

The Fujisaki–Okamoto Transform in the Context of Quantum Threats*

Alicja SALAMON and Mariusz JURKIEWICZ

Military University of Technology, 2 Gen. S. Kaliski St., Warsaw, Poland

Correspondence should be addressed to: Alicja SALAMON, alicja.salamon@wat.edu.pl

* Presented at the 46th IBIMA International Conference, 26-27 November 2025, Ronda, Spain

Abstract

This work examines the Fujisaki–Okamoto (FO) transform, the standard technique for upgrading a public key encryption scheme secure only against chosen plaintext attacks to one achieving full chosen ciphertext security (CCA2). The FO transform has been central to the design of key encapsulation mechanisms in the NIST Post Quantum Cryptography competition, and its security has been established in the classical Random Oracle Model (ROM). A crucial challenge arises in the post quantum setting, since the ROM differs fundamentally from the Quantum Random Oracle Model (QROM), where adversaries can issue superposition queries. This subtle distinction, often overlooked outside a narrow circle of experts, calls into question the validity of classical proofs when extended to quantum capable adversaries. Our study addresses this gap by analyzing the FO transform explicitly within the QROM framework.

On the practical side, we implement and evaluate two variants of the CRYSTALS-Kyber scheme: the CPA secure encryption scheme and the CCA secure key encapsulation mechanism derived via the FO transform. The experiments demonstrate that while the transform introduces negligible overhead in key generation and encapsulation, it substantially increases the cost of decapsulation due to the additional re-encryption step. These results highlight the trade off between achieving strong CCA2 security and maintaining efficiency. As future work, we aim to explore refinements of CRYSTALS-Kyber that achieve full compatibility with the QROM, thereby contributing to the development of encryption schemes that combine rigorous post quantum security with practical performance.

Keywords: Fujisaki-Okamoto transform, quantum computational model, post-quantum cryptography

Introduction

Modern cryptography is at a critical juncture. For decades, the security of widely deployed public key systems has rested on the presumed intractability of certain number theoretic problems, most prominently integer factorization and discrete logarithms in both finite fields and elliptic curve groups. These hardness assumptions have withstood classical attacks and, as a consequence, have enabled the design of practical, efficient, and trusted cryptographic standards. However, the rapid progress in quantum computation challenges this long-standing foundation. Shor's algorithm demonstrates that a sufficiently powerful quantum computer would solve both integer factorization and discrete logarithms in polynomial time, thereby rendering RSA, Diffie–Hellman, and elliptic curve cryptography insecure.

Although large-scale, fault-tolerant quantum computers have not yet been realized, research in quantum engineering is advancing at an unprecedented pace. Leading technology companies and research consortia are investing heavily in the development of scalable architectures, error correction techniques, and quantum algorithms. This trajectory makes it prudent to anticipate the eventual deployment of quantum devices capable of breaking today’s classical public key infrastructure. The prospect of a “cryptographic transition period” is particularly concerning, since adversaries may already be recording encrypted communications today in anticipation of future decryption once quantum resources become available, an attack scenario commonly referred to as harvest now, decrypt later.

In view of this threat, the research community has turned toward post-quantum cryptography, namely, cryptographic schemes that can withstand adversaries equipped with quantum computation. The U.S. National Institute of Standards and Technology (NIST) has recognized the urgency of this transition, initiating a multi-year standardization process aimed at identifying, analyzing, and eventually adopting quantum-resistant algorithms. The focus has largely been on problems believed to be resistant to quantum speedups, such as those arising from lattices, codes, multivariate polynomials, and hash-based constructions. Among these, lattice-based cryptography has emerged as a particularly strong candidate, offering both efficiency and conjectured hardness in the face of quantum algorithms.

Within this landscape, the question of how to achieve strong security guarantees in the quantum setting remains central. In public key encryption, indistinguishability under chosen ciphertext attack (CCA) is widely accepted as the appropriate security target, since it provides robustness against powerful adversaries who may interact adaptively with decryption oracles. However, many post-quantum proposals achieve security only under chosen plaintext attacks (CPA). To bridge this gap, the Fujisaki–Okamoto (FO) transform provides a generic and well-studied methodology for upgrading CPA secure schemes to CCA security.

A further point about the Fujisaki–Okamoto transform deserves emphasis. The transform is often described as assumption neutral. It is a black box compiler that upgrades chosen plaintext security to chosen ciphertext security while remaining agnostic to the specific hardness assumption that underlies the base encryption scheme. In particular, the construction does not rely on the algebraic details of factorization, discrete logarithms, or lattice problems, and at first glance it appears unrelated to the quantum threat. This intuition, however, is misleading. The subtlety arises not from the computational assumption, but from the proof methodology itself. The transform fundamentally exploits the random oracle model in its classical form. The security proof programs and samples an ideal hash function and reasons about an adversary that submits classical queries to this oracle.

In the presence of a quantum adversary, the abstraction changes in a material way. The quantum random oracle model allows queries in superposition, which alters both what the adversary can learn and how a reduction can control the oracle. Classical proof techniques that are routine for Fujisaki–Okamoto, such as lazy sampling at single points, classical reprogramming of the oracle after the fact, and black box rewinding that assumes classical transcripts, do not directly carry over when the adversary can make coherent oracle queries. It has been established that the classical random oracle model and its quantum analogue differ in their power and in the admissible proof techniques. As a consequence, a proof that establishes security in the classical model does not automatically imply security against quantum adversaries.

The practical implication is clear. Even though the transform is independent of the particular computational assumption and even though its syntax looks unchanged in the quantum era, one must re-examine its security in a model that explicitly accounts for quantum access to the oracle. The literature has made progress by developing quantum-aware proof tools, for example measure-and-reprogram arguments and carefully structured decryption checks that limit how information can leak through the oracle. Yet a completely general and assumption-free theorem that covers all natural instantiations of Fujisaki–Okamoto in the quantum random oracle model remains out of reach. For this reason, any claim that a concrete construction achieves chosen ciphertext security in the presence of quantum adversaries should rest on an analysis within the quantum random oracle model, or on a security reduction that avoids oracle programmability altogether.

The present work contributes to this discussion by examining the Fujisaki–Okamoto transform both theoretically and practically as a tool for strengthening post-quantum encryption schemes. From a theoretical perspective, we investigate its status in the classical random oracle model (ROM) as well as in the quantum random oracle model (QROM), which more accurately captures the capabilities of quantum adversaries. While security in the ROM has been established, the situation in the QROM remains unsettled: a general proof of security is lacking, and only certain special cases have been analyzed. From a practical perspective, we assess the computational impact of applying the FO transform to a representative lattice-based scheme, CRYSTALS-

KYBER, which has been selected as part of the NIST post-quantum standardization effort. Our implementation compares CPA secure and CCA secure variants, with the goal of isolating the concrete efficiency costs attributable to the FO transform. The scheme itself is not the main subject of investigation; rather, it serves as a testbed to explore the trade-off between theoretical rigor and practical efficiency in the design of post-quantum secure systems.

Preliminaries

The LWE and MLWE Problems

One of the fundamental challenges in modern post-quantum cryptography is the **Learning With Errors** problem (LWE). This problem generalizes earlier constructions based on SIS and ISIS and plays a central role in the design of many post-quantum cryptographic schemes.

In its general form, an LWE instance is defined as follows. For a fixed matrix $A \in \mathbb{Z}_q^{n \times m}$, an unknown secret vector $s \in \mathbb{Z}_q^m$, and a random noise vector $e \in \mathbb{Z}^n$, one obtains

$$y \equiv As + e \pmod{q}.$$

The vector e introduces small errors, typically sampled from a discrete Gaussian or a discrete binomial distribution. Although the equation is linear, the added noise makes recovering s computationally intractable. *Computational variant.* In the computational version, the adversary is given a set of independent samples (A_i, y_i) , each satisfying $y_i \equiv A_i s + e_i \pmod{q}$. The goal is to recover the secret vector s . *Decisional variant.* In the decisional version, the task is to distinguish between samples generated according to the above distribution and uniformly random pairs, where y is chosen uniformly at random from \mathbb{Z}_q^n .

Generalization: the MLWE problem. A widely used extension of LWE in practical post-quantum schemes (e.g., KYBER) is the **Module Learning With Errors** problem (MLWE). Here, given a parameter $k \in \mathbb{N}$, a ring R , the problem is to distinguish pairs

$$(\mathbf{A}, \mathbf{t} = \mathbf{A}\mathbf{s} + \mathbf{e}) \in R^{k \times k} \times R^k,$$

where $\mathbf{s}, \mathbf{e} \in R^k$ are secret and error vectors, from uniformly random pairs (\mathbf{A}, \mathbf{t}) sampled from $R^{k \times k} \times R^k$. In this work, we consider the polynomial space $R_q = \mathbb{Z}_q[x]/(x^n + 1)$ where \mathbb{Z}_q is a ring.

The advantage of MLWE is that it retains the hardness of LWE while improving efficiency and compactness of representation, which makes it especially suitable for practical cryptographic schemes.

Reduction from MLWE to LWE. It is important to note that MLWE generalizes LWE and can be reduced to it. Specifically, an MLWE instance can be transformed into a larger LWE instance by expanding the coefficients in R_q into coefficient vectors over \mathbb{Z}_q^n . Consequently, if there exists an efficient algorithm for solving LWE for certain parameters, it can also be applied to solve MLWE with appropriately scaled parameters.

Conversely, the hardness of MLWE inherits from that of LWE. The absence of efficient algorithms for LWE implies that MLWE instances with suitable parameters are also hard. From a practical perspective, this enables the construction of more efficient schemes while preserving the same security foundation.

Random Oracle Model

The Random Oracle Model (ROM) is a widely used abstraction in cryptography, introduced to facilitate security proofs that would be difficult or even intractable to carry out in the standard model. In this framework, a hash function H is idealized as a truly random function, accessible only through queries to an oracle. For each fresh query, the oracle returns a value chosen uniformly at random from its range, while ensuring consistency for repeated queries. The oracle can therefore be regarded as a “black box” mechanism: its internal structure is hidden, but its responses are consistent and appear unpredictable.

The strength of the ROM lies in its ability to enable security reductions under the strong assumption that H behaves as a perfect source of randomness. Proofs in this model typically rely on three fundamental properties of the oracle:

1. **Uniformity:** until a query x is made, the value $H(x)$ is completely undetermined and distributed uniformly at random.
2. **Consistency:** once $H(x)$ has been defined, the same value is always returned for subsequent queries on x .
3. **Programmability:** in reduction based proofs, the simulator can assign values to $H(x)$ provided the resulting distribution remains consistent with uniformity.

These properties make the ROM a versatile and powerful tool for the design and analysis of cryptographic primitives. At the same time, the model has inherent limitations. In practice, the random oracle must be instantiated with a concrete hash function, and no real function can perfectly emulate the behavior of a true random oracle. Consequently, it is not guaranteed that security proven in the ROM always translates to the standard model. Indeed, there exist constructions that are provably secure in the ROM but demonstrably insecure under every instantiation with an actual hash function.

Despite these shortcomings, the ROM has been extremely influential. It is best viewed as a heuristic framework: security in the ROM provides strong evidence that a scheme has no obvious structural weaknesses, even if it does not constitute a proof in the strictest sense. For this reason, the ROM continues to serve as the foundation for the security analysis of many widely deployed protocols, including signature schemes and encryption systems, and remains one of the most important tools in modern cryptographic practice.

Quantum Random Oracle Model

Most modern cryptographic schemes are proven secure in the classical ROM, where the adversary is restricted to queries consisting of classical inputs. Concretely, given access to a random oracle H , the adversary can only evaluate $y = H(x)$ for some classical input $x \in \{0,1\}^n$.

In practice, however, the random oracle is instantiated by a fixed hash function that is known both to the honest participants of the protocol and to the adversary. Within this setting, a quantum adversary gains a decisive advantage: the ability to query the oracle on a superposition of inputs rather than on a single classical string. This capability fundamentally distinguishes the quantum setting from the classical one. Formally, a quantum query is represented by the state

$$\sum_{x \in \{0,1\}^n} a_x |x\rangle|0\rangle,$$

where the amplitudes $a_x \in \mathbb{C}$ satisfy the normalization condition

$$\sum_{x \in \{0,1\}^n} |a_x|^2 = 1.$$

The quantum oracle implements the unitary transformation

$$|x\rangle|y\rangle \mapsto |x\rangle|y \oplus H(x)\rangle,$$

so that its response to the superposition query is

$$\sum_{x \in \{0,1\}^n} a_x |x\rangle|H(x)\rangle.$$

In this way, the adversary obtains a coherent superposition of all pairs $(x, H(x))$, thereby accessing information that would require exponentially many queries in the classical model. It is essential to stress, however, that the cryptographic scheme itself remains entirely classical: both the interaction between protocol participants and the formulation of their queries are restricted to classical bit strings. The quantum capability lies exclusively with the adversary's access to the oracle.

Security Proofs in the QROM

Proof techniques in the QROM often rely on *quantum-accessible pseudorandom functions* (QPRFs). These allow an efficient simulation of a quantum random oracle while maintaining security against a quantum-enabled

adversary. The approach assumes the existence of a QPRF, which can be used by the simulator to construct a reduction that emulates the random oracle in a way that is indistinguishable from real quantum oracle access.

A **quantum-accessible pseudorandom function** is an efficiently computable function satisfying the following condition for all efficient quantum algorithms \mathcal{D} :

$$|\Pr[\mathcal{D}^{\text{QPRF}}(1^n) = 1] - \Pr[\mathcal{D}^{\text{H}}(1^n) = 1]| \leq \text{negl}(n).$$

One method of constructing such functions for use in security proofs is based on **t -wise independent functions**. A function $f: \{0,1\}^n \rightarrow \{0,1\}^n$ is called **t -wise independent** if, for t distinct inputs x_1, x_2, \dots, x_t , the outputs $f(x_1), f(x_2), \dots, f(x_t)$ are mutually independent and uniformly distributed.

The simplest example of such a function is a polynomial of degree $t - 1$:

$$f(x) = a_0 + a_1x + \dots + a_{t-1}x^{t-1},$$

where $a_i \in GF(2^n), i \in \{0, \dots, t - 1\}$. (Zhandry, 2012) showed that $2q$ -wise independent functions are indistinguishable from a truly random oracle by quantum algorithms making at most q queries. Below we recall two key lemmas commonly used in QROM proofs, which will later be applied in the analysis of the modified Fujisaki–Okamoto transform:

Lemma 1 (O2H, *One Way to Hiding Lemma*, see (Unruh, 2014)). *Let $H: \{0,1\}^n \rightarrow \{0,1\}^m$ be a random oracle. Consider an algorithm \mathcal{A} making at most q_H queries to H . Define:*

$$P_{\mathcal{A}}^1 = \Pr \left[\mathbf{b}' = 1: x \xleftarrow{\$} \{0,1\}^n, \mathbf{b}' \leftarrow \mathcal{A}^{\text{H}}(x, H(x)) \right],$$

$$P_{\mathcal{A}}^2 = \Pr \left[\mathbf{b}' = 1: x \xleftarrow{\$} \{0,1\}^n, y \xleftarrow{\$} \{0,1\}^m, \mathbf{b}' \leftarrow \mathcal{A}^{\text{H}}(x, y) \right],$$

where \mathcal{A}^{H} indicates that \mathcal{A} may query H .

Now consider an algorithm \mathcal{B} which, given input x , operates as follows:

- choose $i \xleftarrow{\$} \{1, \dots, q_H\}$,
- choose $y \xleftarrow{\$} \{0,1\}^m$,
- provide x, y to $\mathcal{A}^{\text{H}}(x, y)$,
- just before the i -th query of \mathcal{A} to H , measure the query register and record x' ,
- if \mathcal{A} makes fewer than i queries, output $\perp \notin \{0,1\}^n$,
- define the event that \mathcal{B} wins if $x = x'$, with probability

$$P_{\mathcal{B}} = \Pr \left[x = x': x \xleftarrow{\$} \{0,1\}^n, x' \leftarrow \mathcal{B}(x) \right].$$

Then,

$$|P_{\mathcal{A}}^1 - P_{\mathcal{A}}^2| \leq 2q_H \sqrt{P_{\mathcal{B}}}.$$

This lemma is analogous to the use of the event in the classical ROM, where an adversary queries the oracle at a critical point x . In the quantum setting, the measurement performed by \mathcal{B} plays a similar role, and bounding $P_{\mathcal{B}}$ allows us to control the difference $|P_{\mathcal{A}}^1 - P_{\mathcal{A}}^2|$.

Lemma 2 (Adaptive O2H, *One Way to Hiding Lemma, Adaptive*, see (Unruh, 2014)). *Let $H: \{0,1\}^* \rightarrow \{0,1\}^n$ be a random oracle. Consider an algorithm \mathcal{A}_0 making at most q_0 queries to H and an algorithm \mathcal{A}_1 , which uses the final state of \mathcal{A}_0 and makes at most q_1 queries to H . Define:*

$$P_{\mathcal{A}}^1 = \Pr \left[b' = 1 : m \leftarrow \mathcal{A}_0^H, x \stackrel{\$}{\leftarrow} \{0,1\}^l, b' \leftarrow \mathcal{A}_1^H(x, H(x \parallel m)) \right],$$

$$P_{\mathcal{A}}^2 = \Pr \left[b' = 1 : m \leftarrow \mathcal{A}_0^H, x \stackrel{\$}{\leftarrow} \{0,1\}^l, C \stackrel{\$}{\leftarrow} \{0,1\}^n, b' \leftarrow \mathcal{A}_1^H(x, C) \right].$$

Now define an algorithm \mathcal{B} , which, given (j, C, x) , runs $\mathcal{A}_1^H(x, C)$ and measures the query register just before the j -th query of \mathcal{A}_1 to H , returning $x' \parallel m'$. If \mathcal{A}_1 makes fewer than j queries, \mathcal{B} outputs $\perp \notin \{0,1\}^l$. Define the event $x = x' \wedge m = m'$ with probability

$$P_{\mathcal{B}} = \Pr[x = x' \wedge m = m' : m \leftarrow \mathcal{A}_0^H, x \stackrel{\$}{\leftarrow} \{0,1\}^l, C \stackrel{\$}{\leftarrow} \{0,1\}^n, \\ j \leftarrow \{1, \dots, q_1\}, x' \parallel m' \leftarrow \mathcal{B}^H(j, C, x)],$$

then

$$|P_{\mathcal{A}}^1 - P_{\mathcal{A}}^2| \leq 2q_1 \sqrt{P_{\mathcal{B}}} + q_0 \cdot 2^{-l/2+2}.$$

The Fujisaki–Okamoto Transform in the ROM

Hybrid Encryption Scheme

The Fujisaki–Okamoto transform, introduced in the following subsection, is an example of a **hybrid encryption scheme**. Such schemes combine both asymmetric and symmetric cryptography, thereby leveraging the advantages of each approach. Hybrid encryption techniques are widely used, as they naturally address the inherent limitations of the individual paradigms.

Asymmetric schemes, while providing a high level of security, are computationally expensive, making them inefficient for encrypting long messages. On the other hand, symmetric schemes are significantly faster but require a secure key exchange between the communicating parties. Hybrid schemes mitigate these drawbacks by integrating both methods into a single construction.

The simplest example of a hybrid scheme is a construction that combines an asymmetric encryption scheme \mathcal{B}_{asy} and a symmetric encryption scheme \mathcal{B}_{sym} , where the key spaces satisfy $\mathcal{M}^{\text{asy}} = \mathcal{K}^{\text{sym}}$. The resulting hybrid scheme $\mathcal{B}_{\text{hy}} = (\text{KGen}^{\text{hy}}, \text{Enc}^{\text{hy}}, \text{Dec}^{\text{hy}})$ is defined as follows:

- $\text{KGen}^{\text{hy}}(1^n)$:
 1. $(\text{sk}^{\text{asy}}, \text{pk}^{\text{asy}}) \stackrel{\$}{\leftarrow} \text{KGen}^{\text{asy}}(1^n)$,
 2. $\text{sk} \leftarrow \text{sk}^{\text{asy}}, \text{pk} \leftarrow \text{pk}^{\text{asy}}$,
 3. output the key pair $(\text{sk}, \text{pk}) \leftarrow \text{KGen}^{\text{hy}}(1^n)$.
- $\text{Enc}^{\text{hy}}(\text{pk}, m)$:
 1. $k \stackrel{\$}{\leftarrow} \mathcal{K}^{\text{sym}}$,
 2. $c \stackrel{\$}{\leftarrow} \text{Enc}^{\text{sym}}(k, m)$,
 3. $e \stackrel{\$}{\leftarrow} \text{Enc}^{\text{asy}}(\text{pk}, k)$,
 4. output the ciphertext $(c, e) \leftarrow \text{Enc}^{\text{hy}}(m, \text{pk})$.
- $\text{Dec}^{\text{hy}}(\text{sk}, (c, e))$:
 1. $k' \leftarrow \text{Dec}^{\text{asy}}(\text{sk}, e)$,
 2. $m' \leftarrow \text{Dec}^{\text{sym}}(k, c)$,
 3. output the message $m' \leftarrow \text{Dec}^{\text{hy}}(\text{sk}, (c, e))$.

The commonly accepted security standard in modern public-key encryption is chosen ciphertext security (CCA). In this context, a significant limitation of the above construction becomes apparent. For the hybrid scheme \mathcal{B}_{hy} to achieve CCA security, it is necessary that both \mathcal{B}_{asy} and \mathcal{B}_{sym} are themselves CCA-secure. This requirement substantially restricts the range of feasible constructions.

The Fujisaki–Okamoto Transform

Consider a scenario where we have an asymmetric encryption scheme that is secure against chosen plaintext attacks (CPA), and a symmetric encryption scheme that provides semantic security. Applying the Fujisaki–Okamoto transform enables the construction of a hybrid scheme that achieves chosen ciphertext security (CCA) in the random oracle model. Let the asymmetric encryption scheme be defined as $\mathcal{B}_{asy} = (\text{KGen}^{asy}, \text{Enc}^{asy}, \text{Dec}^{asy})$ over $(\mathcal{M}^{asy}, \mathcal{C}^{asy}, \mathcal{R}^{asy})$, the symmetric encryption scheme as $\mathcal{B}_{sym} = (\text{KGen}^{sym}, \text{Enc}^{sym}, \text{Dec}^{sym})$ over $(\mathcal{K}^{sym}, \mathcal{M}^{sym}, \mathcal{C}^{sym})$, and let there be two hash functions:

- $G: \mathcal{M}^{asy} \rightarrow \mathcal{K}^{sym}$,
- $H: \mathcal{M}^{asy} \times \mathcal{C}^{sym} \rightarrow \mathcal{R}^{asy}$.

We construct the hybrid scheme $\mathcal{B}_{fo} = (\text{KGen}^{fo}, \text{Enc}^{fo}, \text{Dec}^{fo})$ over $(\mathcal{M}^{fo}, \mathcal{C}^{fo})$ as follows:

- $\text{KGen}^{fo}(1^n)$ is the key generation algorithm for producing an asymmetric key pair (pk, sk) , given the security parameter. The algorithm proceeds as follows:
 1. $(sk^{asy}, pk^{asy}) \leftarrow \text{KGen}^{asy}(1^n)$,
 2. $sk \leftarrow sk^{asy}, pk \leftarrow pk^{asy}$,
 3. output the key pair $(sk, pk) \leftarrow \text{KGen}^{fo}(1^n)$.
- $\text{Enc}^{fo}(pk, m)$ is the encryption algorithm, which takes as input a public key pk and a message $m \in \mathcal{M}^{fo} (= \mathcal{M}^{sym})$. The algorithm proceeds as follows:
 1. $\sigma \xrightarrow{\$} \mathcal{M}^{asy}$,
 2. $k \leftarrow G(\sigma)$,
 3. $c \leftarrow \text{Enc}^{sym}(k, m)$,
 4. $h \leftarrow H(\sigma, c)$,
 5. $e \leftarrow \text{Enc}^{asy}(pk, \sigma; h)$,
 6. output the ciphertext $(e, c) \leftarrow \text{Enc}^{fo}(pk, m)$.
- $\text{Dec}^{fo}(sk, (e, c))$ is the decryption algorithm, which takes as input the private key sk and a ciphertext (e, c) . The algorithm proceeds as follows:
 1. Parse (e, c) into e and c . If parsing fails, return the error symbol \perp .
 2. If $e \notin \mathcal{C}^{asy}$ or $c \notin \mathcal{C}^{sym}$, return \perp .
 3. $\hat{\sigma} \leftarrow \text{Dec}^{asy}(sk, e)$,
 4. If $\hat{\sigma} \notin \mathcal{M}^{asy}$, return \perp .
 5. $\hat{k} \leftarrow G(\hat{\sigma})$,
 6. $\hat{h} \leftarrow H(\hat{\sigma}, c)$,
 7. Verify that $e = \text{Enc}^{asy}(pk, \hat{\sigma}; \hat{h})$. Otherwise, return \perp .

8. $y \leftarrow \text{Dec}^{\text{sy}}(\hat{k}, c)$,
9. output the message $y \leftarrow \text{Dec}^{\text{fo}}(\text{sk}, (e, c))$.

We now show that the scheme \mathcal{B}_{fo} satisfies the correctness condition. For a key pair $(\text{sk}, \text{pk}) \leftarrow \text{KGen}^{\text{fo}}(1^n)$, a message $m \in \mathcal{M}^{\text{sym}}$, and a valid ciphertext (e, c) such that it can be parsed into $e \in \mathcal{C}^{\text{asy}}$ and $c \in \mathcal{C}^{\text{sym}}$, we have:

$$c = \text{Enc}^{\text{sym}}(G(\sigma), m) \text{ and } e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma; H(\sigma, c)).$$

Decrypting $\sigma' \leftarrow \text{Dec}^{\text{asy}}(\text{sk}, e)$, by the correctness of the scheme \mathcal{B}_{asy} , we obtain:

$$\sigma' = \text{Dec}^{\text{asy}}(\text{sk}, \text{Enc}^{\text{asy}}(\text{pk}, \sigma; H(\sigma, c))) = \sigma.$$

Consequently, since $\sigma' = \sigma$ implies $G(\sigma') = G(\sigma)$, decrypting $m' \leftarrow \text{Dec}^{\text{sym}}(G(\sigma'), c)$ yields, by the correctness of the scheme \mathcal{B}_{sym} :

$$m' = \text{Dec}^{\text{sym}}(G(\sigma), \text{Enc}^{\text{sym}}(G(\sigma), m)) = m.$$

Security Proof of the Fujisaki–Okamoto Scheme

Let $\mathcal{B}_{\text{asy}} = (\text{KGen}^{\text{asy}}, \text{Enc}^{\text{asy}}, \text{Dec}^{\text{asy}})$ be an asymmetric encryption scheme defined over $(\mathcal{M}^{\text{asy}}, \mathcal{C}^{\text{asy}}, \mathcal{R}^{\text{asy}})$ that is secure in the CPA model. Let \mathcal{A}_{asy} be a PPT adversary attacking \mathcal{B}_{asy} in this model, i.e.,

$$\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) \leq \text{negl}(n).$$

Assume further that $\mathcal{B}_{\text{sym}} = (\text{KGen}^{\text{sym}}, \text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ is a symmetric encryption scheme defined over $(\mathcal{K}^{\text{sym}}, \mathcal{M}^{\text{sym}}, \mathcal{C}^{\text{sym}})$ that satisfies semantic security. Hence the advantage of any adversary \mathcal{A}_{sym} is negligible:

$$\text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) \leq \text{negl}(n).$$

Theorem 1. *Let $\mathcal{E}^{\text{fo}} = (\text{KGen}^{\text{fo}}, \text{Enc}^{\text{fo}}, \text{Dec}^{\text{fo}})$ over $(\mathcal{K}^{\text{fo}}, \mathcal{M}^{\text{fo}}, \mathcal{C}^{\text{fo}})$ be the hybrid scheme obtained by applying the Fujisaki–Okamoto transform to \mathcal{B}_{asy} and \mathcal{B}_{sym} . Under the above assumptions, there exists an adversary \mathcal{A}_{fo} allowed to make q_{H} queries to the random oracles and q_{Dec} queries to the decryption oracle such that its advantage satisfies $\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{E}^{\text{fo}}}^{\text{CCA}}(n) = 2 \cdot q_{\text{H}} \cdot \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) + 2 \cdot q_{\text{Dec}} \cdot 2^{-\gamma}$.*

The original security proof of the Fujisaki–Okamoto hybrid scheme can be found in (Fujisaki & Okamoto, 1999). However, due to several modifications in this construction, we provide a complete proof here, as it differs from the original version.

Proof. The proof proceeds via a standard sequence-of-games argument: we define games G_1, G_2, \dots, G_l , all over the same probability space, and move from one game to the next by changing only how queries to the oracles are answered.

Remarks.

- In each game, the challenger (in G_0) or simulator (in subsequent games) samples a bit $b \in \{0,1\}$, and the adversary outputs a bit $b' \in \{0,1\}$ at the end.
- Variables sampled by the simulator are marked with the superscript “*”. The challenge returned by the simulator to the adversary (i.e., the output of Enc^{fo}) is denoted (e^*, c^*) , where

$$c^* = \text{Enc}^{\text{sym}}(G(\sigma^*), m_b) \quad \text{and} \quad e^* = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*, H(\sigma^*, c^*)).$$

- S_i , for $0 \leq i \leq l$, is the event that in game G_i the equality $b' = b$ holds (i.e., the adversary wins).
- $\text{ASK}_{\sigma_i^*}$, for $0 \leq i \leq l$, is the event in game G_i that either
 1. \mathcal{A}_{fo} makes a query σ^* to the random oracle G , or
 2. \mathcal{A}_{fo} makes a query (σ^*, c^*) to the random oracle H .

Game G_0

Game G_0 is the baseline CCA attack game. Its flow against \mathcal{E}^{fo} is:

1. The challenger generates a key pair $(\text{sk}, \text{pk}) \xleftarrow{\$} \text{KGen}^{\text{fo}}(1^n)$.
2. The adversary receives pk and access to the oracles G, H , and Dec^{fo} .
3. The adversary submits a pair of messages $(m_0, m_1) \in \mathcal{M}^{\text{fo}}$.
4. The challenger samples $b \xleftarrow{\$} \{0,1\}$ and encrypts, producing the challenge $e^* \parallel c^* \leftarrow \text{Enc}^{\text{fo}}(\text{pk}, m_b)$.
5. The adversary receives (e^*, c^*) , regains access to G, H and Dec^{fo} , and finally outputs b' .

By definition of CCA security,

$$\text{Adv}_{\mathcal{A}^{\text{fo}}, \mathcal{E}^{\text{fo}}}^{\text{CCA}}(n) = \left| \Pr \left[\text{Exp}_{\mathcal{E}^{\text{fo}}, \mathcal{A}^{\text{fo}}}^{G_0}(1^n) = 1 \right] - \frac{1}{2} \right| = \left| \Pr[S_0] - \frac{1}{2} \right|.$$

Our goal is to upper bound $\text{Adv}_{\mathcal{A}^{\text{fo}}, \mathcal{E}^{\text{fo}}}^{\text{CCA}}(n)$. In subsequent steps we modify the event S_i and bound the differences $\Pr[S_{i-1}] - \Pr[S_i]$. Once $\Pr[S_i]$ is determined exactly, we obtain the desired upper bound on $\text{Adv}_{\mathcal{A}^{\text{fo}}, \mathcal{E}^{\text{fo}}}^{\text{CCA}}(n)$.

Game G_1

Game G_1 modifies G_0 by changing the decryption oracle: the simulator does not use the secret key sk .

During the simulation, all queries to G and H , made by both the adversary and the simulator, are recorded together with their responses. We maintain the lists

$$\begin{aligned} Q_G &= [(\sigma_i, k_i)], \\ Q_H &= [(\sigma_i, c_i, H_i)]. \end{aligned}$$

The flow of G_1 is as follows:

1. The simulator generates $(\text{sk}, \text{pk}) \xleftarrow{\$} \text{KGen}^{\text{fo}}(1^n)$.
2. The adversary receives pk , access to G and H , and to the *simulated* decryption oracle. Upon a decryption query (e, c) , the simulator checks whether there is an entry (σ, c, h) in Q_H with $\sigma \in \mathcal{M}^{\text{asy}}$ and $e = \text{Enc}^{\text{fo}}(\text{pk}, \sigma; h)$.
 - If not, it returns \perp .
 - Otherwise, it computes $k = G(\sigma)$ and $m' \leftarrow \text{Dec}^{\text{sym}}(k, c)$.

If the pair (e, c) has been queried before, the same output as previously is returned.

3. The adversary submits $(m_0, m_1) \in \mathcal{M}^{\text{fo}}$.
4. The simulator samples $b \xleftarrow{\$} \{0,1\}$ and forms the challenge $c^* \leftarrow \text{Enc}^{\text{fo}}(\text{pk}, m_b)$.
5. The adversary receives c^* and regains access to G, H and the simulated decryption oracle, then outputs b' .

Define BAD as the event that the adversary submits a decryption query $(e_{\text{BAD}}, c_{\text{BAD}})$ for which there is no triple $(\sigma, c_{\text{BAD}}, h)$ in Q_H satisfying the simulation's consistency conditions. In G_0 this is not an issue as long as $(e_{\text{BAD}}, c_{\text{BAD}})$ is a valid ciphertext (since the real decryption algorithm is used), whereas in G_1 the response is \perp .

From the adversary's perspective, outside of BAD there is no difference between interacting with the real scheme in G_0 and the simulation in G_1 . Hence,

$$\begin{aligned}
S_0 \wedge \neg \text{BAD} &= S_1 \wedge \neg \text{BAD}, \\
\Pr[S_0 \wedge \neg \text{BAD}] &= \Pr[S_1 \wedge \neg \text{BAD}], \\
\Pr[S_0] - \Pr[S_1] &\leq \Pr[\text{BAD}].
\end{aligned}$$

Assume $e \in \mathcal{C}^{\text{asy}}$, i.e., there exist $\sigma \in \mathcal{M}^{\text{asy}}$ and $r \in \mathcal{R}^{\text{asy}}$ such that $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma, r)$, and that the simulator's challenge is $e^* \parallel c^*$ with $\sigma^* = \text{Dec}^{\text{asy}}(\text{sk}, e^*)$. Then:

1. If $e = e^*$, necessarily $c \neq c^*$ because $(e^*, c^*) \neq (e, c)$ (the adversary may not submit the challenge for decryption),

or

2. if $e \neq e^*$, then $c \neq c^*$ or $\sigma \neq \sigma^*$; otherwise $H(\sigma, c) = H(\sigma^*, c)$, hence $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma; H(\sigma, c)) = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c^*)) = e^*$, a contradiction.

Therefore every decryption query satisfies

$$(\sigma, c) \neq (\sigma^*, c^*) \Leftrightarrow \sigma \neq \sigma^* \vee c \neq c^*.$$

Since H is modeled as a random oracle, the answer to a query (σ, c) is independent of the answer to (σ^*, c^*) for any $(\sigma, c) \neq (\sigma^*, c^*)$. The adversary attempts to form a valid ciphertext e by guessing a triple (σ, c, h) such that $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma; h)$, without first querying H on (σ, c) . Under the random-oracle assumption, the value $h = H(\sigma, c)$ is uniformly random from the adversary's viewpoint. Hence,

$$\Pr[e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma, h) \wedge h = H(\sigma, c)] \leq \Pr\left[e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma, h): h \xleftarrow{\$} \mathcal{R}^{\text{asy}}\right].$$

For any pk and any $\sigma \in \mathcal{M}^{\text{asy}}$ there are at least 2^γ possible ciphertexts of the form $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma, h)$ as h ranges uniformly. Thus,

$$\begin{aligned}
\Pr\left[e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma, h): h \xleftarrow{\$} \mathcal{R}^{\text{asy}}\right] &\leq \frac{q_{\text{Dec}}}{2^\gamma}, \\
\Pr[\text{BAD}] &\leq q_{\text{Dec}} \cdot 2^{-\gamma}.
\end{aligned}$$

Since $\Pr[S_0] - \Pr[S_1] \leq \Pr[\text{BAD}]$ and $\Pr[\text{BAD}] \leq q_{\text{Dec}} \cdot 2^{-\gamma}$, we conclude that

$$\Pr[S_0] - \Pr[S_1] \leq q_{\text{Dec}} \cdot 2^{-\gamma}.$$

Game G_2

In game G_2 we modify game G_1 by replacing the values $G(\sigma^*)$ and $H(\sigma^*, c^*)$ with independently sampled elements, denoted a^* and r^* , respectively. This change is purely conceptual, since a^* and r^* have the same distributions as $G(\sigma^*)$ and $H(\sigma^*, c^*)$, while $(a^* \neq G(\sigma^*))$ and $(r^* \neq H(\sigma^*, c^*))$.

- The simulator sends the challenge (e^*, c^*) such that $c^* = \text{Enc}^{\text{sym}}(a^*, m_b)$, $e^* = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; r^*)$.
- If the adversary queries the random oracle G at σ^* , the simulator returns a^* . For a query (σ^*, c^*) to H , it returns r^* .
- This modification can be viewed as replacing the oracles G and H with new oracles G' and H' that return exactly the same answers as the originals for all queries except at the inputs σ^* and (σ^*, c^*) .
- The joint distribution of answers under G, H is identical to that under G', H' , hence

$$\Pr[S_1] - \Pr[S_2] = 0.$$

Game G_3

In game G_3 we modify game G_2 by reusing the same challenge that the simulator generated for the adversary, namely

$$(e^*, c^*) \text{ such that } c^* = \text{Enc}^{\text{sym}}(a^*, m_b), \quad e^* = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; r^*),$$

where $r^* \neq H(\sigma^*, c^*)$ and $a^* \neq G(\sigma^*)$. However, the modified oracles G', H' are replaced back by the original oracles G, H . Thus, their answers to the queries σ^* and (σ^*, c^*) are no longer a^*, r^* but rather $G(\sigma^*)$ and $H(\sigma^*, c^*)$.

Note that from the adversary's perspective, the simulations in G_2 and G_3 are indistinguishable, except when the adversary queries the oracles at σ^* or (σ^*, c^*) . Hence, a difference can arise only if the events $\text{ASK}_{\sigma_2^*}$ or $\text{ASK}_{\sigma_3^*}$ occur. We therefore have

$$\begin{aligned} S_2 \wedge \neg \text{ASK}_{\sigma_2^*} &= S_3 \wedge \neg \text{ASK}_{\sigma_3^*}, \\ \text{ASK}_{\sigma_2^*} &= \text{ASK}_{\sigma_3^*}. \end{aligned}$$

Consequently,

$$\Pr[S_2] - \Pr[S_3] \leq \Pr[\text{ASK}_{\sigma_3^*}].$$

Game G_4

In game G_4 we modify game G_3 by changing the way random oracle responses are generated. Instead of selecting the random functions G and H in advance, their outputs are determined lazily, at the time of each query.

- Q_G is the list of G queries and responses. Initially empty, upon a query σ , a value $a \xleftarrow{\$} \mathcal{K}^{\text{sym}}$ is sampled, returned as the response, and recorded in Q_G as the pair (σ, a) .
- Q_H is the list of H queries and responses. Initially empty, upon a query (σ, c) , a value $h \xleftarrow{\$} \mathcal{R}^{\text{asy}}$ is sampled, returned as the response, and recorded in Q_H as the triple (σ, c, h) .
- If the adversary repeats a query that already appears in the list, the previously stored value is returned.

Since this change is purely conceptual, we have

$$\begin{aligned} \Pr[S_4] - \Pr[S_3] &= 0, \\ \Pr[\text{ASK}_{\sigma_4^*}] - \Pr[\text{ASK}_{\sigma_3^*}] &= 0. \end{aligned}$$

Bounding the Advantage of Adversary \mathcal{A}_{f_0}

We obtain

$$\Pr[S_0] - \Pr[S_4] \leq q_{\text{Dec}} \cdot 2^{-\gamma} + \Pr[\text{ASK}_{\sigma_4^*}].$$

We now show that $\Pr[\text{ASK}_{\sigma_4^*}] \leq q_H \cdot \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n)$. The argument proceeds via a reduction: given an efficient adversary \mathcal{A}_{f_0} attacking the CCA security of \mathcal{E}^{f_0} with advantage $\text{Adv}_{\mathcal{A}_{f_0}, \mathcal{B}_{f_0}}^{\text{CCA}}(n)$, we can construct an efficient adversary \mathcal{A}_{asy} that attacks the CPA security of \mathcal{B}_{asy} .

Since for asymmetric encryption schemes semantic security implies CPA security, and semantic security also implies resistance to message recovery attacks with the same advantage, we may assume that \mathcal{A}_{asy} attacks the CPA security of \mathcal{B}_{asy} in the message-recovery game.

Adversary \mathcal{A}_{asy} simulates the interaction of \mathcal{A}_{f_0} in game G_4 , including the random oracles and the decryption oracle. The game proceeds as follows:

1. The challenger \mathcal{S} in the message-recovery game computes $(\text{sk}, \text{pk}) \xleftarrow{\$} \text{KGen}^{\text{asy}}(1^n)$, samples $\sigma^* \xleftarrow{\$} \mathcal{M}^{\text{asy}}$, sets $e^* \xleftarrow{\$} \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*)$, and sends e^* to \mathcal{A}_{asy} .
2. \mathcal{A}_{asy} forwards the public key pk to \mathcal{A}_{f_0} , initiating the simulation of game G_4 .
3. \mathcal{A}_{f_0} submits two messages (m_0, m_1) .
4. \mathcal{A}_{asy} constructs the challenge for \mathcal{A}_{f_0} by sampling $b \xleftarrow{\$} \{0, 1\}$, choosing a one-time key $k^* \xleftarrow{\$} \mathcal{K}^{\text{sym}}$, and computing $c^* \xleftarrow{\$} \text{Enc}^{\text{sym}}(k^*, m_b)$. It then sends the challenge (e^*, c^*) to \mathcal{A}_{f_0} .

5. When \mathcal{A}_{f_0} outputs a bit b' , adversary \mathcal{A}_{asy} selects uniformly at random one of the queries from $Q_G \parallel Q_H$:

$$q \stackrel{\$}{\leftarrow} [Q_G \parallel Q_H],$$

and sends σ' (extracted from q) as its answer to the challenger.

Thus, q is either of the form $(\sigma', G(\sigma'))$ or $(\sigma', c', H(\sigma', c'))$. The value σ' is returned to \mathfrak{S} . Since q was chosen uniformly at random, we have

$$\begin{aligned} \Pr[\sigma^* = \sigma' | \text{ASK}_{\sigma_4^*}] &= \frac{1}{q_H}, \\ \Pr[\sigma^* = \sigma'] &= \frac{1}{q_H} \cdot \Pr[\text{ASK}_{\sigma_4^*}]. \end{aligned}$$

Because \mathcal{B}_{asy} is CPA-secure, it follows that

$$\Pr[\sigma^* = \sigma'] = \frac{1}{q_H} \cdot \Pr[\text{ASK}_{\sigma_4^*}] \leq \text{Adv}_{\mathcal{A}_{asy}, \mathcal{B}_{asy}}^{\text{CPA}}(n),$$

which yields the bound

$$\Pr[\text{ASK}_{\sigma_4^*}] \leq q_H \cdot \text{Adv}_{\mathcal{A}_{asy}, \mathcal{B}_{asy}}^{\text{CPA}}(n).$$

Therefore,

$$\begin{aligned} \Pr[S_0] - \Pr[S_4] &\leq q_{\text{Dec}} \cdot 2^{-\gamma} + \Pr[\text{ASK}_{\sigma_4^*}] \\ &= q_{\text{Dec}} \cdot 2^{-\gamma} + q_H \cdot \text{Adv}_{\mathcal{A}_{asy}, \mathcal{B}_{asy}}^{\text{CPA}}(n). \end{aligned}$$

To bound $\Pr[S_0]$, it remains to estimate $\Pr[S_4]$. We now show that

$$\Pr[S_4] \leq \frac{1}{2} + \text{Adv}_{\mathcal{A}_{sym}, \mathcal{B}_{sym}}^{\text{SEM}}(n).$$

As before, we use adversary \mathcal{A}_{f_0} against the CCA security of \mathcal{E}^{f_0} to construct an adversary \mathcal{A}_{sym} against the semantic security of \mathcal{B}_{sym} . By assumption, \mathcal{A}_{sym} has advantage at most $\text{Adv}_{\mathcal{A}_{sym}, \mathcal{B}_{sym}}^{\text{SEM}}(n)$.

Adversary \mathcal{A}_{sym} simulates \mathcal{A}_{f_0} as in game G_4 , including the random oracles and the decryption oracle. The game proceeds as follows:

1. \mathcal{A}_{sym} generates a key pair $(sk, pk) \leftarrow \text{KGen}^{f_0}$ and sends pk to \mathcal{A}_{f_0} .
2. \mathcal{A}_{f_0} submits $(m_0, m_1) \in \mathcal{M}^{f_0}$. \mathcal{A}_{sym} forwards (m_0, m_1) to its challenger \mathfrak{S} in the semantic security game.
3. \mathfrak{S} generates $k \leftarrow \text{KGen}^{\text{sym}}$, samples $b \stackrel{\$}{\leftarrow} \{0,1\}$, and computes $c^* \leftarrow \text{Enc}^{\text{sym}}(k, m_b)$, which is sent to \mathcal{A}_{sym} .
4. \mathcal{A}_{sym} constructs the challenge for \mathcal{A}_{f_0} as follows:
 - sample $\sigma^* \stackrel{\$}{\leftarrow} \mathcal{M}^{\text{asy}}$,
 - sample $h^* \stackrel{\$}{\leftarrow} \mathcal{R}^{\text{asy}}$,
 - compute $e^* \leftarrow \text{Enc}^{\text{asy}}(pk, \sigma^*; h^*)$,
 - form the challenge (e^*, c^*) and send it to \mathcal{A}_{f_0} .
5. \mathcal{A}_{f_0} returns a bit b' to \mathcal{A}_{sym} , which is then forwarded to \mathfrak{S} .

Adversary \mathcal{A}_{sym} wins iff $\mathbf{b}' = \mathbf{b}$, which occurs precisely when \mathcal{A}_{fo} succeeds. Hence

$$\Pr[\mathbf{b}' = \mathbf{b}] = \frac{1}{2} + \text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) \leq \frac{1}{2} + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n).$$

From the perspective of \mathcal{A}_{fo} , this experiment is identical to game G_4 . Therefore

$$\Pr[S_4] = \frac{1}{2} + \text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) \leq \frac{1}{2} + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n).$$

Recalling that

$$\Pr[S_0] - \Pr[S_4] \leq q_{\text{Dec}} \cdot 2^{-\gamma} + q_{\text{H}} \cdot \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n),$$

we obtain the upper bound

$$\Pr[S_0] \leq q_{\text{Dec}} \cdot 2^{-\gamma} + q_{\text{H}} \cdot \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + \frac{1}{2} + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n).$$

It follows that

$$\begin{aligned} \text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) &= \left| \Pr[S_0] - \frac{1}{2} \right| \\ &\leq q_{\text{Dec}} \cdot 2^{-\gamma} + q_{\text{H}} \cdot \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n). \end{aligned}$$

This completes the proof.

Finally, since

$$\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) \leq \text{negl}(n) \quad \text{and} \quad \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) \leq \text{negl}(n),$$

we conclude that $\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) \leq \text{negl}(n)$. Thus the theorem is proven.

The Fujisaki–Okamoto Transform in the QROM

In the previous section we presented a complete security proof of the Fujisaki–Okamoto (FO) transform in the classical ROM. However, this proof cannot be directly transferred to the quantum setting. We now highlight the key obstacles that arise when considering a quantum adversary:

- In the ROM proof, we constructed a sequence of games. In game 1 (3.3.2), the decryption oracle was simulated by answering the adversary’s queries not through the secret key, but by using lists of queries and answers to the random oracles involved in the scheme. For a quantum adversary, who may query the oracle in superposition over all possible inputs, constructing and maintaining such lists becomes problematic for the simulator.
- In game 2 (3.3.3), we replaced the values $G(\sigma^*)$ and $H(\sigma^*, c^*)$ with random values a^* and r^* of the same distribution. In game 3 (3.3.4), we reintroduced $G(\sigma^*)$ and $H(\sigma^*, c^*)$ as oracle responses, while still using a^* and r^* in the challenge ciphertext. From the adversary’s perspective, the games are indistinguishable—except when the adversary queries the oracle at σ^* or (σ^*, c^*) . In the classical model, the probability of this event was shown to be negligible. In the quantum model, however, the adversary can always include such queries within its superposition, making the argument substantially more difficult.

At present, no general technique is known for overcoming these issues in the FO transform. Following (Targhi & Unruh, 2015), we consider a modified variant of FO that introduces an additional hash function $H': \mathcal{M}^{\text{asy}} \rightarrow \mathcal{M}^{\text{asy}}$.

The encryption algorithm is modified as follows, by adding step 4:

1. $\sigma \xleftarrow{\$} \mathcal{M}^{\text{asy}},$

2. $c \leftarrow \text{Enc}^{\text{sym}}(G(\sigma), m)$,
3. $e \leftarrow \text{Enc}^{\text{asy}}(\text{pk}, \sigma; H(\sigma, c))$,
4. $d \leftarrow H'(\sigma)$,
5. output ciphertext $(e, c, d) \leftarrow \text{Enc}^{\text{hy}}(\text{pk}, m)$.

The decryption algorithm is modified by adding step 5 to verify the value of H' on $\hat{\sigma}$:

1. $\hat{\sigma} \leftarrow \text{Dec}^{\text{asy}}(\text{sk}, e)$,
2. if $\hat{\sigma} \notin \mathcal{M}^{\text{asy}}$, return \perp ,
3. $\hat{h} \leftarrow H(\hat{\sigma}, c)$,
4. if $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; \hat{h})$, return \perp ,
5. if $d \neq H'(\hat{\sigma})$, return \perp ,
6. $\hat{k} \leftarrow G(\hat{\sigma})$,
7. $y \leftarrow \text{Dec}^{\text{sy}}(\hat{k}, c)$,
8. output message $y \leftarrow \text{Dec}^{\text{hy}}(\text{sk}, (e, c, d))$.

This modified construction differs from the standard FO transform only by the introduction of the additional function H' . As a result, the ciphertext is extended by one component, and the algorithms are augmented with the computation of $H'(\sigma)$ during encryption and its verification during decryption. All other components of the scheme remain unchanged. The assumptions on the underlying building blocks, namely \mathcal{B}_{asy} and \mathcal{B}_{sym} , are also preserved.

Security Proof of the Fujisaki–Okamoto Scheme in the QROM

Let $\mathcal{B}_{\text{asy}} = (\text{KGen}^{\text{asy}}, \text{Enc}^{\text{asy}}, \text{Dec}^{\text{asy}})$ be an asymmetric encryption scheme, defined over $(\mathcal{M}^{\text{asy}}, \mathcal{C}^{\text{asy}}, \mathcal{R}^{\text{asy}})$, which is secure in the CPA model. Let \mathcal{A}_{asy} denote an adversary attacking scheme \mathcal{B}_{asy} in this model. This means that the following condition holds:

$$\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) \leq \text{negl}(n).$$

Furthermore, let $\mathcal{B}_{\text{sym}} = (\text{KGen}^{\text{sym}}, \text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ be a symmetric encryption scheme, defined over $(\mathcal{K}^{\text{sym}}, \mathcal{M}^{\text{sym}}, \mathcal{C}^{\text{sym}})$, which satisfies semantic security. Thus, the advantage of an adversary \mathcal{A}_{sym} is negligible:

$$\text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) \leq \text{negl}(n).$$

Theorem 2. *Let $\mathcal{E}^{\text{fo}} = (\text{KGen}^{\text{fo}}, \text{Enc}^{\text{fo}}, \text{Dec}^{\text{fo}})$ over $(\mathcal{K}^{\text{fo}}, \mathcal{M}^{\text{fo}}, \mathcal{C}^{\text{fo}})$ be the hybrid scheme obtained via the Fujisaki–Okamoto transform applied to \mathcal{B}_{asy} and \mathcal{B}_{sym} . Under the above assumptions, there exists a quantum adversary \mathcal{A}_{fo} , making at most $q_{\text{fo}} = q_{\text{G}} + q_{\text{H}} + q_{\text{H}'} + q_{\text{Dec}} + 1$ queries to the random oracles G, H, H' , to the decryption oracle, and to the encryption oracle (the additional +1 accounts for the mandatory encryption query used to obtain the challenge), such that its advantage satisfies:*

$$\begin{aligned} \text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) &\leq O\left(\frac{q_{\text{fo}}^{9/5}}{2^{\gamma/5}}\right) + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) \\ &\quad + 2q_{\text{G} \times \text{H}'} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + 2q_{\text{H}1} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + q_{\text{H}0} \cdot 2^{-l/2+2}}}. \end{aligned}$$

Here $q_{\text{G} \times \text{H}'} = q_{\text{G}} + q_{\text{H}'} + 2q_{\text{Dec}}$, $q_{\text{H}0}$ denotes the number of queries made by \mathcal{A}_{fo} to H before receiving the challenge, and $q_{\text{H}1}$ the number of queries made to H after receiving the challenge.

Proof. The analysis proceeds by stating a few technical assumptions and fixing notation to be used throughout:

- As in the classical random oracle proof, we construct a sequence of games and derive upper bounds on the differences between the adversary's winning probabilities.
- We denote by S_i (for $0 \leq i \leq l$) the event that the adversary succeeds in game G_i , consistent with the classical approach.
- Random variables sampled by the simulator are annotated with index $*$, following the convention used earlier.

Game G_0

Analogously to the classical proof, Game G_0 represents the baseline version of a CCA attack on the scheme Enc^{fo} by the adversary \mathcal{A}_{fo} , who has quantum access to the oracles H, G, H' and classical access to the decryption oracle:

1. The challenger generates a key pair $(sk, pk) \xleftarrow{\$} \text{KGen}^{\text{fo}}(1^n)$.
2. The adversary receives the public key pk and quantum access to the oracles G, H, H' , as well as classical access to Dec^{fo} .
3. The adversary submits two messages: $(m_0, m_1) \in \mathcal{M}^{\text{fo}}$.
4. The challenger samples a bit $b \xleftarrow{\$} \{0,1\}$ and encrypts the message, generating the challenge ciphertext: $(e^*, c^*, d^*) \leftarrow \text{Enc}^{\text{fo}}(pk, m_b)$.
5. The adversary receives (e^*, c^*, d^*) along with renewed access to G, H, H' and Dec^{fo} . Finally, it outputs a bit b' .

We say that the adversary wins if $b = b'$ (event S_0). By definition, we have:

$$\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) = \left| \Pr[S_0] - \frac{1}{2} \right|.$$

In the subsequent steps, we will gradually modify the game G_i and establish upper bounds on the differences between the adversary's success probabilities, i.e., $\Pr[S_{i-1}] - \Pr[S_i]$. Once the final probability $\Pr[S_l]$ is determined, we will obtain an upper bound on $\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n)$ and demonstrate that it is negligible.

Game G_1

In Game G_1 , we introduce a modification with respect to G_0 , consisting of replacing the decryption algorithm Dec^{fo} with Dec^* , which for each query of the form (e, c, d) performs the following steps:

1. If $e = e^*$ it returns \perp .
2. If $\hat{e} \leftarrow \text{Dec}^{\text{asy}}(sk, e) \notin \mathcal{M}^{\text{asy}}$, it sends a query $H'(\sigma^* \oplus 1)$ and returns \perp .
3. If $e \neq \text{Enc}^{\text{asy}}(pk, \hat{e}; H(\hat{e}, c))$, it sends a query $H'(\sigma^* \oplus 1)$ and returns \perp .
4. If $d \neq H'(\hat{e})$ it returns \perp .
5. Otherwise, it outputs $\text{Dec}^{\text{sy}}(G(\hat{e}), c)$.

The most important change from the adversary's perspective is step 1: the new algorithm always returns the error symbol when the asymmetric part of the ciphertext matches that of the challenge. Additionally, in steps 2 and 3, before returning an error, the algorithm queries the oracle H' with input $\sigma^* \oplus 1$. The reason for this operation will become clear later in the proof, during the analysis of the differences between Games G_4 and G_5 .

Our goal is to show that the difference $|\Pr[S_0] - \Pr[S_1]|$ is negligible. To this end, let us analyze situations in which the adversary could distinguish the two games by issuing a decryption query (e, c, d) to the oracle Dec^{fo} in Game G_0 and to the oracle Dec^* in Game G_1 .

Observe that if $e \neq e^*$, both decryption oracles behave identically. Thus, the only cases worth considering are those where $e = e^*$ (since Dec^* always returns \perp in that case):

- $c = c^*, d = d^*$ - such a query cannot occur, since the adversary is not allowed to request decryption of the challenge ciphertext.
- $c = c^*, d \neq d^*$ - here Dec^{fo} also returns \perp .
- $c \neq c^*, d = d^*$ - this implies that $\text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c)) = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c^*))$, in which case the algorithm Dec^{fo} would output the decryption result $\text{Dec}^{\text{fo}}(\text{sk}, (e, c, d))$.

To bound the probability of such a situation, we use the lemma derived in :

Lemma 3. *Let $f: \{0,1\}^{n_1} \rightarrow \{0,1\}^{n_2}$ be a function with min-entropy k , and let $H: \{0,1\}^* \rightarrow \{0,1\}^{n_1}$ be a random oracle. Then any quantum algorithm \mathcal{A} making q queries to H will encounter a collision with probability at most $O\left(\frac{q^{9/5}}{2^{k/5}}\right)$.*

Since in the scheme Enc^{asy} encrypting any message may yield at least 2^γ possible ciphertexts (i.e., the encryption algorithm has min-entropy γ), and the event $\text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c)) = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c^*))$ can be treated as a collision, we obtain:

$$|\Pr[S_0] - \Pr[S_1]| \leq O\left(\frac{q_{\text{fo}}^{9/5}}{2^{\gamma/5}}\right) \leq \text{negl}(n).$$

Games G_2 and G_3

In Game G_2 , we modify Game G_1 by replacing the values $G(\sigma^*)$ and $H'(\sigma^*)$ with random values $a^* \in \mathcal{K}^{\text{sym}}$ and $d^* \in \mathcal{M}^{\text{asy}}$, respectively. The interaction of the adversary \mathcal{A}_{fo} with the simulator proceeds analogously to Game G_1 , with the following differences:

- the simulator uses a^* instead of $G(\sigma^*)$ as the symmetric encryption key, and
- the value d^* as the ciphertext component (e^*, c^*, d^*) instead of $H'(\sigma^*)$.

The randomness used during asymmetric encryption ($H(\sigma^*, c^*)$) remains unchanged.

Using the assumption of semantic security of Enc^{sym} , we can show that

$$\Pr[S_2] = \frac{1}{2} + \text{negl}(n)$$

(see proof in).

Let us now bound the difference $|\Pr[S_1] - \Pr[S_2]|$. In the classical random oracle model, we defined the event BAD and analyzed probability differences via the estimate of $\Pr[\text{BAD}]$. In the quantum random oracle model, however, this approach cannot be applied directly. Instead, we rely on the O2H Lemma ([O2H]).

Assume $\mathcal{A}^{\text{G} \times \text{H}'}$ is a quantum adversary with access to the random oracle $G \times H'$ (where $G \times H'(\sigma) = (G(\sigma), H'(\sigma))$). Upon receiving input (σ^*, a^*, d^*) , it performs the following steps:

- selects a function $H: \mathcal{M}^{\text{asy}} \times \mathcal{C}^{\text{sym}} \rightarrow \mathcal{R}^{\text{asy}}$, generates $(\text{sk}, \text{pk}) \leftarrow \text{KGen}^{\text{asy}}(1^n)$, samples a bit $b \xleftarrow{\$} \{0,1\}$, and gives the public key pk to \mathcal{A}_{fo} ,
- \mathcal{A}_{fo} issues queries to the oracles H, G, H', Dec^* via interaction with $\mathcal{A}^{\text{G} \times \text{H}'}$, and eventually submits a challenge request (m_0, m_1) ,
- $\mathcal{A}^{\text{G} \times \text{H}'}$ computes $c^* \leftarrow \text{Enc}^{\text{sym}}(a^*, m_b)$, $e^* \leftarrow \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*, H(\sigma^*, c^*))$, and sends the challenge (e^*, c^*, d^*) to \mathcal{A}_{fo} ,
- \mathcal{A}_{fo} issues further queries to the oracles H, G, H', Dec^* and finally outputs a bit b' ,

- $\mathcal{A}^{G \times H'}$ outputs whether $b = b'$.

We can estimate that the maximum number of queries $\mathcal{A}^{G \times H'}$ makes to random oracles (on behalf of \mathcal{A}_{f_0}) is $q_{G \times H'} = q_G + q_{H'} + 2q_{Dec}$ (since the algorithm Dec^* makes two oracle calls to $G \times H'$).

Now, define algorithm \mathcal{B} , which - upon receiving σ^* - operates according to the following Game G_3 :

- samples $i \xleftarrow{\$} \{1, \dots, q_{G \times H'}\}$, $a^* \xleftarrow{\$} \mathcal{K}^{sym}$, and $d^* \xleftarrow{\$} \mathcal{M}^{asy}$,
- sends (σ^*, a^*, d^*) to $\mathcal{A}^{G \times H'}$,
- before $\mathcal{A}^{G \times H'}$ makes its i -th query to the oracle $G \times H'$, \mathcal{B} measures the query argument and obtains $\hat{\sigma}$ (if $\mathcal{A}^{G \times H'}$ makes fewer than i queries, \mathcal{B} returns \perp),
- \mathcal{B} wins in G_3 if $\hat{\sigma} = \sigma^*$, in which case it outputs 1.

Table 1 summarizes the behavior of $\mathcal{A}^{G \times H'}$ and \mathcal{B} in Game G_3 .

Notice that for the probabilities defined in the O2H Lemma, we have:

$$P_{\mathcal{A}}^1 = \Pr[1 \leftarrow G_1] = \Pr[S_1],$$

$$P_{\mathcal{A}}^2 = \Pr[1 \leftarrow G_2] = \Pr[S_2],$$

with respect to adversary $\mathcal{A}^{G \times H'}$, and:

$$P_{\mathcal{B}} = \Pr[1 \leftarrow G_3] = \Pr[S_3],$$

with respect to adversary \mathcal{B} . Based on the analysis and the O2H Lemma, we obtain:

$$|\Pr[S_1] - \Pr[S_2]| \leq 2q_{G \times H'} \sqrt{S_3}.$$

It remains to estimate $\Pr[S_3]$, which we address in the next game.

Table 1: Operation of algorithms $\mathcal{A}^{G \times H'}$ and \mathcal{B} in Game G_3

$\mathcal{A}^{G \times H'}(\sigma^*, a^*, d^*)$	$\mathcal{B}(\sigma^*)$:
1. $H: \mathcal{M}^{asy} \times \mathcal{C}^{sym} \rightarrow \mathcal{R}^{asy}$ $(sk, pk) \leftarrow KGen^{asy}(n)$ $b \xleftarrow{\$} \{0, 1\}$	1. $i \xleftarrow{\$} \{1, \dots, q_1\}$, $a^* \xleftarrow{\$} \mathcal{K}^{sym}$ $d^* \xleftarrow{\$} \mathcal{M}^{asy}$
2. $(m_0, m_1) \leftarrow \mathcal{A}_{f_0}^{H, G, H', Dec^*}(pk)$.	2. sends (σ^*, a^*, d^*) to $\mathcal{A}^{G \times H'}$
3. $c^* \leftarrow Enc^{sym}(a^*, m_b)$, $e^* \leftarrow Enc^{asy}(pk, \sigma^*, H(\sigma^*, c^*))$	3. before $\mathcal{A}^{G \times H'}$ makes its i -th query to $G \times H'$, \mathcal{B} measures the query argument obtaining $\hat{\sigma}$
4. $b' \leftarrow \mathcal{A}_{f_0}^{H, G, H', Dec^*}((e^*, c^*, d^*))$	4. if $\mathcal{A}^{G \times H'}$ makes fewer than i queries, \mathcal{B} returns \perp
5. outputs whether $b = b'$	5. outputs whether $\hat{\sigma} = \sigma^*$

Game G_4

In game G_4 , we modify game G_3 by replacing the random oracle H' with a $2(q_{H'} + q_{Dec})$ -wise independent function (as defined in [wiseindependent]). This function is constructed as a random polynomial of degree $2(q_{H'} + q_{Dec}) - 1$ with coefficients in the field $GF(\mathcal{M}^{asy})$. Since a $2(q_{H'} + q_{Dec})$ -wise independent function is

indistinguishable from a truly random oracle from the perspective of a quantum algorithm making at most $q_{H'} + q_{\text{Dec}}$ queries, we obtain:

$$\Pr[S_3] = \Pr[S_4].$$

Game G_5

In game G_5 , we once again modify the decryption algorithm Dec^* , replacing it with the algorithm Dec^{**} , which:

- Has full knowledge of the polynomial used to simulate the oracle H' . In what follows, we will continue to use the notation H' , keeping in mind that in the context of game G_5 it refers to the polynomial rather than an ideal random oracle.
- Upon receiving a query of the form (e, c, d) , performs the following steps:
 1. If $e = e^*$, return \perp .
 2. Construct the set Ω consisting of all roots of the polynomial $H' - d$.
 3. If Ω contains $\hat{\sigma} \neq \sigma^*$ such that $e = \text{Enc}^{\text{asy}}(\text{sk}, \hat{\sigma}; H(\hat{\sigma}, c))$:
 - Compute $H'_{\text{poly}}(\hat{\sigma})$.
 - Compute $\hat{a} \leftarrow G(\hat{\sigma})$.
 - Return $\text{Dec}^{\text{sym}}(\hat{a}, c)$.
 4. If $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$, compute H' for a random value $\sigma \xleftarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and return \perp .
 5. If $H'(\sigma^*) = d$:
 - Compute $\hat{a} \leftarrow G(\sigma^*)$.
 - Return $\text{Dec}^{\text{sym}}(\hat{a}, c)$.
 6. Otherwise, return \perp .

Since game G_5 differs from game G_4 only in the decryption algorithm, it suffices to show that the algorithms Dec^* and Dec^{**} behave identically with respect to any possible query. For clarity, Table 2 compares the steps of algorithms Dec^* and Dec^{**} .

Table 2: Comparison of the behavior of algorithms Dec^* and Dec^{}**

$\text{Dec}^*(e, c, d)$:	$\text{Dec}^{**}(e, c, d)$:
1. If $e = e^*$ return \perp	1. If $e = e^*$, return \perp .
2. If $\hat{\sigma} \leftarrow \text{Dec}^{\text{asy}}(\text{sk}, e) \notin \mathcal{M}^{\text{asy}}$, query $H'(\hat{\sigma} \oplus 1)$ and return \perp	2. Construct the set Ω consisting of all roots of the polynomial $H' - d$.
3. If $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; H(\hat{\sigma}, c))$, query $H'(\hat{\sigma} \oplus 1)$ and return \perp	3. If Ω contains $\hat{\sigma} \neq \sigma^*$ such that $e = \text{Enc}^{\text{asy}}(\text{sk}, \hat{\sigma}; H(\hat{\sigma}, c))$:
4. If $d \neq H'(\hat{\sigma})$, return \perp	<ul style="list-style-type: none"> ○ Query $H'(\hat{\sigma})$ ○ Compute $\hat{a} \leftarrow G(\hat{\sigma})$ ○ Return $\text{Dec}^{\text{sym}}(\hat{a}, c)$
5. Return $\text{Dec}^{\text{sym}}(G(\hat{\sigma}), c)$	4. If $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$, compute H' for a random value $\sigma \xleftarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and return \perp

$\text{Dec}^*(e, c, d):$	$\text{Dec}^{**}(e, c, d):$
	5. If $H'(\sigma^*) = d$: <ul style="list-style-type: none"> ○ Compute $\hat{a} \leftarrow G(\sigma^*)$ ○ Return $\text{Dec}^{\text{sym}}(\hat{a}, c)$ 6. Return \perp

Suppose that the adversary submits a query to the decryption oracle of the form (c, e, d) and let $\hat{\sigma} \leftarrow \text{Dec}^{\text{asy}}(\text{sk}, e)$. We consider the following cases:

1. $\hat{\sigma} = \perp$:
Both algorithms issue a query $H'(\sigma)$ for some $\sigma \neq \sigma^*$ and return \perp .
2. $\hat{\sigma} \neq \perp, \hat{\sigma} \neq \sigma^*, H'(\hat{\sigma}) \neq d$:
Two cases arise:
 - $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; H(\hat{\sigma}, c))$: then Dec^* queries $H'(\sigma^* \oplus 1)$ and returns \perp , while Dec^{**} computes H' for a random value $\sigma \xrightarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and also returns \perp .
 - $e = \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; H(\hat{\sigma}, c))$: then Dec^* returns \perp since $H'(\hat{\sigma}) \neq d$, and Dec^{**} also returns \perp because $\hat{\sigma} \notin \Omega$, which likewise follows from $H'(\hat{\sigma}) \neq d$.
3. $\hat{\sigma} \neq \perp, \hat{\sigma} \neq \sigma^*, H'(\hat{\sigma}) = d$:
Two cases arise:
 - $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; H(\hat{\sigma}, c))$: then Dec^* queries $H'(\sigma^* \oplus 1)$ and returns \perp , while Dec^{**} computes H' for a random value $\sigma \xrightarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and also returns \perp .
 - $e = \text{Enc}^{\text{asy}}(\text{pk}, \hat{\sigma}; H(\hat{\sigma}, c))$: then Dec^* queries $H'(\hat{\sigma})$ and returns $\text{Dec}^{\text{sym}}(G(\hat{\sigma}), c)$, and Dec^{**} does the same.
4. $\hat{\sigma} = \sigma^*, H'(\hat{\sigma}) \neq d$:
Three cases arise:
 - If $e^* = e$: then both algorithms return \perp .
 - If $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$: then Dec^* queries $H'(\sigma^* \oplus 1)$ and returns \perp , while Dec^{**} computes H' for a random value $\sigma \xrightarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and returns \perp .
 - $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$: then Dec^* returns \perp since $H'(\hat{\sigma}) \neq d$, and Dec^{**} also returns \perp because $\hat{\sigma} \notin \Omega$, which likewise follows from $H'(\hat{\sigma}) \neq d$.
5. $\hat{\sigma} = \sigma^*, H'(\hat{\sigma}) = d$:
Three cases arise:
 - If $e^* = e$: then both algorithms return \perp .
 - If $e \neq \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$: then Dec^* queries $H'(\sigma^* \oplus 1)$ and returns \perp , while Dec^{**} computes H' for a random value $\sigma \xrightarrow{\$} \mathcal{M}^{\text{asy}} \setminus \{\sigma^*\}$ and returns \perp .
 - $e = \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; H(\sigma^*, c))$: then both algorithms query $H'(\sigma^*)$ and return $\text{Dec}^{\text{sym}}(G(\sigma^*), c)$.

Based on the above analysis, we can write:

$$\Pr[S_4] = \Pr[S_5].$$

Games G_6 and G_7

The value $H(\sigma^*, c^*)$ has so far served as the randomness seed in the asymmetric encryption algorithm. We now replace it with actual randomness from \mathcal{R}^{asy} . Table 3 presents the operation of algorithms $\mathcal{A}^{G \times H'}$ and \mathcal{B} in game G_6 :

Table 3: Operation of algorithms $\mathcal{A}^{G \times H'}$ and \mathcal{B} in game G_6

$\mathcal{A}^{G \times H'}(\sigma^*, a^*, d^*)$	$\mathcal{B}(\sigma^*)$:
<ol style="list-style-type: none"> 1. $H: \mathcal{M}^{\text{asy}} \times \mathcal{C}^{\text{sym}} \rightarrow \mathcal{R}^{\text{asy}}$ $(\text{sk}, \text{pk}) \leftarrow \text{KGen}^{\text{asy}}(n)$ $b \leftarrow \{0,1\}$ 2. $(m_0, m_1) \leftarrow \mathcal{A}_{f_0}^{(H,G,H',\text{Dec}^{**})}(\text{pk})$. 3. $c^* \leftarrow \text{Enc}^{\text{sym}}(a^*, m_b)$ $e^* \leftarrow \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*)$ 4. $b' \leftarrow \mathcal{A}_{f_0}^{H,G,H',\text{Dec}^{**}}(e^*, c^*, d^*)$ 5. return the result of the comparison $b = b'$ 	<ol style="list-style-type: none"> 1. $i \xleftarrow{\\$} \{1, \dots, q_1\}$ $a^* \xleftarrow{\\$} \mathcal{K}^{\text{sym}}$ $d^* \xleftarrow{\\$} \mathcal{M}^{\text{asy}}$ 2. send (σ^*, a^*, d^*) to $\mathcal{A}^{G \times H'}$ 3. before the i-th query of $\mathcal{A}^{G \times H'}$ to the oracle $G \times H'$, \mathcal{B} performs a measurement of the query argument and obtains $\hat{\sigma}$ 4. if $\mathcal{A}^{G \times H'}$ makes fewer than i queries, \mathcal{B} returns \perp 5. return the result of the comparison: $\hat{\sigma} = \sigma^*$

To estimate the upper bound of the difference $|\Pr[S_5] - \Pr[S_6]|$, we use Lemma O2HA (2).

Assume that adversary \mathcal{A}_{f_0} makes q_{bf} queries to the oracle $G \times H'$ before submitting the encryption query.

We create two adversaries with access to the random oracle H : \mathcal{A}_{H_0} and \mathcal{A}_{H_1} by “splitting” adversary $\mathcal{A}^{G \times H'}$. \mathcal{A}_{H_0} samples $i \xleftarrow{\$} \{1, \dots, q_{G \times H'}\}$ and performs the following:

- $a^* \xleftarrow{\$} \mathcal{K}^{\text{sym}}, d^* \xleftarrow{\$} \mathcal{M}^{\text{asy}},$
- $(m_0, m_1) \leftarrow \mathcal{A}_{f_0}^{(H,G,H',\text{Dec}^{**})}(\text{pk}),$
- $c^* \leftarrow \text{Enc}^{\text{sym}}(a^*, m_b).$

Before the i -th query to the oracle $G \times H'$, \mathcal{A}_{H_0} outputs c^* .

Adversary \mathcal{A}_{H_1} has access to the parameters and final state of \mathcal{A}_{H_0} . Upon receiving (σ^*, h^*) as input, it performs the following:

- If $i > q_{bf}$ (i.e., \mathcal{A}_{f_0} has already submitted the encryption query):
 - $e^* \leftarrow \text{Enc}^{\text{asy}}(\text{pk}, \sigma^*; h^*),$
 - $b' \leftarrow \mathcal{A}_{f_0}^{H,G,H',\text{Dec}^{**}}(e^*, c^*, d^*).$
- It measures the argument of the i -th query to $G \times H'$, obtaining $\hat{\sigma}$.
- Returns the result of the comparison $\hat{\sigma} = \sigma^*$.

Assuming that \mathcal{A}_{H_0} makes q_{H_0} queries to oracle H and \mathcal{A}_{H_1} makes q_{H_1} queries to oracle H , we define algorithm \mathcal{B} , which—after receiving σ^* —operates according to the following game G_7 :

- Sample $j \xleftarrow{\$} \{1, \dots, q_{H1}\}$, $h \xleftarrow{\$} \mathcal{R}^{\text{asy}}$.
- Send (σ^*, h^*) to \mathcal{A}_{H1} .
- Before \mathcal{A}_{H1} issues its j -th query to oracle H , \mathcal{B} measures the query argument and obtains $\hat{\sigma}$ (if \mathcal{A}_{H1} makes fewer than j queries, \mathcal{B} returns \perp).
- \mathcal{B} wins in G_7 if $\hat{\sigma} = \sigma^*$, in which case it returns 1 as the outcome of the game.

Now note that for the probabilities defined in Lemma O2HA we have:

$$P_{\mathcal{A}}^1 = \Pr[1 \leftarrow G_5] = \Pr[S_5],$$

$$P_{\mathcal{A}}^2 = \Pr[1 \leftarrow G_6] = \Pr[S_6],$$

with respect to adversary \mathcal{A}_{H1} , and:

$$P_{\mathcal{B}} = \Pr[1 \leftarrow G_7] = \Pr[S_7],$$

with respect to adversary \mathcal{B} . Based on the analysis and Lemma O2HA, we obtain:

$$|\Pr[S_5] - \Pr[S_6]| \leq 2q_{H1}\sqrt{\Pr[S_7]} + q_{H0} \cdot 2^{-l/2+2}.$$

where l is the bit-length of elements from \mathcal{M}^{asy} .

Using the assumption of CPA-security of \mathcal{B}_{asy} , we can show that

$$\Pr[S_6] \leq \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) \text{ and } \Pr[S_7] \leq \text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n)$$

(proof in (Targhi & Unruh, 2015)).

Bounding the Adversary's Advantage

Based on the established bounds, we obtain

$$\Pr[S_0] \leq O\left(\frac{q_{\text{fo}}^{9/5}}{2^{\gamma/5}}\right) + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n)$$

$$+ 2q_{G \times H'} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + 2q_{H1} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + q_{H0} \cdot 2^{-l/2+2}}.$$

Recalling that

$$\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) = \left| \Pr[S_0] - \frac{1}{2} \right|,$$

we can upper-bound $\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n)$ as

$$\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) \leq O\left(\frac{q_{\text{fo}}^{9/5}}{2^{\gamma/5}}\right) + \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n)$$

$$+ 2q_{G \times H'} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + 2q_{H1} \sqrt{\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) + q_{H0} \cdot 2^{-l/2+2}}.$$

This proves the claimed bound. Since

$$\text{Adv}_{\mathcal{A}_{\text{asy}}, \mathcal{B}_{\text{asy}}}^{\text{CPA}}(n) \leq \text{negl}(n) \quad \text{and} \quad \text{Adv}_{\mathcal{A}_{\text{sym}}, \mathcal{B}_{\text{sym}}}^{\text{SEM}}(n) \leq \text{negl}(n),$$

it follows that $\text{Adv}_{\mathcal{A}_{\text{fo}}, \mathcal{B}_{\text{fo}}}^{\text{CCA}}(n) \leq \text{negl}(n)$. This completes the proof of the theorem.

Evaluating the Fujisaki-Okamoto Transform with CRYSTALS-Kyber

CRYSTALS-Kyber PKE

Before presenting the individual algorithms that define the CRYSTALS-KYBER scheme, it is essential to first introduce the underlying definitions and assumptions on which the construction relies.

- All computations in the scheme are carried out over integers modulo q and polynomials from the ring $R_q = \mathbb{Z}_q[x]/(x^n + 1)$.
- For further considerations, we define the notion of symmetric modular reduction. For an integer n , its symmetric representation modulo q is denoted as

$$n^* = n \bmod^s q,$$

where

$$n^* \in \begin{cases} \left(-\frac{q}{2}, \frac{q}{2}\right] & \text{if } q \text{ is even,} \\ \left(-\frac{q-1}{2}, \frac{q-1}{2}\right] & \text{if } q \text{ is odd.} \end{cases}$$

When the symmetric representation is not required, we apply the standard modular reduction: $n^* = n \bmod^s q$, where $n^* \in [0, q)$.

- The notation $[x]$ denotes the operation of rounding the rational number x to the nearest integer.
- A key part of the scheme is the generation of polynomials with *small* coefficients, which are used in constructing the secret key and the error vectors. For this, we define a centered binomial distribution B_η , where η is a positive integer. Sampling from B_η proceeds as follows:

1. $\{a_i, b_i\} \stackrel{\$}{\leftarrow} \{0, 1\}^2$ for $i \in \{0, \dots, \eta - 1\}$,
2. return $\sum_{i=0}^{\eta-1} (a_i - b_i)$.

Sampling a polynomial v with coefficients from B_η is denoted $v \leftarrow \beta_\eta$. This distribution, similar to a discrete Gaussian centered at zero, enables the efficient construction of small error vectors.

- The Kyber scheme employs compression and decompression functions primarily to reduce the size of transmitted parameters. Additionally, compression is essential in recovering the message during decryption.

The compression function $\text{Compr}_q(x, d)$, where $x \in \mathbb{Z}_q$ and $d < \lceil \log_2 q \rceil$, maps x to an integer in $\{0, \dots, 2^d - 1\}$. The decompression function partially reverses this process, possibly introducing a bounded error:

$$x' = \text{Decompr}_q(\text{Compr}_q(x, d), d),$$

such that

$$|x' - x| \bmod^s q \leq \left\lfloor \frac{q}{2^{d+1}} \right\rfloor.$$

The functions are defined as

$$\begin{aligned} \text{Compr}_q(x, d) &= \left\lfloor \left(\frac{2^d}{q}\right) \cdot x \right\rfloor \bmod 2^d, \\ \text{Decompr}_q(x, d) &= \left\lfloor \left(\frac{q}{2^d}\right) \cdot x \right\rfloor. \end{aligned}$$

When applied to polynomials or vectors of polynomials, these functions are evaluated coefficient-wise.

- The notation $\text{Sample}()$ denotes a deterministic function which, given a binary seed x , outputs y of arbitrary length according to a specified distribution \mathcal{S} , written as $y \leftarrow \mathcal{S} := \text{Sample}(x)$.

We now describe the public-key encryption scheme $\text{KyberCPA} = (\text{KGen}^{\text{KCPA}}, \text{Enc}^{\text{KCPA}}, \text{Dec}^{\text{KCPA}})$. The parameters are: k, d_u, d_t, d_v, η , and the security parameter $n = 256$. The message space is $\mathcal{M} = \{0,1\}^n$. A message $m \leftarrow \mathcal{M}$ is interpreted as a polynomial of degree $n - 1$ with coefficients in $\{0,1\}$.

- $\text{KGen}^{\text{KCPA}}(1^n)$ – key generation algorithm, producing a pair (pk, sk) :
 1. Sample two seeds $\rho, \sigma \leftarrow \{0,1\}^n$,
 2. Generate a matrix $A^{k \times k} \in R_q$ from seed $\rho: A^{k \times k} \leftarrow R_q := \text{Sample}(\rho)$,
 3. Generate two vectors of k polynomials from $\beta_\eta: (s, e) \leftarrow \beta_\eta^k \times \beta_\eta^k := \text{Sample}(\sigma)$,
 4. Compute $t = A \cdot s + e$, then compress with parameter $d_t: t = \text{Compr}_q(A \cdot s + e, d_t)$,
 5. Output keys $(\text{pk}, \text{sk}) = ((t, \rho), s)$.
- $\text{Enc}^{\text{KCPA}}(\text{pk}, m \in \mathcal{M})$ – encryption algorithm, taking public key $\text{pk} = (t, \rho)$ and message $m \in \mathcal{M}$, outputs ciphertext $c = (u, v)$:
 1. Sample randomness $\text{rand} \leftarrow \{0,1\}^n$,
 2. Decompress $t: t' = \text{Decompr}_q(t, d_t)$,
 3. Reconstruct matrix $A \leftarrow R_q := \text{Sample}(\rho)$,
 4. Generate vectors $r, e_1 \in \beta_\eta^k$ and polynomial $e_2 \leftarrow \beta_\eta$ from $\text{rand}: (r, e_1, e_2) \leftarrow \beta_\eta^k \times \beta_\eta^k \times \beta_\eta$
 5. Compute $u = \text{Compr}_q(A^T \cdot r + e_1, d_u)$,
 6. Compute $v = \text{Compr}_q(q, d_v)$,
 7. Output ciphertext $c = (u, v)$.
- $\text{Dec}^{\text{KCPA}}(\text{sk}, c)$ – decryption algorithm, taking secret key $\text{sk} = s$ and ciphertext $c = (u, v)$, outputs recovered message m' :
 1. Decompress $u: u' = \text{Decompr}_q(u, d_u)$,
 2. Decompress $v: v' = \text{Decompr}_q(v, d_v)$,
 3. Compute $m' = \text{Compr}_q(v - s^T \cdot u, 1)$,
 4. Output m' .

The authors of (Joppe, et al., 2017) prove that the above scheme is secure against chosen-plaintext attacks. However, as noted earlier, such a security level is insufficient for most practical applications. Modern cryptographic protocols require security against chosen-ciphertext attacks, which also applies to the Kyber scheme.

Key Encapsulation Mechanism

The CRYSTALS-KYBER scheme, in its CCA-secure variant, is designed in the form of a **Key Encapsulation Mechanism (KEM)**. A KEM scheme consists of three algorithms $\mathcal{B}_{\text{KEM}} = (\text{KGen}, \text{Encaps}, \text{Decaps})$:

1. The key generation algorithm $\text{KGen}(1^n)$ outputs a key pair (pk, sk) .
2. The encapsulation algorithm $\text{Encaps}(\text{pk})$ returns a pair (K, c) , where K denotes the established shared key and c its encapsulation.

3. The decapsulation algorithm $\text{Decaps}(sk, c)$ returns the shared key K or a designated error symbol.

KEM schemes are not intended for direct encryption of messages. Instead, they enable the secure establishment of a shared key, which can then be used in other cryptographic schemes (e.g., symmetric encryption). CRYSTALS-KYBER has been approved as the standardized key encapsulation mechanism for post-quantum cryptography.

CRYSTALS-Kyber KEM

To construct a scheme secure against CCA attacks in the presence of quantum adversaries, the designers of KYBER employed the Fujisaki–Okamoto transformation in the variant introduced in (Targhi & Unruh, 2015) (see also Section 4, where this transformation and its security proof are discussed), together with the key encapsulation mechanism construction described in (Hofheinz, et al., 2017).

The key encapsulation scheme $\text{KyberCCA} = (\text{KGen}^{\text{KCCA}}, \text{Encaps}^{\text{KCCA}}, \text{Decaps}^{\text{KCCA}})$ employs hash functions $G: \{0,1\}^{2n} \rightarrow \{0,1\}^{2n}$, $H: \{0,1\}^n \rightarrow \{0,1\}^n$, and $H': \{0,1\}^{2n} \rightarrow \{0,1\}^{2n}$ (the latter used for deriving the shared key). It operates as follows:

- $\text{KGen}^{\text{KCCA}}(1^n)$ key generation algorithm producing a key pair (pk, sk) :
 1. sample $z \leftarrow \{0,1\}^n$,
 2. invoke the key generation algorithm $\text{KGen}^{\text{KCPA}}(1^n)$ to produce $(pk^{\text{KCPA}}, sk^{\text{KCPA}})$,
 3. output the key pair (pk, sk) , where $pk = pk^{\text{KCPA}}$ and $sk = (sk^{\text{KCPA}}, z)$.
- $\text{Encaps}^{\text{KCCA}}(pk)$ — encapsulation algorithm taking as input a public key $pk = (t, \rho)$ and outputting a ciphertext $c = (u, v)$ together with a key K :
 1. sample $m \leftarrow \{0,1\}^n$,
 2. compute $(\widehat{K}, r) \leftarrow G(H(pk), m)$,
 3. invoke the encryption algorithm Enc^{KCPA} deterministically with randomness r : $c \leftarrow \text{Enc}^{\text{KCPA}}(pk, m; r)$,
 4. compute $K \leftarrow H'(\widehat{K}, H(c))$,
 5. output the ciphertext c and key K .
- $\text{Decaps}^{\text{KCCA}}(sk, pk, c)$ - decapsulation algorithm taking as input a secret key $sk = (s, z)$, a public key $pk = (t, \rho)$, and a ciphertext $c = (u, v)$, returning the key K :
 1. use the decryption algorithm Dec^{KCPA} to recover m : $m' = \text{Dec}^{\text{KCPA}}(sk, c)$,
 2. compute $(\widehat{K}', r') \leftarrow G(H(pk), m')$,
 3. invoke the encryption algorithm Enc^{KCPA} deterministically with randomness r' : $c' \leftarrow \text{Enc}^{\text{KCPA}}(pk, m'; r')$,
 4. if $c = c'$, output $K \leftarrow H'(\widehat{K}', H(c))$,
 5. if $c \neq c'$, output $K \leftarrow H'(z, H(c))$.

Implementation and Experimental Results

This section presents the design and results of the implementation of the KyberCPA and KyberCCA schemes. The objective of this implementation is not to deliver an optimized version of the CRYSTALS-KYBER scheme, but rather to evaluate the effect of applying the Fujisaki–Okamoto (FO) transformation within its construction. In this setting, the role of CRYSTALS-KYBER is that of a representative framework used to assess the transformation, rather than the principal object of study itself.

For this reason, the implementation was carried out in a straightforward manner, without the use of additional optimization techniques, in order to reproduce as faithfully as possible the mathematical structure and logical flow of the KyberCPA and KyberCCA schemes.

The entire implementation was developed in the SageMath environment, an open source platform for symbolic and numerical computations that integrates seamlessly with the Python programming language. This environment was chosen because of its extensive support for algebraic operations, which facilitates a transparent representation of computations on polynomials, matrices, finite fields, and modular structures. Moreover, its high level syntax and flexibility enabled a direct and faithful translation of the theoretical assumptions underlying the schemes into executable code.

It should be emphasized, however, that SageMath **is not optimized for low-level performance**. As an interpreted system running on top of Python, it introduces significant overhead, stemming from limited control over data representation and the absence of advanced optimization techniques, such as automatic avoidance of data copying, SIMD operations, or manual management of buffers and CPU cache.

For instance, in the conducted tests, the encapsulation time was around 1.29 seconds, whereas an analogous Python implementation (kyber-py), run on a modern processor (Intel i7-9750H), achieved times on the order of 2.9 milliseconds. This represents more than a 400-fold slowdown, resulting not from the algorithm itself, but from the characteristics of the environment.

Therefore, it must be stressed that the **performance results presented here are purely auxiliary** and serve only to support the analysis of the correctness of the encryption scheme construction with the applied FO transformation based on Kyber. All performance-related observations should be regarded as indicative only and not binding in the context of real-world, production-grade applications.

Implementation Details

The implementation was carried out in the SageMath environment, with the goal of faithfully reproducing the structure of KyberCPA and KyberCCA without applying additional optimizations. Polynomials are represented as arrays of coefficients, enabling precise control over arithmetic operations and a direct implementation of the NTT transformation.

Matrices and vectors of polynomials are generated from random seeds using SHAKE256 as an extendable output function, ensuring the required number of coefficients. Polynomials with small coefficients are sampled from the distribution B_η , following the scheme's specification.

Hash functions were instantiated as follows: G with SHA3-512, H with SHA3-256, and H' with SHAKE256. Compression and decompression procedures were implemented exactly as defined in the Kyber specification, applied coefficient-wise to polynomials and vectors.

The program accepts as input the modulus $q = 7681$, security parameter $n = 256$, dimension parameter $k = 4$, error parameter $\eta = 3$, compression parameters $d_u = 11$, $d_v = 3$, $d_t = 11$, and primitive roots $\psi = 62$, $\omega = 3844$ used in NTT/INTT computations.

Based on these components, pseudocode was developed for the key generation, encryption, and decryption algorithms (or equivalently, encapsulation and decapsulation in the KEM variant).

Input: q, n, k, dt
Output: $sk = s, pk = (t, rho)$

```

1      rho ← {0,1}n
2      sigma ← {0,1}n
3      A ← Sample(rho, q, n, k)
4      s, e ← Sample(eta, q, n, k)
5      t ← INTT(NTT(A, s))+e
6      compress t
7      return (sk = s, pk = (t, rho))
```

Fig 1. Algorithm KyberCPA.KGen

Input: $pk = (t, \rho), q, n, k, \eta, m, dt, du, dv, \text{ran} = \text{None}$
Output: $c = (u, v)$

```
1      decompress t
2      if ran = None then
3          ran  $\leftarrow \{0,1\}^n$ 
4       $A \leftarrow \text{Sample}(\rho, q, n, k)$ 
5       $r, e_1, e_2 \leftarrow \text{Sample}(\text{ran}, n, k, \eta)$ 
6       $u \leftarrow \text{INTT}(\text{NTT}(A^T, r)) + e_1$ 
7      compress u
8       $v \leftarrow \text{INTT}(\text{NTT}(t^T, r) + e_2 + \text{round}(q/2) * m)$ 
9      compress v
10     return  $c = (u, v)$ 
```

Fig 2. KyberCPA.Enc

Input: $sk = s, c = (u, v), q, n, du, dv$
Output: m

```
1      decompress u
2      decompress v
3       $m \leftarrow v - \text{INTT}(\text{NTT}(s^T, u))$ 
4      compress m
5      return m
```

Fig 3. KyberCPA.Dec

Input: q, n, k, dt
Output: $sk = (sk_{\text{PKE}}, s), pk = pk_{\text{PKE}}$

```
1      generate  $(sk_{\text{PKE}}, pk_{\text{PKE}}) \leftarrow \text{KyberCPA.KGen}(q, n, k, dt)$ 
2      sample  $z \leftarrow \{0,1\}^n$ 
3      return  $(sk = (sk_{\text{PKE}}, s), pk = (pk_{\text{PKE}}))$ 
```

Fig 4. KyberCCA.KGen

Input: $pk = (t, \rho), q, n, k, du, dv, dt, \eta$
Output: $c = (u, v), K$

```
1       $m \leftarrow \{0,1\}^n$ 
2       $K_1 \parallel \text{ran} \leftarrow G(H(pk) \parallel m)$ 
3       $c \leftarrow \text{KyberCPA.Enc}(pk, q, n, k, \eta, m, dt, dt, dv, \text{ran})$ 
4       $K \leftarrow H'(K_1 \parallel H(c))$ 
5      return  $c, K$ 
```

Fig 5. KyberCCA.Encaps

Input: $sk = (z, s)$, $pk = (t, \rho)$, $c = (u, v)$, $q, n, k, du, dt, dv, \eta$
Output: K

```

1      m1 ← KyberCPA.Dec(sk, c, q, n, du, dv)
2      K11 || ran1 ← G(H(pk) || m1)
3      u1, v1 ← KyberCPA.Enc(pk, q, n, k, η, m1, dt, du, dv, ran1)
4      if u1 = u then
5          if v1 = v then
6              K ← H'(K11 || H(c))
7          if v1 ≠ v then
8              K ← H'(z || H(c))
9      if u1 ≠ u then
10         K ← H'(z || H(c))
11     return K

```

Fig 6. KyberCCA.Decaps

Implementation Results

To evaluate the performance impact of applying the Fujisaki–Okamoto transform, 1000 independent measurements were conducted for each operation in KyberCPA (PKE) and KyberCCA (KEM). Tables 4–6 summarize the average execution times, standard deviations, and observed extrema for key generation, encryption/encapsulation, and decryption/decapsulation.

Table 4: Key generation times for KyberCPA (PKE) and KyberCCA (KEM).

Metric	PKE	KEM
Mean [s]	0.9007	0.9032
Std. deviation [s]	0.0191	0.0283
Maximum [s]	0.9715	1.1026
Minimum [s]	0.8755	0.8670

Table 5: Encryption (PKE) vs. encapsulation (KEM).

Metric	Encryption	Encapsulation
Mean [s]	1.2930	1.2957
Std. deviation [s]	0.0219	0.0199
Maximum [s]	1.6457	1.4523
Minimum [s]	1.2565	1.2589

Table 6: Decryption (PKE) vs. decapsulation (KEM).

Metric	Decryption	Decapsulation
Mean [s]	0.3226	1.6179
Std. deviation [s]	0.0077	0.0184
Maximum [s]	0.4074	1.7468
Minimum [s]	0.3110	1.5707

The results indicate that:

- Key generation times remain virtually identical in both schemes, as the FO transform does not alter this phase apart from adding a single random value to the secret key.
- Encapsulation in KyberCCA shows almost no overhead compared to encryption in KyberCPA, confirming that additional hash computations are negligible in practice.
- Decapsulation is the dominant cost: its average runtime is roughly the sum of encryption and decryption in KyberCPA, due to the need for an additional re-encryption step required by the FO transform.

Conclusions

In this work we analyzed the Fujisaki–Okamoto (FO) transform as a mechanism for upgrading public key encryption schemes to chosen ciphertext (CCA) security, with particular emphasis on the setting of quantum adversaries. On the theoretical side, we reviewed ROM and its quantum counterpart (QROM), highlighting the extent to which the security guarantees of the FO transform carry over to both frameworks. On the practical side, we implemented and evaluated two variants of the CRYSTALS–KYBER scheme: the CPA secure version and the CCA secure version obtained through the application of the FO transform.

Our performance evaluation shows that the FO transform introduces only negligible overhead in key generation and encapsulation, but imposes a significant cost in decapsulation. In particular, the average runtime of decapsulation in Kyber–FO is comparable to the combined cost of encryption and decryption in Kyber–CPA, primarily due to the additional encryption step required by the transform. These findings confirm that the FO transform remains a robust and conceptually simple method for strengthening the security of lattice based cryptography, though at the expense of increased computational effort in decapsulation.

As future work, it would be valuable to investigate modifications of CRYSTALS–KYBER that achieve full compatibility with the QROM framework. Such refinements hold promise for the development of quantum resilient schemes that combine rigorous theoretical guarantees with practical efficiency, and they will constitute the focus of our subsequent research.

References

- Ajtai, M., 1996. Generating hard instances of lattice problems (extended abstract). Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing.
- Boneh, D. et al., 2011. Random oracles in a quantum world. International conference on the theory and application of cryptography and information security.
- Boneh, D. & Shoup, V., 2023. A Graduate Course in Applied Cryptography.. s.l.:s.n.
- Fujisaki, E. & Okamoto, T., 1999. Secure Integration of Asymmetric and Symmetric Encryption Schemes. Springer Berlin Heidelberg.
- Guo, F., Susilo, W. & Mu, Y., 2018. Introduction to Security Reduction.. s.l.:s.n.

- Hofheinz, D., Hövelmanns, K. & Kiltz, E., 2017. A Modular Analysis of the Fujisaki-Okamoto Transformation. Cryptology ePrint Archive, Paper 2017/604.
- Joppe, B. et al., 2017. CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. Cryptology ePrint Archive, Paper 2017/634.
- Katz, J. & Lindell, Y., 2015. Introduction to Modern Cryptography.. s.l.:s.n.
- Lippert, J., Blömer, J. & Domik, G., 2014. The Fujisaki-Okamoto Transformation.. s.l.:s.n.
- Lyubashevsky, V., Peikert, C. & Regev, O., 2010. On Ideal Lattices and Learning with Errors over Rings. Springer Berlin Heidelberg.
- National Institute of Standards and Technology, 2023. Post-Quantum Cryptography: Current State and Quantum Mitigation Strategies. s.l.:s.n.
- Regev, O., 2005. On lattices, learning with errors, random linear codes, and cryptography. Association for Computing Machinery.
- Satriawan, A., Mareta, R. & Lee, H., 2024. A Complete Beginner Guide to the Number Theoretic Transform (NTT). Cryptology ePrint Archive, Paper 2024/585.
- Targhi, E. E. & Unruh, D., 2015. Quantum Security of the Fujisaki-Okamoto and OAEP Transforms.. Cryptology ePrint Archive, Paper 2015/1210.
- Unruh, D., 2014. Quantum position verification in the random oracle model.. Cryptology ePrint Archive, Paper 2014/118.
- Zhandry, M., 2012. Secure Identity-Based Encryption in the Quantum Random Oracle Model.. Cryptology ePrint Archive, Paper 2012/076.