road ahead*, arXiv preprint arXiv:2503.04806 (2025).

[157] Nathalie Lang, Jannis Leuther, and Stefan Lucks, *Generic composition: From classical to quantum security*, Post-Quantum Cryptography (Cham) (Magali Bardet and Ruben Niederhagen, eds.), Springer Nature Switzerland, 2026, pp. 37–70.

[158] Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu, *How (not) to build quantum pke in minicrypt*, Advances in Cryptology – CRYPTO 2024 (Cham) (Leonid Reyzin and Douglas Stebila, eds.), Springer Nature Switzerland, 2024, pp. 152–183.

[159] Qinyi Li, Lise Millerjord, and Colin Boyd, *TEAKEX: TESLA-Authenticated Group Key Exchange*, Information Security and Privacy, vol. 15658, Springer Nature Singapore, Singapore, 2025, Series Title: Lecture Notes in Computer Science, pp. 289–309.

[160] Z Li, Y Zhang, F Zhang, and Y Yang, *Ideal lattice-based oblivious transfer protocol of provably secure*, Appl Res Comput **34** (2017), no. 1, 242–245.

[161] Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang, *Adaptive Oblivious Transfer with Access Control from Lattice Assumptions*, Advances in Cryptology – ASIACRYPT 2017 (Cham) (Tsuyoshi Takagi and Thomas Peyrin, eds.), Springer International Publishing, 2017, pp. 533–563.

[162] Qipeng Liu, *Non-uniformity and quantum advice in the quantum random oracle model*, Advances in Cryptology – EUROCRYPT 2023 (Cham) (Carmit Hazay and Martijn Stam, eds.), Springer Nature Switzerland, 2023, pp. 117–143.

[163] Qipeng Liu and Mark Zhandry, *Revisiting post-quantum fiat-shamir*, Cryptology ePrint Archive, Paper 2019/262, 2019.

[164] Torsten Lodderstedt, Kristina Yasuda, Tim Looker, and Peter Bastian, *Openid for verifiable credential issuance*, `https://openid.github.io/OpenID4VCI/openid-4-verifiable-credential-issuance-1_1-wg-draft.html`, June 2026, OpenID Foundation Digital Credentials Protocols Working Group, Working Group Draft, accessed 27 June 2026.

[165] ———, *Openid for verifiable presentations*, `https://openid.github.io/OpenID4VP/openid-4-verifiable-presentations-1_0.html`, June 2026, OpenID Foundation Digital Credentials Protocols Working Group, accessed 27 June 2026.

[166] Alex Lombardi, Fermi Ma, and Nicholas Spooner, *Post-quantum zero knowledge, revisited or: How to do quantum rewinding undetectably*, 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS), 2022, pp. 851–859.

[167] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan, *Algebraic methods for interactive proof systems*, J. ACM **39** (1992), no. 4, 859–868.

[168] Chujiao Ma, Luis Colon, Joe Dera, Bahman Rashidi, and Vaibhav Garg, *Caraf: crypto agility risk assessment framework*, Journal of Cybersecurity **7** (2021), no. 1, tyab013.

[169] Christian Majenz and Jaya Sharma, *Security of the fischlin transform in quantum random oracle model*, 2026.

[170] Victor Maram and Keita Xagawa, *Post-quantum anonymity of kyber*, Public-Key Cryptography – PKC 2023 (Alexandra Boldyreva and Vladimir Kolesnikov, eds.), Lecture Notes in Computer Science, vol. 13940, Springer, Cham, 2023.

[171] John Preuß Mattsson, Erik Thormarker, Masoud Asadi, and Sini Ruohomaa, *Migrating telecom to quantum-resistant cryptography on a global scale*, Ericsson Technology Review (2026).

[172] Vasileios Mavroeidis, Kamer Vishi, Mateusz D. Zych, and Audun Jøsang, *The Impact of Quantum Computing on Present Cryptography*, International Journal of Advanced Computer Science and Applications **9** (2018), no. 3, arXiv:1804.00200 [cs.CR].

[173] Dominic Mayers, *Unconditionally secure quantum bit commitment is impossible*, Phys. Rev. Lett. **78** (1997), 3414–3417.

[174] Bo Mi, Darong Huang, Shaohua Wan, Yu Hu, and Kim-Kwang Raymond Choo, *A post-quantum light weight 1-out-n oblivious transfer protocol*, Computers & Electrical Engineering **75** (2019), 90–100.

[175] Microsoft Security Blog, *Building your cryptographic inventory: A customer strategy for cryptographic posture management*, April 2026, Accessed: 2026-05-14.

[176] Dustin Moody, Ray Perlner, Andrew Regenscheid, Angela Robinson, and David Cooper, *Transition to post-quantum cryptography standards*, NIST Internal Report (IR) NIST IR 8547 ipd, National Institute of Standards and Technology, 2024.

[177] Splunk Nasreddine Bencherchali, *Detection: Cisco secure firewall - blacklisted ssl certificate fingerprint*, `https://research.splunk.com/network/c43f7b49-2dab-4e76-892e-7f971c2f20f1/`, Accessed: 2026-06-26.

[178] Christian Näther, Daniel Herzinger, Stefan-Lukas Gazdag, Jan-Philipp Steghöfer, Simon Daum, and Daniel Loebenberger, *Migrating software systems toward post-quantum cryptography-a systematic literature review*, IEEE access **12** (2024), 132107–132126.

[179] National Cyber Security Centre (UK), *Timelines for migration to post-quantum cryptography*, 2023, Accessed: 24 April 2026.

[180] National Institute of Standards and Technology, *Digital signature standard (dss)*, Federal Information Processing Standards Publication NIST FIPS 186-5, National Institute of Standards and Technology, Gaithersburg, MD, February 2023.

[181] ———, *Module-lattice-based key-encapsulation mechanism standard*, Federal Information Processing Standards (FIPS) 203, National Institute of Standards and Technology, Gaithersburg, MD, August 2024.

[182] Tristan Nemoz, Zoé Amblard, and Aurélien Dupin, *Characterizing the qind-qcpa (in)security of the cbc, cfb, ofb and ctr modes of operation*, Post-Quantum Cryptography (Cham) (Thomas Johansson and Daniel Smith-Tone, eds.), Springer Nature Switzerland, 2023, pp. 445–475.

[183] Ngoc Khanh Nguyen and Gregor Seiler, *Greyhound: Fast polynomial commitments from lattices*, Advances in Cryptology – CRYPTO 2024 (Cham) (Leonid Reyzin and Douglas Stebila, eds.), Springer Nature Switzerland, 2024, pp. 243–275.

[184] Wilson Nguyen and Srinath Setty, *Neo and superneo: Post-quantum folding with pay-per-bit costs over small fields*, Cryptology ePrint Archive (2026).

[185] Markku-Juhani O. Saarinen, *Mobile energy requirements of the upcoming nist post-quantum cryptography standards*, 2020 8th IEEE International Conference on Mobile Cloud Computing, Services, and Engineering (MobileCloud), 2020, pp. 23–30.

[186] National Institute of Standards and Technology, *Migration to post-quantum cryptography: Preparation for considering the implementation and adoption of quantum safe cryptography*, NIST Special Publication (NCCoE Practice Guide) SP 1800-38A, National Institute of Standards and Technology, 2023, Draft.

[187] Office of Management and Budget, *Report on Post-Quantum Cryptography*, Tech. report, Executive Office of the President of the United States, July 2024.

[188] Akinlemi Olushola and S. P. Meenakshi, *Design and implementation of an authenticated post-quantum session protocol using ML-KEM (Kyber), ML-DSA (Dilithium), and AES-256-GCM*, Frontiers in Physics **13** (2026), 1723966.

[189] A. Omoseebi et al., *Ml-assisted side-channel analysis on pqc algorithms*, Proceedings of ResearchGate Preprint, 2025.

[190] David Ott, Kenny Paterson, and Dennis Moreau, *Where is the research on cryptographic transition and agility?*, Communications of the ACM **66** (2023), no. 4, 29–32.

[191] David Ott, Christopher Peikert, et al., *Identifying research challenges in post quantum cryptography migration and cryptographic agility*, arXiv preprint arXiv:1909.07353 (2019).

[192] M. Ounsworth, J. Gray, M. Pala, J. Klaussner, and S. Fluhrer, *Composite ml-dsa for use in x.509 public key infrastructure*, `https://www.ietf.org/archive/id/draft-ietf-lamps-pq-composite-sigs-13.html`, Accessed: 2026-06-26.

[193] OWASP, *Scwe-054: Signature malleability*, `https://scs.owasp.org/SCWE/SCSVS-CRYPTO/SCWE-054/`, Accessed: 2026-06-26.

[194] OWASP CycloneDX Project, *Authoritative guide to cbom: Implement cryptography bill of materials for post-quantum systems and applications*, OWASP Foundation, 2024, Informative guidance; the normative standard is ECMA-424.

[195] Thanmai Sai P., K S Shailesh, Sanjay C Hiremath, Nithin Ramakrishnan, Anooja Ali, and Ajil A, *Drop-in-replaceability analysis using post quantum cryptography algorithms*, 2023 Fourth International Conference on Smart Technologies in Computing, Electrical and Electronics (ICSTCEE), 2023, pp. 1–6.

[196] Trevor Perrin, Moxie Marlinspike, and Rolfe Schmidt, *The double ratchet algorithm*, Specification, Signal Messenger, 2025.

[197] Peter Pessl and Robert Primas, *More Practical Single-Trace Attacks on the Number Theoretic Transform*, Progress in Cryptology – LATINCRYPT, LNCS, vol. 11774, Springer, 2019, pp. 130–149.

[198] Ponemon Institute LLC, *2026 global state of post-quantum and cryptographic security trends*, Technical report, Entrust, January 2026, Sponsored by Entrust.

[199] Post-Quantum Cryptography Coalition (PQCC), *Post-quantum cryptography (pqc) migration roadmap*, Tech. report, MITRE Corporation, May 2025, Approved for public release. Distribution unlimited. 24-03931-7.

[200] QRAMM, *Cryptographic inventory guide*, December 2024, Accessed: 2026-05-18.

[201] ———, *Pqc migration planning*, December 2024, Accessed: 2026-05-18.

[202] Prasanna Ravi, Shivam Bhasin, Sujoy Sinha Roy, and Anupam Chattopadhyay, *On exploiting message leakage in (few) nist pqc candidates for practical message recovery attacks*, IEEE Transactions on Information Forensics and Security **17** (2022), 684–699.

[203] Prasanna Ravi, Anupam Chattopadhyay, Jan-Pieter D'Anvers, and Anubhab Baksi, *Side-Channel and Fault-Injection Attacks over Lattice-Based Post-Quantum Schemes (Kyber, Dilithium): Survey and New Results*, ACM Trans. on Embedded Computing Systems (TECS) **23** (2024), no. 2, 1–54.

[204] Ueno Rei, Xagawa Keita, Tanaka Yutaro, Ito Akira, Takahashi Junko, and Homma Naofumi, *Curse of re-encryption: A generic power/em analysis on post-quantum kems*, IACR Transactions on Cryptographic Hardware and Embedded Systems (2021), 296–322.

[205] Eric Rescorla, *The transport layer security (tls) protocol version 1.3*, Tech. Report 8446, IETF, 2018.

[206] Sebastien Riou, Jong-Yeong Park, Liga Anwar, Axel Poschmann, and Michael Hutter, *Apples, oranges, and signatures: Pitfalls and methodology in ML-DSA benchmarking*, `https://github.com/sebastien-riou/BTVG-MLDSA/blob/main/2026-05-10%20-%20MAgiCS%20-%20MLDSA%20Sign%20Benchmarking.pdf`, Accessed: 2026-06-26.

[207] M.-J. O. Saarinen, *Introduction to side-channel security of nist pqc standards*, NIST PQC Seminar, 2023.

[208] Manuel B. Santos, Armando N. Pinto, and Paulo Mateus, *Quantum and classical oblivious transfer: A comparative analysis*, IET Quantum Communication **2** (2021), no. 2, 42–53.

[209] Sophie Schmieg, *Hashml-dsa considered harmful*, `https://keymaterial.net/2024/11/05/hashml-dsa-considered-harmful/`, Accessed: 2026-06-25.

[210] Nicolai Schmitt, Johanna Henrich, Dominik Heinz, Nouri Alnahawi, and Alexander Wiesmaier, *On criteria and tooling for cryptographic inventories*, Sicherheit 2024, Gesellschaft für Informatik eV, 2024, pp. 49–63.

[211] Juliane Schmüser, Philip Klostermeyer, Kay Friedrich, and Sascha Fahl, *"i'm pretty expert and i still screw it up": Qualitative insights into experiences and challenges of designing and implementing cryptographic library apis*, 2025 IEEE Symposium on Security and Privacy (SP), IEEE, 2025, pp. 2322–2340.

[212] Maximilian Schöffel, Frederik Lauer, Carl C. Rheinländer, and Norbert Wehn, *Secure iot in the era of quantum computers—where are the bottlenecks?*, Sensors **22** (2022), no. 7.

[213] Srinath Setty, *Spartan: Efficient and general-purpose zksnarks without trusted setup*, Advances in Cryptology – CRYPTO 2020 (Cham) (Daniele Micciancio and Thomas Ristenpart, eds.), Springer International Publishing, 2020, pp. 704–737.

[214] Abdul Basit Shahid, Khwaja Mansoor, Yawar Abbas Bangash, Waseem Iqbal, and Shynar Mussiraliyeva, *Post-quantum cryptographic authentication protocol for industrial iot using lattice-based cryptography*, Scientific Reports (2026).

[215] Tianshu Shan, Jiangxia Ge, and Rui Xue, *Qcca-secure generic transformations in the quantum random oracle model*, Public-Key Cryptography – PKC 2023 (Cham) (Alexandra Boldyreva and Vladimir Kolesnikov, eds.), Springer Nature Switzerland, 2023, pp. 36–64.

[216] Rentaro Shiba and Tetsu Iwata, *Multi-key security in the quantum world: Revisiting tweakable even-mansour and fx*, IACR Transactions on Symmetric Cryptology **2026** (2026), no. 2.

[217] Dimitrios Sikeridis, David Ott, Sean Huntley, Shivali Sharma, Vasantha Kumar Dhanasekar, Megha Bansal, Akhilesh Kumar, Anwitha UN, Daniel Beveridge, and Sairam Veeraswamy, *Elca: Introducing enterprise-level cryptographic agility for a post-quantum era*, Cryptology ePrint Archive (2023).

[218] Estevenson Solano, *Cryptographic inventory: Key to visibility, digital trust and post-quantum readiness*, nov 2025, Accessed: 2026-05-18.

[219] Patrick Struck and Maximiliane Weishäupl, *A framework for advanced signature notions*, Public-Key Cryptography – PKC 2026 (Cham) (Shi Bai and Edoardo Persichetti, eds.), Springer Nature Switzerland, 2026, pp. 156–188.

[220] Hong-Wei Sun, Fei Gao, Rong-Xue Xu, Dan-Dan Li, Zhen-Qiang Li, and Ke-Jia Zhang, *Quantum key-recovery attacks on permutation-based pseudorandom functions*, IEEE Internet of Things Journal **12** (2025), no. 17, 35692–35704.

[221] YanJin Tan, JunTao Gao, and XueLian Li, *Post-quantum security of the sum of even-mansour*, Cryptology ePrint Archive, Paper 2025/2274, 2025.

[222] Ehsan Ebrahimi Targhi and Dominique Unruh, *Post-quantum security of the fujisaki-okamoto and oaep transforms*, Theory of Cryptography (Berlin, Heidelberg) (Martin Hirt and Adam Smith, eds.), Springer Berlin Heidelberg, 2016, pp. 192–216.

[223] George Tasopoulos, Charis Dimopoulos, Apostolos P. Fournaris, Raymond K. Zhao, Amin Sakzad, and Ron Steinfeld, *Energy consumption evaluation of post-quantum tls 1.3 for resource-constrained embedded devices*, Proceedings of the 20th ACM International Conference on Computing Frontiers (New York, NY, USA), CF '23, Association for Computing Machinery, 2023, p. 366–374.

[224] M. Tehranipoor et al., *Pqc-sep: Power side-channel evaluation platform for post-quantum cryptography algorithms*, Springer (2024).

[225] Stefano Tessaro and Chenzhi Zhu, *Revisiting bbs signatures*, Advances in Cryptology – EUROCRYPT 2023 (Cham) (Carmit Hazay and Martijn Stam, eds.), Springer Nature Switzerland, 2023, pp. 691–721.

[226] The Tor Project, *Tor protocol specification*, 2026.

[227] TNO, *Guidelines for migrating to post-quantum cryptography*, 2024, Accessed: 24 April 2026.

[228] Callum Turino, William J. Buchanan, Owen Lo, and Christoph Thümmler, *Pqc-leo: An evaluation framework for post-quantum cryptographic algorithms*, 2025 IEEE 7th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA), 2025, pp. 237–247.

[229] Dominique Unruh, *Universally composable quantum multi-party computation*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2010, pp. 486–505.

[230] ———, *Revocable quantum timed-release encryption*, J. ACM **62** (2015), no. 6.

[231] ———, *Computationally binding quantum commitments*, Advances in Cryptology – EUROCRYPT 2016 (Berlin, Heidelberg) (Marc Fischlin and Jean-Sébastien Coron, eds.), Springer Berlin Heidelberg, 2016, pp. 497–527.

[232] Samet Ünsal, *A comparative performance evaluation of kyber, sntrup761, and frodokem for post-quantum cryptography*, 2025.

[233] Mohib Ur Rehman, *Why cryptographic discovery matters for post-quantum security*, may 2026, Accessed: 2026-05-14.

[234] Jelle Vos, Stanislaw Jarecki, and Christopher A. Wood, *Hybrid post-quantum password authenticated key exchange*, `https://datatracker.ietf.org/doc/draft-vos-cfrg-pqpake/`, Accessed: 2026-06-26.

[235] Jelle Vos, Stanislaw Jarecki, Christopher A. Wood, Cathie Yun, Steve Myers, and Yannick Sierra, *A hybrid asymmetric password-authenticated key exchange in the random oracle model*, Cryptology ePrint Archive, Paper 2025/1343, 2025.

[236] Alexander Wiesmaier, Nouri Alnahawi, Tobias Grasmeyer, Julian Geißler, Alexander Zeier, Pia Bauspieß, and Andreas Heinemann, *On pqc migration and crypto-agility*, arXiv preprint arXiv:2106.09599 (2021).

[237] Tim Wood, Keerthi Thomas, Matthew Dean, Swaminathan Kannan, and Robert Learney, *Towards post-quantum verifiable credentials*, Proceedings of the 19th International Conference on Availability, Reliability and Security (New York, NY, USA), ARES '24, Association for Computing Machinery, 2024.

[238] Keita Xagawa, *Signatures with memory-tight security in the quantum random oracle model*, Advances in Cryptology – EUROCRYPT 2024 (Cham) (Marc Joye and Gregor Leander, eds.), Springer Nature Switzerland, 2024, pp. 30–58.

[239] Yufei Xing and Shuguo Li, *A Compact Hardware Implementation of CCA-Secure Key Exchange Mechanism CRYSTALS-Kyber on FPGA*, IACR Trans. Cryptographic Hardware and Embedded Systems (TCHES) **2021** (2021), no. 2, 328–356.

[240] Vijay Kumar Yadav, Nitish Andola, Shekhar Verma, and S. Venkatesan, *A Survey of Oblivious Transfer Protocol*, ACM Computing Surveys **54** (2022), no. 10s, 1–37.

[241] Takashi Yamakawa and Mark Zhandry, *Classical vs quantum random oracles*, Advances in Cryptology – EUROCRYPT 2021 (Cham) (Anne Canteaut and François-Xavier Standaert, eds.), Springer International Publishing, 2021, pp. 568–597.

[242] ———, *Verifiable quantum advantage without structure*, J. ACM **71** (2024), no. 3.

[243] Yu-Guang Yang, Bo Zhou, Guang-Bao Xu, Dong-Huan Jiang, Dan Li, Yi-Hua Zhou, Wei-Min Shi, and Ting-Ting Song, *Quantum oblivious transfer for universal function*, Advanced Quantum Technologies **9** (2026), no. 4, e00796.

[244] Tatu Ylonen and Chris Lonvick, *The secure shell (ssh) authentication protocol*, Tech. Report 4252, IETF, 2006.

[245] ———, *The secure shell (ssh) connection protocol*, Tech. Report 4254, IETF, 2006.

[246] ———, *The secure shell (ssh) transport layer protocol*, Tech. Report 4253, IETF, 2006.

[247] Quan Yuan, Chao Sun, and Tsuyoshi Takagi, *Revisiting the security of fiat-shamir signature schemes under superposition attacks*, Information Security and Privacy (Singapore) (Tianqing Zhu and Yannan Li, eds.), Springer Nature Singapore, 2024, pp. 164–184.

[248] Mark Zhandry, *How to record quantum queries, and applications to quantum indifferentiability*, Advances in Cryptology – CRYPTO 2019 (Cham) (Alexandra Boldyreva and Daniele Micciancio, eds.), Springer International Publishing, 2019, pp. 239–268.

[249] ———, *How to construct quantum random functions*, J. ACM **68** (2021), no. 5.

[250] ———, *A Note on Quantum-Secure PRPs*, Quantum **9** (2025), 1696.

[251] Mengyuan Zhang, Wenling Wu, and Han Sui, *Quantum ind-cpa security notions for aead*, Post-Quantum Cryptography (Cham) (Ruben Niederhagen and Markku-Juhani O. Saarinen, eds.), Springer Nature Switzerland, 2025, pp. 195–230.

[252] ———, *Quantum security of the combined feedback mode*, Chinese Journal of Electronics **34** (2025), no. 6, 1787–1796.

[253] Chuan Zhao, Shengnan Zhao, Minghao Zhao, Zhenxiang Chen, Chong-Zhi Gao, Hongwei Li, and Yu an Tan, *Secure multi-party computation: Theory, practice and applications*, Information Sciences **476** (2019), 357–372.

[254] Biming Zhou, Haodong Jiang, and Yunlei Zhao, *Cpa-secure kems are also sufficient for post-quantum tls 1.3*, Advances in Cryptology – ASIACRYPT 2024 (Singapore) (Kai-Min Chung and Yu Sasaki, eds.), Springer Nature Singapore, 2025, pp. 433–464.

[255] Yaoyun Zhou and Qian Wang, *Hero-sign: Hierarchical tuning and efficient compiler-time gpu optimizations for sphincs+ signature generation*, 2026 IEEE International Symposium on High Performance Computer Architecture (HPCA), 2026, pp. 1–13.