Anonymous, Robust Post-Quantum Public Key Encryption

Paul Grubbs[1], Varun Maram[2][0000−0002−1607−9062], and Kenneth G. Paterson[2]

[1] University of Michigan, USA.
[2] Department of Computer Science, ETH Zurich, Switzerland.
paulgrubbs12@gmail.com, vmaram@inf.ethz.ch, kenny.paterson@inf.ethz.ch

Abstract. A core goal of the NIST PQC competition is to produce PKE schemes which, even if attacked with a large-scale quantum computer, maintain the security guarantees needed by applications. The main security focus in the NIST PQC context has been IND-CCA security, but other applications demand that PKE schemes provide *anonymity* (Bellare *et al.*, ASIACRYPT 2001), and *robustness* (Abdalla *et al.*, TCC 2010). Examples of such applications include anonymous cryptocurrencies, searchable encryption, and auction protocols. However, almost nothing is known about how to build post-quantum PKE schemes offering these security properties. In particular, the status of the NIST PQC candidates with respect to anonymity and robustness is unknown.

This paper initiates a systematic study of anonymity and robustness for post-quantum PKE schemes. Firstly, we identify implicit rejection as a crucial design choice shared by most post-quantum KEMs, show that implicit rejection renders prior results on anonymity and robustness for KEM-DEM PKEs inapplicable, and transfer prior results to the implicit-rejection setting where possible. Secondly, since they are widely used to build post-quantum PKEs, we examine how the Fujisaki-Okamoto (FO) transforms (Fujisaki and Okamoto, Journal of Cryptology 2013) confer robustness and enhance weak anonymity of a base PKE.

We then leverage our theoretical results to study the anonymity and robustness of three NIST KEM finalists—Saber, Kyber, and Classic McEliece—and one alternate, FrodoKEM. Overall, our findings for robustness are definitive: we provide positive robustness results for Saber, Kyber, and FrodoKEM, and a negative result for Classic McEliece. Our negative result stems from a striking property of KEM-DEM PKE schemes built with the Classic McEliece KEM: for any message m, we can construct a single hybrid ciphertext c which decrypts to the chosen m under *any* Classic McEliece private key.

Our findings for anonymity are more mixed: we identify barriers to proving anonymity for Saber, Kyber, and Classic McEliece. We also found that in the case of Saber and Kyber, these barriers lead to issues with their IND-CCA security claims. We have worked with the Saber and Kyber teams to fix these issues, but they remain unresolved. On the positive side, we were able to prove anonymity for FrodoKEM and a variant of Saber introduced by D'Anvers *et al.* (AFRICACRYPT 2018). Our analyses of these two schemes also identified technical gaps in their IND-CCA security claims, but we were able to fix them.

1 Introduction

The increasingly real threat of quantum computers breaking all widely-deployed public-key cryptography has driven research in new paradigms for building core public-key primitives like signatures, public-key encryption (PKE), and key encapsulation mechanisms (KEMs) from problems that are computationally intractable even for quantum computers. An umbrella term for this is *Post-Quantum Cryptography* (PQC). The US National Institute of Standards and Technology (NIST) is in the process of selecting new standards which will be used for decades to come. The process has reached its third round with four finalist candidates and five alternate candidates in the KEM/PKE category. The main security target of evaluation for these schemes until now has been IND-CCA security. This was appropriate as a starting point because it suffices for many important use cases. But we argue that the time has now come for a broader study of the candidates' fitness for emerging applications where security properties other than IND-CCA are required.

Two important security properties that go beyond IND-CCA security are *anonymity* (or key privacy) and *robustness*. Anonymity was first formalised in the public key setting by [8]. Roughly, a PKE scheme is anonymous if a ciphertext does not leak anything about which public key was used to create it; strong forms of anonymity equip the adversary with a decryption oracle. Anonymous PKE is a fundamental component of several deployed anonymity systems, most notably anonymous cryptocurrencies like Zcash [10]. It is also important in building anonymous broadcast encryption schemes [6, 29], anonymous credential systems [12] and auction protocols [35]. Robustness for PKE, first formalised in [2], goes hand-in-hand with anonymity. Suppose a party equipped with a private key receives a ciphertext for an anonymous PKE scheme. In the absence of other information, how does a party decide that it is the intended receiver of that ciphertext? The standard approach is to perform trial decryption. Robustness provides an assurance that this process does not go wrong – that the receiver is not fooled into accepting a plaintext intended for someone else. Robustness is also important for maintaining consistency in searchable encryption [1] and ensuring auction bid correctness [35]. Various robustness notions for PKE were studied in [2], while stronger notions were introduced in [16]; the symmetric setting was treated in [17, 21, 15, 28].

To date, there is almost no work that shows how to build anonymous, robust post-quantum PKE schemes. Nor is it known whether the NIST candidates meet these extended notions. The only directly relevant work is by Mohassel [32], who showed a number of foundational results on anonymity and robustness of hybrid PKEs built via the KEM-DEM paradigm ("DEM" being an abbreviation for "data encapsulation mechanism"). Our work is influenced by Mohassel's general approach; however, Mohassel only considers KEMs that are directly constructed from strongly-secure PKEs via sampling a random message from the PKE scheme's message space and then PKE-encrypting it. This makes the results of [32] inapplicable to NIST candidates, for a few reasons. First, the NIST candidates are all KEMs, not PKEs, so there is a basic syntactic mismatch. Sec-

ond, the base PKEs used within the candidate KEMs are only weakly (e.g. OW-CPA) secure, but [32] relies on the starting PKE having (e.g.) IND-CCA security. Finally, [32] only analyzes explicit-rejection KEMs, for which decapsulation can fail, but all the NIST candidates except the alternate candidate HQC [31] are actually implicit-rejection KEMs that never output $\perp$. This means, e.g., the NIST *finalist* KEMs cannot be even weakly robust, while the constructions of [32] all start from robust KEMs.

One of the negative results of [32] is that even if a KEM enjoys a strong anonymity property, the hybrid PKE scheme that results from applying the standard KEM-DEM construction may not be anonymous. This is concerning, since it indicates that if one only focuses on KEMs in the NIST competition, rather than the PKE schemes that will inevitably be built from them using the standard KEM-DEM approach, then there is no guarantee that desired security properties will actually carry over. Thus, one must dig into a KEM's internals if the target is to achieve anonymous hybrid PKE.

In fact, all the NIST candidates in the KEM/PKE category are constructed using variants of the Fujisaki-Okamato (FO) transform [18–20]. The FO transform takes a weakly secure PKE scheme (e.g. one that is OW-CPA or IND-CPA secure) and elevates it to a KEM that is IND-CCA secure. The FO transform and variants of it have recently been heavily analysed, [24, 34, 26, 37, 25], in the Random Oracle Model (ROM) and the Quantum ROM (QROM) [11], but insofar as we are aware, only with a view to establishing IND-CCA security of the resulting KEMs. Only one prior work [23] studies the relationship between FO transforms and anonymity; it shows that the original FO transform enhances anonymity in the ROM. But this result does not tell us whether the modern FO variants used by the NIST finalists also enhance (or even preserve) robustness and anonymity properties; notably, the results of [23] are not in the QROM.

Anonymity and robustness for the KEM-DEM paradigm. Our first main contribution is a modular theory of anonymity and robustness for PKE schemes built via the KEM-DEM paradigm. This extends the work of [32] to general KEMs (instead of those built only from PKEs). An interesting aspect that emerges is a fundamental separation between our results for implicit- and explicit-rejection KEMs. At a high level, KEMs that perform implicit rejection do not in general transfer anonymity and robustness to PKEs obtained via the KEM-DEM paradigm from the KEM component, whilst KEMs that offer explicit rejection, and that also satisfy a mild robustness property, do. Our positive result for explicit rejection KEMs relies on a relatively weak anonymity notion for KEMs which we introduce here, wANO-CCA security. Our negative results for the implicit rejection case are proved through the construction of specific counterexamples and are surprisingly strong. For example, an implicit rejection KEM cannot be robust, but can achieve a strong form of collision freeness (SCFR-CCA, that we define here). This is in some sense the next best thing to robustness. We show that even this property is not sufficient, by exhibiting an implicit rejection KEM that is ANO-CCA, IND-CCA and SCFR-CCA secure, and a DEM that is AE (authenticated encryption) secure and satisfies a strong robustness property

(XROB, from [17]), but where the PKE scheme resulting from composing this KEM and DEM is not ANO-CCA secure.

Anonymity and robustness from FO transforms. Since all the NIST finalists are KEMs of the implicit rejection type and we have a strong negative result there, we must dig deeper if we wish to assure ourselves that anonymity and robustness will be obtained for PKEs built from those KEMs. This introduces our second main contribution, wherein we analyse how the FO transform (and its variants) lift anonymity and robustness properties from a starting weakly-secure PKE scheme, first to the strongly-secure KEM built by the FO transform, and then to the hybrid PKE scheme constructed using the KEM-DEM paradigm.

For explicit-rejection KEMs, we show that for a slight variant of the $\mathrm{HFO}^{\perp}$ transform of [24], the base PKE's weak anonymity and robustness are enhanced to strong (ANO-CCA) anonymity and strong (SROB-CCA) robustness, as long as an intermediate deterministic PKE used in the transform is collision-free. For implicit-rejection KEMs, we show that the $\mathrm{FO}^{\not\perp}$ transform of [24] similarly enhances anonymity and collision-freeness. The culmination of this analysis is showing that KEMs and PKEs built via FO-type transforms can bypass our negative result for implicit rejection KEMs.

Application to NIST candidates. We then apply our above generic analysis for implicit-rejection KEMs to specific schemes related to the NIST PQC competition which employ a transform close to $\mathrm{FO}^{\not\perp}$. In particular, we focus on the NIST finalist Classic McEliece [3], a simplified version of the NIST finalist Saber [7] from [14] that we call "proto-Saber", and the NIST alternate candidate FrodoKEM [4]. The reason we consider proto-Saber instead of the actual Saber scheme is that the IND-CCA security claims made for Saber in its NIST third round specification [7] seem to have been taken from those of proto-Saber in [14] *without modification*. However, the actual technical specification of Saber in [7, Section 8] and the reference implementation of Saber differ from proto-Saber in crucial ways that impact on its formal security analysis. We return to this issue in more detail below and in Section 5.

For Classic McEliece, we show that the hybrid PKE resulting from applying the standard KEM-DEM construction is not strongly robust (in the sense defined in [2]). In fact, we can show that, for any plaintext m, it is possible to construct a single ciphertext c such that c always decrypts to m under *any* Classic McEliece private key. The construction of c does not even need the public key! We stress that this property does not indicate any problem with IND-CCA security of Classic McEliece, but it does expose its limitations as a general-purpose KEM for the broad set of applications that can be envisaged for NIST public key algorithms. Since our $\mathrm{FO}^{\not\perp}$-related results on anonymity of KEMs and PKEs built from them depend on robustness properties, Classic McEliece's limitations in this regard present a barrier to establishing its anonymity using our techniques (but do not preclude a direct proof).

For proto-Saber, the news is better. We provide positive results on anonymity and robustness properties of its KEM and the hybrid PKE schemes derived from

it. Towards these results, we have to adapt our analysis on $\mathrm{FO}^{\not\perp}$ to the actual transform used by proto-Saber. In doing so, we were also able to obtain an explicit proof of IND-CCA security for proto-Saber in the QROM that matches the tightness claimed in [14]. This is relevant because despite claims to the contrary in [14], we find that even the IND-CCA security of proto-Saber cannot be directly proved using any of the known results concerning the $\mathrm{FO}^{\not\perp}$ transform. This is due to low-level details of how proto-Saber applies hash functions to intermediate values in its internal computations. These details are crucial given the delicate nature of QROM proofs and invalidate the direct application of known results on "standard" FO transforms in the QROM.

FrodoKEM uses an FO-type transform that is *identical* to that of proto-Saber. Hence, our positive results on tight IND-CCA security, anonymity and robustness of proto-Saber also apply to FrodoKEM in a similar fashion.

Saber and Kyber [5] both implement the same transform, one which hashes even more intermediate values than proto-Saber does. This creates barriers in applying the proof strategies that we used for proto-Saber when trying to establish anonymity of Saber and Kyber. Interestingly, as we explain in detail, these extra hashes also act as barriers in proving even the IND-CCA security of these two finalists in the QROM with the bounds as claimed in their respective specifications. We consider this an important finding given the centrality of IND-CCA security as the design target in the NIST competition. On a positive note, we show that our robustness analysis of proto-Saber can be extended to Saber and Kyber, which implies that these two NIST finalists lead to strongly robust hybrid PKE schemes. Finally, we suggest small modifications to Saber and Kyber that would bring their FO-type transforms closer to that of proto-Saber and allow us to overcome the aforementioned problems.

Subsequent Work. The NIST finalist NTRU [13] uses altogether a different transform, namely $\mathrm{FO}_m^{\not\perp}$ [24], that differs from $\mathrm{FO}^{\not\perp}$ in a way which makes it difficult to extend our analysis of $\mathrm{FO}^{\not\perp}$ to NTRU. However, in subsequent work to ours, Xagawa [39] has established the anonymity and robustness properties of NTRU by utilizing a stronger property of its base PKE scheme, namely the so-called *strong disjoint-simulatability*.

Paper organisation. Section 2 contains preliminary definitions. Section 3 contains our anonymity and robustness definitions for KEMs, and analysis of generic KEM-DEM composition. Section 4 contains our study of anonymity and robustness enhancement for FO-type transforms, and the security of hybrid PKE built from FO-type KEMs. Section 5 contains our study of the NIST candidate KEMs.

2 Preliminaries

In this section, we briefly define the preliminaries necessary for the main body. We begin with defining the syntax of primitives of interest.

Primitives. A key encapsulation mechanism (KEM) $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ is a tuple of algorithms. The randomized key generation algorithm KGen takes no input and outputs a pair $(\mathsf{pk}, \mathsf{sk})$ of a public encapsulation key pk and a private decapsulation key sk. The randomized encapsulation algorithm Encap takes as input the encapsulation key pk, and outputs a pair (C, k) where C is a ciphertext and k is a bit string. The determinstic decapsulation algorithm Decap takes as input the encapsulation key pk, the decapsulation key sk, and the ciphertext C. If decapsulation can output either a key k or an error symbol $\perp$, we call the KEM an *explicit-rejection* KEM. If decapsulation can only output a key k, we call the KEM an *implicit-rejection* KEM.

A public-key encryption (PKE) scheme $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is a tuple of algorithms. The algorithm KGen is the same as above for KEMs. (It is conventional to call KGen's outputs the encryption/public and decryption/private key, respectively, instead of "encapsulation"/"decapsulation" keys.) The randomized encryption algorithm Enc takes as input the public key pk, and message m, and outputs a ciphertext C. Below, we will sometimes use a modified syntax for encryption, where instead of sampling internal randomness, the algorithm is deterministic and takes random coins as an additional input. Letting r be a string of random bits, we will write $\mathsf{Enc}(\mathsf{pk}, m; r)$ to denote the output of Enc when run with randomness r. Finally, the deterministic decryption algorithm Dec takes as input the public key pk, the secret key sk, and a ciphertext C, and outputs a message m or an error symbol $\perp$.

We assume the reader is familiar with the syntax for authenticated encryption with associated data (AEAD or AE) schemes and message authentication codes (MACs), along with the *correctness* and γ*-spreadness* properties of PKE schemes and KEMs. We provide the corresponding formal definitions in the full version of this paper [22].

Associated to each algorithm that comprises a primitive above is one or more input spaces (e.g. sets of possible keys $\mathcal{K}$ and messages $\mathcal{M}$) and an output space (e.g. the set of possible ciphertexts $\mathcal{C}$). We assume each algorithm checks that each of inputs is in this set, and aborts if not. To reduce notational clutter, we will not make these input/output spaces explicit below, except where necessary.

The KEM-DEM framework. Composing a KEM and a data encapsulation mechanism (DEM) is a standard way to build PKE. Schemes built this way are often called "hybrid" PKE. For completeness, we describe the hybrid PKE built via KEM-DEM composition. Let KEM be a KEM, and DEM be an authenticated encryption scheme. (Below, we will use "DEM" and "AEAD" synonymously.) The hybrid PKE $\mathsf{PKE}^{\mathrm{hy}} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is built as follows. The algorithm $\mathsf{PKE}^{\mathrm{hy}}.\mathsf{KGen}$ is the same as $\mathsf{KEM}.\mathsf{KGen}$. The algorithm $\mathsf{PKE}^{\mathrm{hy}}.\mathsf{Enc}$ takes as input the encapsulation key pk and a message m. It first runs $(C_0, \mathsf{k}) \leftarrow\!\!\$\ \mathsf{KEM}.\mathsf{Encap}(\mathsf{pk})$, the computes $C_1 \leftarrow\!\!\$\ \mathrm{AEAD}.\mathsf{Enc}(\mathsf{k}, m)$ and outputs ciphertext (C_0, C_1). The algorithm $\mathsf{PKE}^{\mathrm{hy}}.\mathsf{Dec}$ first uses sk to decapsulate C_0 and get k or possibly an error symbol $\perp$. Unless decapsulation failed, the algorithm completes by running $\mathrm{AEAD}.\mathsf{Dec}(\mathsf{k}, C_1)$, outputting either m or an error symbol $\perp$.

The Fujisaki-Okamoto transform. Classical results of Fujisaki and Okamoto [18–20] show how to amplify (in the random oracle model, or ROM) the security of public-key encryption, from one-wayness (OW) or indistinguishability (IND) under chosen-plaintext attack (CPA) to indistinguishability under chosen-ciphertext attack (IND-CCA). In this work we will mostly be interested in modern variants of this so-called "FO transform" studied first by Hofheinz et al. [24] in the classical ROM and QROM; extensions in the QROM were then given by [26, 37, 34]. Details of these transforms can be found in Section 4.

2.1 Security Definitions

Next we state several standard security notions which we will use below. In this work we use the "concrete" security paradigm, which explicitly measures the success probability and resource usage of specific adversaries, which we specify using the code-based game-playing framework of Bellare and Rogaway [9]. We will not relate quantities of interest, such as runtime or oracle queries, to a security parameter. We define relevant security notions for PKE (upper box), AEAD and MAC (lower box) in Figure 1.

PKE security notions are given for chosen-ciphertext attacks. All adversaries have access to a decryption oracle D that takes a ciphertext and (where relevant, i.e., in games with *two* key-pairs) a bit that selects which secret key to use. In ANO-CCA and IND-CCA games, the decryption oracle $D_{\not\in}$ disallows queries for the challenge ciphertext. For each PKE notion, the corresponding definition for chosen-plaintext attacks can be obtained by simply removing the decryption oracle. In INT-CTXT, the adversary has an encryption (resp., decryption) oracle that takes associated data and a message (resp., ciphertext); flag win is set to true if the adversary submits a query to its decryption oracle that returns non-$\perp$, but was not returned from an encryption query. In SUF-CMA, the oracle TagO's inputs and outputs are stored in the table $\mathbf{T}$ after each query. In otROR-CCA, the oracles $E_1, \$_1$ are one-time encryption and random-bits oracles, respectively. The many-time security definition ROR-CCA is identical to otROR-CCA, but without this restriction. As for PKE above, CPA variants can be obtained by removing decryption oracles.

For any game G in Figure 1, we define an associated advantage measure for an adversary $\mathcal{A}$ and primitive P, denoted $\mathbf{Adv}^{G}_{P}(\mathcal{A})$, to be either $\Pr\left[G^{\mathcal{A}}_{P} \Rightarrow \mathsf{true}\right]$ or the absolute difference between that quantity and 1/2, if the game G is a bit-guessing game like IND-CCA.

3 Anonymity and Robustness of KEMs

In [32], Mohassel studied the anonymity and robustness of KEMs. However, all of his definitions and results apply only to the special case of KEMs that are constructed from PKE schemes in a restricted way, namely KEMs in which the encapsulation algorithm selects a random message for the PKE scheme and encrypts it using the PKE scheme's encryption algorithm. With this limitation,

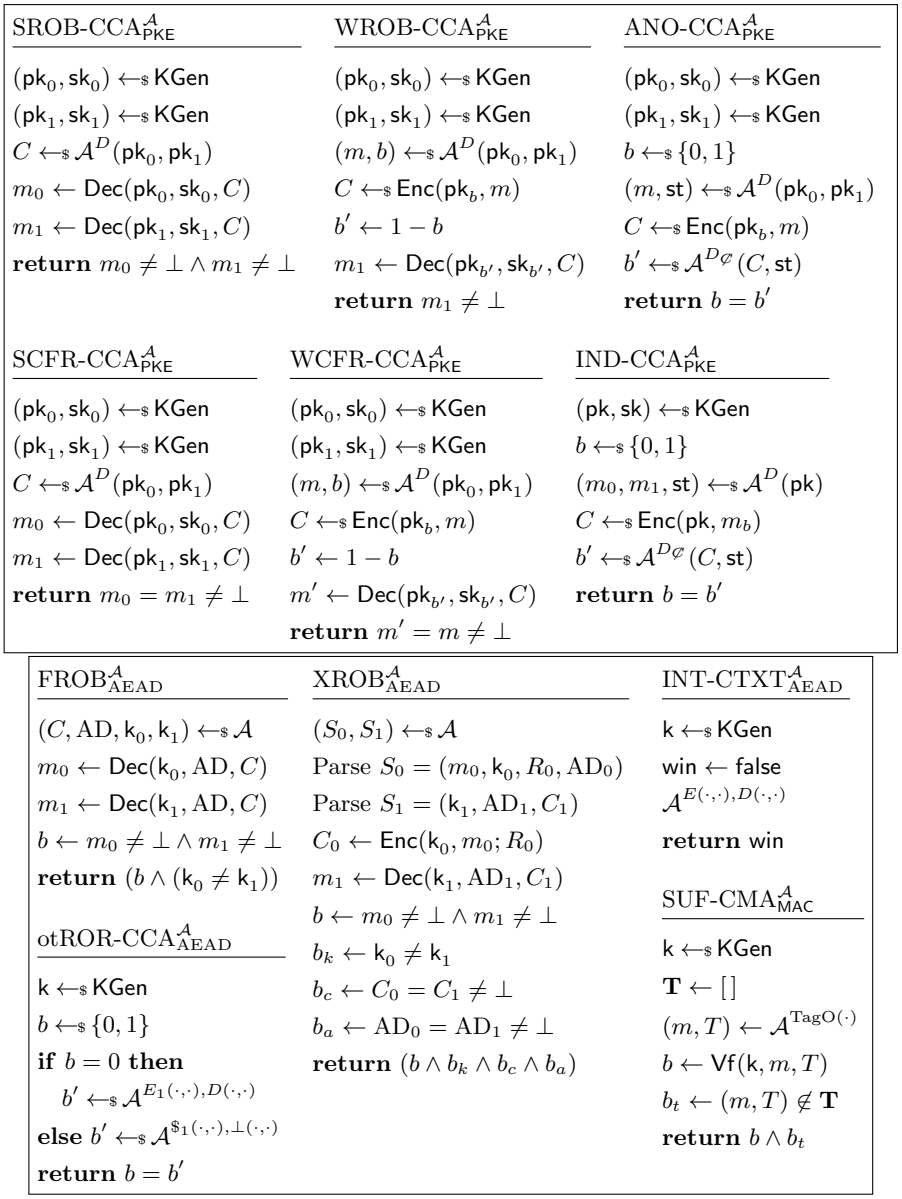

Fig. 1. Security games used in this paper w.r.t. PKE $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ (upper box), and AEAD $\mathrm{AEAD} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ and $\mathsf{MAC} = (\mathsf{KGen}, \mathsf{Tag}, \mathsf{Vf})$ (lower box). In all games associated with PKE above *except* IND-CCA, the decryption oracle D (and $D_{\not\subset}$ in ANO-CCA) also takes as input a *bit* that denotes which secret key (sk_0 or sk_1) to use to decrypt the queried ciphertext. Also, see Section 2.1 for more details.

Mohassel provided a number of interesting results (positive and negative) concerning the anonymity and robustness of KEMs and of PKEs constructed from them via the KEM-DEM framework.

In this section, we bridge the definitional gap left by Mohassel's work by first considering fully general definitions for KEM anonymity and robustness, and then revisiting his results on these properties in the context of the KEM-DEM framework. As we shall see, how much can be recovered depends in a critical way on the KEM's behaviour with respect to rejection of invalid encapsulations.

We first define ANO-CCA security of a KEM $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ via the security game between an adversary and a challenger, as described in Figure 2. Note that the security game differs from the AI-ATK game defined for so-called *general encryption schemes* in [2], where in the latter, an adversary can have access to multiple public-keys (and some corresponding secret keys which will not result in a trivial win for the adversary). Since we are only considering PKE schemes and KEMs in this paper, it is not hard to show that the two security notions are equivalent up to a factor depending on the number of secret key queries an adversary could make (as already discussed in [2]).

An analogous ANO-CPA definition can be obtained simply by removing decapsulation queries in the above game. An adversary $\mathcal{A}$'s advantage in the ANO-{CPA,CCA} game is then defined to be:

$$\mathbf{Adv}^{\mathrm{ANO}-\{\mathrm{CPA,CCA}\}}_{\mathsf{KEM}}(\mathcal{A}) = |\Pr[\mathsf{G}^{\mathcal{A}} = 1] - 1/2|$$

where $\mathsf{G}^{\mathcal{A}}$ refers to $\mathcal{A}$ playing in the appropriate version of the anonymity game,

In the context of KEM-DEM framework for constructing PKE schemes, we will find it sufficient to work with an even weaker notion of anonymity for KEMs, that we refer to as *weak* anonymity. Here, the security game above is modified by giving the adversary only C^* in response to its challenge query, instead of (C^*, k^*); see Figure 2. We then refer to wANO-{CPA,CCA} security and define adversarial advantages as above.

We also define weak robustness (WROB) and strong robustness (SROB) security notions for general KEMs. The security games described in Figure 2 define both notions via two different finalisation steps. Note that the security game for WROB has a subtle difference from the corresponding WROB-ATK game defined for general encryption schemes in [2] (in addition to the fact that, in the latter game, an adversary can have access to multiple public-keys). The difference is that in our notion, an adversary outputs a bit b that determines which of the two public-keys $(\mathsf{pk}_0, \mathsf{pk}_1)$ will be used for encapsulation. This is required because the weak robustness notion is inherently *asymmetric* w.r.t. the two challenge public-keys, since one key is used for encapsulation (resp. encryption in case of PKE schemes) and the other for decapsulation (resp. decryption in case of PKE schemes).

Again, analogous WROB-CPA and SROB-CPA definitions can be obtained simply by removing decapsulation queries in the above games. The advantage of an adversary $\mathcal{A}$ in the {WROB,SROB}-{CPA,CCA} game is then defined as:

$$\mathbf{Adv}^{\{\mathrm{WROB,SROB}\}-\{\mathrm{CPA,CCA}\}}_{\mathsf{KEM}}(\mathcal{A}) = \Pr[\mathsf{G}^{\mathcal{A}} = 1]$$

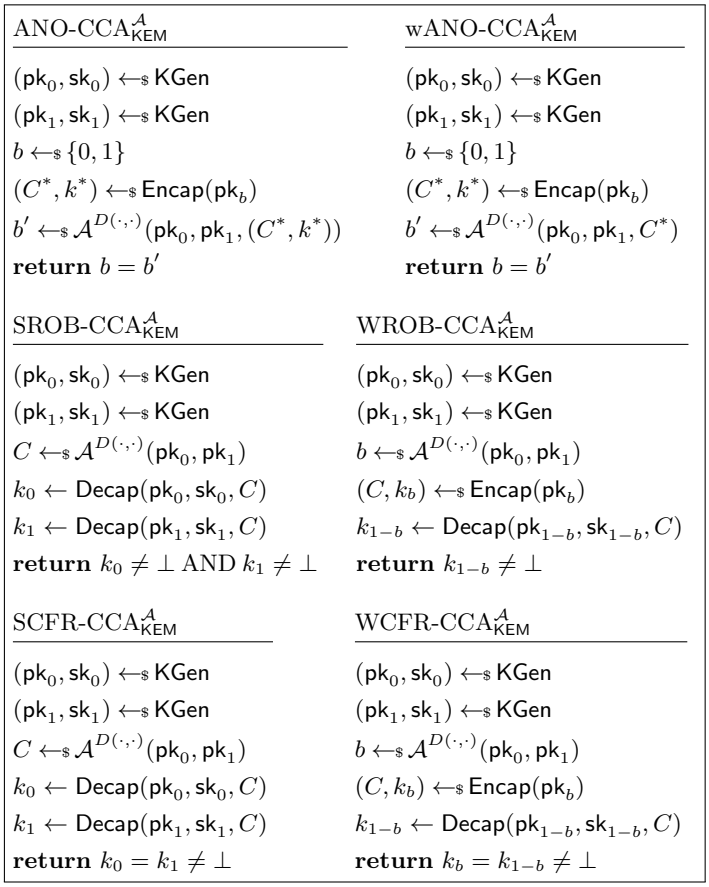

Fig. 2. KEM security notions for chosen-ciphertext attacks. All adversaries have access to a decryption oracle D that takes a ciphertext and (where relevant) a bit that selects which secret key to use. In ANO-CCA and wANO-CCA games, the decryption oracle disallows queries for the challenge ciphertext. For each notion, the corresponding definition for chosen-plaintext attacks can be obtained by simply removing the decryption oracle.

where $\mathsf{G}^{\mathcal{A}}$ refers to $\mathcal{A}$ playing in the appropriate version of the robustness game.

Note that these robustness definitions apply mainly for KEMs that have explicit rejection on decapsulation errors. KEMs that offer only implicit rejection can never satisfy even the WROB-CPA notion.

With these anonymity and robustness notions in hand, it is straightforward to extend the result of [32, Claim 3.3] concerning anonymity preservation from the specific case of KEMs constructed directly from PKEs to fully general KEMs (with a non-zero decapsulation error probability); in fact, we can also show the robustness of hybrid PKE schemes constructed from robust KEMs via the KEM-DEM framework. Namely, we have the following:

Theorem 1. *Let* $\mathsf{PKE}^{hy} = (\mathsf{KGen}, \mathsf{Enc}^{hy}, \mathsf{Dec}^{hy})$ *be a hybrid encryption scheme obtained by composing a KEM* $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ *with a one-time secure authenticated encryption scheme* $\mathsf{DEM} = (\mathsf{Enc}, \mathsf{Dec})$*. If* KEM *is* δ*-correct, then:*

1. *For any* ANO-CCA *adversary* $\mathcal{A}$ *against* PKE^{hy}*, there exist* wANO-CCA *adversary* $\mathcal{B}$*,* IND-CCA *adversary* $\mathcal{C}$ *and* WROB-CPA *adversary* $\mathcal{D}$ *against* KEM*, and* INT-CTXT *adversary* $\mathcal{E}$ *against* DEM *such that*

$$\begin{aligned}\mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) \leq & \mathbf{Adv}^{\text{wANO-CCA}}_{\mathsf{KEM}}(\mathcal{B}) + 2\mathbf{Adv}^{\text{IND-CCA}}_{\mathsf{KEM}}(\mathcal{C}) \\ & + \mathbf{Adv}^{\text{WROB-CPA}}_{\mathsf{KEM}}(\mathcal{D}) + \mathbf{Adv}^{\text{INT-CTXT}}_{\mathsf{DEM}}(\mathcal{E}) + \delta\,.\end{aligned}$$

   The running times of $\mathcal{B}$*,* $\mathcal{C}$ *and* $\mathcal{E}$ *are the same as that of* $\mathcal{A}$*. The running time of* $\mathcal{D}$ *is independent (and less than that) of the running time of* $\mathcal{A}$*.*

2. *For any* WROB-ATK *(resp.* SROB-ATK*) adversary* $\mathcal{A}$ *against* PKE^{hy}*, there exists* WROB-ATK *(resp.* SROB-ATK*) adversary* $\mathcal{B}$ *against* KEM *such that*

$$\begin{aligned}\mathbf{Adv}^{\text{WROB-ATK}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) &\leq \mathbf{Adv}^{\text{WROB-ATK}}_{\mathsf{KEM}}(\mathcal{B})\,, \\ \mathbf{Adv}^{\text{SROB-ATK}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) &\leq \mathbf{Adv}^{\text{SROB-ATK}}_{\mathsf{KEM}}(\mathcal{B})\,,\end{aligned}$$

   where $\text{ATK} \in \{\text{CPA}, \text{CCA}\}$ *and the running time of* $\mathcal{B}$ *is that of* $\mathcal{A}$*.*

Proof (sketch). The proof of Theorem 1.1 closely follows that of [32, Claim 3.3] in terms of the sequence of game-hops. Also for certain game-hops, we rely on security notions that are weaker than the corresponding notions considered in the proof of [32, Claim 3.3] (e.g., WROB-CPA, instead of WROB-CCA, security of the underlying KEM). The complete details of the proof can be found in the full version [22].

To sketch a proof for Theorem 1.2, note that an adversary $\mathcal{A}$ wins the WROB-ATK game w.r.t. PKE^{hy} if it returns a pair (m, b) such that $\mathsf{Dec}^{hy}(sk_{1-b}, C) \neq \perp$ where $C(= (C_{\mathsf{KEM}}, C_{\mathsf{DEM}})) \leftarrow\!\!\$\ \mathsf{Enc}^{hy}(pk_b, m)$. Let $(C_{\mathsf{KEM}}, k_b) \leftarrow\!\!\$\ \mathsf{Encap}(\mathsf{pk}_b)$ and $\mathsf{Decap}(\mathsf{sk}_{1-b}, C_{\mathsf{KEM}}) = k_{1-b}$. It is easy to see that $k_{1-b} \neq \perp$, since $\mathsf{Dec}^{hy}(\mathsf{sk}_{1-b}, C) \neq \perp$. This implies that we can return bit b to win the WROB-ATK game w.r.t. KEM. We can use a similar argument for the SROB-ATK case as well. The complete details can again be found in the full version [22].

Note that Theorem 1 is only meaningful for KEMs with explicit rejection, since for implicit rejection KEMs, the term $\mathbf{Adv}^{\text{WROB-ATK}}_{\mathsf{KEM}}(\cdot)$ in the above security bounds can be large.

3.1 Generic Composition for Implicit Rejection KEMs

Robustness: We first consider what can be said about robustness for PKE schemes built from KEMs offering implicit rejection. We begin with a relaxed notion of robustness, namely *collision freeness* (as introduced for the specific case of KEMs obtained from PKEs in [32]). Informally, a scheme is said to be collision-free if a ciphertext always decrypts to two *different* messages under two different secret keys. We consider two variants, weak (WCFR) and strong collision freeness (SCFR). The security games defined in Figure 2 define both notions via two different finalisation steps.

As usual, analogous WCFR-CPA and SCFR-CPA definitions can be obtained by removing decapsulation queries in the above games. Adversary A's advantage in the {WCFR,SCFR}-{CPA,CCA} game is defined to be:

$$\mathbf{Adv}^{\{\text{WCFR,SCFR}\}-\{\text{CPA,CCA}\}}_{\mathsf{KEM}}(\mathcal{A}) := \Pr[\mathsf{G}^{\mathcal{A}} = 1]$$

where G^A refers to A playing in the appropriate version of the CFR game.

Now suppose we have a KEM that is SCFR-CCA (resp. WCFR-CCA) secure and a DEM that is FROB (resp. XROB) secure. (Recall that FROB and XROB are robustness notions for symmetric encryption schemes introduced in [17] and defined in Figure 1.) Then we can show that the hybrid PKE scheme obtained by composing these KEM and DEM schemes is SROB-CCA (resp. WROB-CCA) secure. More formally,

Theorem 2. *Let* $\mathsf{PKE}^{hy} = (\mathsf{KGen}, \mathsf{Enc}^{hy}, \mathsf{Dec}^{hy})$ *be a hybrid encryption scheme obtained by composing a KEM* $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ *with a DEM* $\mathsf{DEM} = (\mathsf{Enc}, \mathsf{Dec})$*. Then for any* SROB-CCA *(resp.* WROB-CCA*) adversary* $\mathcal{A}$ *against* PKE^{hy}*, there exist* SCFR-CCA *(resp.* WCFR-CCA*) adversary* $\mathcal{B}$ *againt* KEM *and* FROB *(resp.* XROB*) adversary* $\mathcal{C}$ *against* DEM *such that*

$$\mathbf{Adv}^{\text{SROB-CCA}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{SCFR-CCA}}_{\mathsf{KEM}}(\mathcal{B}) + \mathbf{Adv}^{\text{FROB}}_{\mathsf{DEM}}(\mathcal{C}),$$
$$\mathbf{Adv}^{\text{WROB-CCA}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{WCFR-CCA}}_{\mathsf{KEM}}(\mathcal{B}) + \mathbf{Adv}^{\text{XROB}}_{\mathsf{DEM}}(\mathcal{C}),$$

where the running times of $\mathcal{B}$ *and* $\mathcal{C}$ *are the same as that of* $\mathcal{A}$*.*

Proof (sketch). Note that an adversary $\mathcal{A}$ wins the SROB-CCA game w.r.t. PKE^{hy} if it returns a ciphertext C $(= (C_{\mathsf{KEM}}, C_{\mathsf{DEM}}))$ such that $\mathsf{Dec}^{hy}(sk_0, C) \neq \perp$ and $\mathsf{Dec}^{hy}(sk_1, C) \neq \perp$. Let $\mathsf{Decap}(\mathsf{sk}_0, C_{\mathsf{KEM}}) = k_0$ and $\mathsf{Decap}(\mathsf{sk}_1, C_{\mathsf{KEM}}) = k_1$. It is easy to see that $k_0 \neq \perp$ and $k_1 \neq \perp$. Now if $k_0 = k_1$, we can return C_{KEM} to win the SCFR-CCA game w.r.t. KEM. If $k_0 \neq k_1$, we can return $(C_{\mathsf{DEM}}, k_0, k_1)$ to win the FROB game w.r.t. DEM. We can do a similar case-distinction to argue about WROB-CCA security as well. The complete details of the proof can be found in the full version [22].

Note that Farshim et al. [17] provide efficient constructions of FROB- and XROB-secure AE schemes, meaning that the requirements for the above theorem can be easily met. At the same time, they showed that a symmetric AE

scheme that achieves the standard ROR-CCA notion of security is also inherently robust, albeit w.r.t. some weaker notions compared to FROB. Namely, such ROR-CCA secure AE schemes were shown to satisfy the so-called *semi-full robustness* (SFROB) notion in [17]. The SFROB notion of robustness for symmetric AE schemes is a (potentially) weaker variant of FROB where, in the corresponding security game, the adversary does not get to choose any keys. Instead, two keys are honestly generated and the adversary is given oracle access to encryption and decryption algorithms under both keys. The adversary is also given access to one of the keys, and the game is won (similar to that of FROB) if the adversary returns a ciphertext that decrypts correctly under both honestly generated keys.

The following theorem shows that a DEM that is only ROR-CCA secure – and that lacks the stronger robustness properties from [17] – is incapable of *generically* transforming strongly collision-free implicit rejection KEMs to strongly robust hybrid PKEs.

Theorem 3. *Suppose there exists a KEM that is simultaneously* SCFR-CCA*,* IND-CCA *and* ANO-CCA *secure. Suppose that there exists a* SUF-CMA*-secure MAC scheme and an* ROR-CPA *secure symmetric encryption scheme (such schemes can be built assuming only the existence of one-way functions). Suppose also that collision-resistant hash functions exist. Then there exists an implicit-rejection KEM that is* SCFR-CCA*,* IND-CCA *and* ANO-CCA *secure and a DEM that is* ROR-CCA *secure, such that the hybrid PKE scheme obtained from their composition is not* SROB-CCA *secure.*

Proof (sketch). Let $\mathsf{MAC} = (\mathsf{Tag}, \mathsf{Vf})$ be an SUF-CMA secure MAC. We construct $\overline{\mathsf{MAC}} = (\overline{\mathsf{Tag}}, \overline{\mathsf{Vf}})$ where the only difference from MAC is that we fix a "faulty" key $\overline{k}$ chosen uniformly at random from the original MAC key-space such that $\overline{\mathsf{Vf}}(\overline{k}, \cdot) = 1$. Note that $\overline{\mathsf{MAC}}$ is also SUF-CMA secure. So by composing $\overline{\mathsf{MAC}}$ with an ROR-CPA secure symmetric encryption scheme SE that *never* rejects invalid ciphertexts via the "Encrypt-then-MAC" construction, we get an AE-secure $\overline{\mathsf{DEM}}$. Now let $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ be a KEM that is SCFR-CCA, IND-CCA and ANO-CCA secure, and H be a collision-resistant hash function with its range being the key-space of SE. We construct $\overline{\mathsf{KEM}} = (\mathsf{KGen}, \mathsf{Encap}, \overline{\mathsf{Decap}})$ where the only difference from KEM is that the ciphertext space is augmented by a "special" bitstring $\overline{c}$ such that $\overline{\mathsf{Decap}}(\mathsf{sk}, \overline{c}) = H(\mathsf{pk})||\overline{k}$, for any KEM key-pair $(\mathsf{pk}, \mathsf{sk})$. It is not hard to see that $\overline{\mathsf{KEM}}$ is also IND-CCA, ANO-CCA secure, and SCFR-CCA secure (relying on the collision-resistance of H). Now the composition of $\overline{\mathsf{KEM}}$ and $\overline{\mathsf{DEM}}$ will not result in an SROB-CCA secure hybrid PKE. Specifically, an adversary can return the ciphertext $(\overline{c}, c'||\sigma')$, where $c'||\sigma'$ is an arbitrary $\overline{\mathsf{DEM}}$ ciphertext, to win the corresponding SROB-CCA game with probability 1. Complete details of the proof can be found in the full version [22].

Anonymity: Now we turn to the question of what can be said about anonymity for PKE schemes built from KEMs offering implicit rejection. We prove a negative result that strengthens an analogous result of [32]. That result showed that

there exist KEMs that are ANO-CCA (and IND-CCA) secure and XROB-secure authenticated encryption schemes, such that the hybrid PKE scheme resulting from their composition is *not* ANO-CCA secure. Thus anonymity is not preserved in the hybrid construction. However the KEM construction that was used to show this negative result in [32] is not SCFR-CCA secure, which might lead one to think that the strong collision freeness of implicit rejection KEMs might be sufficient to preserve anonymity. Here, we show this not to be true.

Theorem 4. *Suppose there exists a KEM that is simultaneously* SROB-CCA, IND-CCA *and* ANO-CCA *secure, a claw-free pair of permutations with domain and range being the encapsulated key-space of the KEM, and a collision-resistant hash function. Suppose also that there exists a DEM that is* ROR-CCA *and* XROB*-secure. Then there exists an implicit-rejection KEM that is* SCFR-CCA, IND-CCA *and* ANO-CCA *secure and a DEM that is* ROR-CCA *and* XROB*-secure, such that the resulting hybrid PKE is not* ANO-CCA *secure.*

Proof (sketch). Let $\mathsf{KEM} = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ be a KEM that is IND-CCA, ANO-CCA and SROB-CCA secure. Let H be a collision-resistant hash function that maps the space of public-keys of KEM to its encapsulated key-space. We now construct $\overline{\mathsf{KEM}} = (\mathsf{KGen}, \overline{\mathsf{Encap}}, \overline{\mathsf{Decap}})$ as follows. For the public parameters of $\overline{\mathsf{KEM}}$, we first generate a pair of claw-free permutations with corresponding fixed public-key PK (see [11, Section 4.2] for a more formal definition) $f_1(\mathsf{PK}, \cdot)$ and $f_2(\mathsf{PK}, \cdot)$ with domain and range being the encapsulated key-space of KEM. Now $\overline{\mathsf{Encap}}(\mathsf{pk})$ returns $(C, \overline{k})$ where $(C, k) \leftarrow_{\$} \mathsf{Encap}(\mathsf{pk})$ and $\overline{k} := f_1(\mathsf{PK}, k)$. $\overline{\mathsf{Decap}}(\mathsf{sk}, C)$ returns $\overline{k}'$ where, for $k' \leftarrow \mathsf{Decap}(\mathsf{sk}, C)$, $\overline{k}' := f_1(\mathsf{PK}, k')$ if $k' \neq \perp$ and $\overline{k}' := f_2(\mathsf{PK}, H(\mathsf{pk}))$ if $k' = \perp$. Using straightforward reductions, it is not hard to show that $\overline{\mathsf{KEM}}$ is also IND-CCA and ANO-CCA secure. In addition, we can show that $\overline{\mathsf{KEM}}$ is SCFR-CCA secure by relying on the SROB-CCA security of KEM, collision-resistance of H and claw-freeness assumption w.r.t. $f_1(\mathsf{PK}, \cdot)$ and $f_2(\mathsf{PK}, \cdot)$.

Now let $\mathsf{DEM} = (\mathsf{Enc}, \mathsf{Dec})$ be an ROR-CCA secure AEAD which is additionally XROB-secure. We now describe an adversary $\mathcal{A}$ against the ANO-CCA security of the hybrid PKE scheme w.r.t. the composition of $\overline{\mathsf{KEM}}$ and DEM. Upon receiving two public-keys pk_0 and pk_1 (along with the public-parameters $f_1(\mathsf{PK}, \cdot)$ and $f_2(\mathsf{PK}, \cdot)$), $\mathcal{A}$ selects an arbitrary message m and forwards m to the ANO-CCA challenger. It then receives the ciphertext $C = (C_{\mathsf{KEM}}, C_{\mathsf{DEM}})$ where $(C_{\mathsf{KEM}}, k) \leftarrow_{\$} \overline{\mathsf{Encap}}(\mathsf{pk}_b)$ and $C_{\mathsf{DEM}} \leftarrow_{\$} \mathsf{Enc}(k, m)$, for bit $b \leftarrow_{\$} \{0, 1\}$. Then, $\mathcal{A}$ asks for the decryption of ciphertext $C' = (C_{\mathsf{KEM}}, C'_{\mathsf{DEM}})$ w.r.t. sk_0 where $C'_{\mathsf{DEM}} = \mathsf{Enc}(\hat{k}, m)$ with $\hat{k} = f_2(\mathsf{PK}, H(\mathsf{pk}_0))$. If the response is $\perp$, then $\mathcal{A}$ outputs 0; else, it outputs 1. We use similar arguments as that of [32, Claim 3.1] to show that $\mathcal{A}$ succeeds with a high probability. Complete details of the proof can be found in the full version [22].

The consequence of the above theorem (and its counterexample) is that, for implicit rejection KEMs, we cannot hope to transfer anonymity properties of the KEM to those of the hybrid PKE scheme resulting from the standard KEM-

$\mathsf{Encap}(\mathsf{pk})$		$\mathsf{Decap}(\mathsf{sk}, c)$	
1 :	$m \leftarrow\!\!\$\ \mathcal{M}$	1 :	Parse $c = (c_1, c_2)$
2 :	$c_1 \leftarrow \mathsf{Enc}(\mathsf{pk}, m; G(m))$	2 :	$m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c_1)$
3 :	$c_2 \leftarrow H'(m)$	3 :	$c_1' \leftarrow \mathsf{Enc}(\mathsf{pk}, m'; G(m'))$
4 :	$\boxed{c_2 \leftarrow H'(m, c_1)}$	4 :	**if** $c_1' = c_1 \wedge H'(m') = c_2$ **then**
5 :	$c \leftarrow (c_1, c_2)$	5 :	$\boxed{\textbf{if } c_1' = c_1 \wedge H'(m', c_1) = c_2 \textbf{ then}}$
6 :	$k = H(m, c)$	6 :	**return** $H(m', c)$
7 :	**return** (c, k)	7 :	**else return** $\perp$

Fig. 3. The KEM $\mathsf{HFO}^{\perp}[\mathsf{PKE}, G, H, H']$. Boxed code shows modifications to $\mathsf{HFO}^{\perp}[\mathsf{PKE}, G, H, H']$ required to obtain scheme $\mathsf{HFO}^{\perp'}[\mathsf{PKE}, G, H, H']$. Both constructed schemes reuse algorithm KGen from PKE.

DEM construction in a fully generic manner. To make further progress in this direction, then, we need to look more closely at specific KEM constructions.

4 Anonymity and Robustness of KEMs Obtained from Fujisaki-Okamoto Transforms in the QROM

Fujisaki and Okamoto [18–20] introduced generic transformations that turn weakly secure PKE schemes (e.g. OW-CPA or IND-CPA secure PKE schemes) into IND-CCA secure KEMs and PKE schemes. Several distinct transforms have emerged, each with slightly different flavours; we broadly follow the naming conventions in [24]. One main distinction is whether the constructed KEM offers implicit rejection ($\mathsf{FO}^{\not\perp}$) or explicit rejection ($\mathsf{QFO}_m^{\perp}$). As we have already seen, this distinction is important in considering robustness, and we divide our analysis of the FO transforms in the same way. Since all NIST PQC candidates in the KEM/PKE category except one alternate candidate offer implicit rejection, we mainly focus on the corresponding $\mathsf{FO}^{\not\perp}$ transform. Also, since we are mainly concerned with the post-quantum setting, our analysis that follows will be in the QROM.

4.1 KEMs With Explicit Rejection

Before we focus on the $\mathsf{FO}^{\not\perp}$ transform, we briefly discuss our results related to explicit-rejection KEMs. The paper [27] presents a variant of the Fujisaki-Okamato transform, namely $\mathsf{HFO}^{\perp}$, that results in IND-CCA secure KEMs in the QROM. Given a PKE scheme $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ (with message space $\mathcal{M}$) and hash functions G, H and H', the resulting $\mathsf{KEM}^{\perp} = \mathsf{HFO}^{\perp}[\mathsf{PKE}, G, H, H'] = (\mathsf{KGen}, \mathsf{Encap}, \mathsf{Decap})$ is described in Figure 3.

KGen'		$\mathsf{Encap}(\mathsf{pk})$		$\mathsf{Decap}(\mathsf{sk}', c)$	
1 :	$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}$	1 :	$m \leftarrow_{\$} \mathcal{M}$	1 :	Parse $\mathsf{sk}' = (\mathsf{sk}, s)$
2 :	$s \leftarrow_{\$} \mathcal{M}$	2 :	$r \leftarrow G(m)$	2 :	$m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$
3 :	$\mathsf{sk}' = (\mathsf{sk}, s)$	3 :	$c \leftarrow \mathsf{Enc}(\mathsf{pk}, m; r)$	3 :	$r' \leftarrow G(m')$
4 :	**return** $(\mathsf{pk}, \mathsf{sk}')$	4 :	$k \leftarrow H(m, c)$	4 :	$c' \leftarrow \mathsf{Enc}(\mathsf{pk}, m'; r')$
		5 :	**return** (c, k)	5 :	**if** $c' = c$ **then**
				6 :	**return** $H(m', c)$
				7 :	**else return** $H(s, c)$

Fig. 4. The KEM $\mathsf{FO}^{\not\perp}[\mathsf{PKE}, G, H]$.

We introduce a slight variant of the above transform, namely $\mathsf{HFO}^{\perp'}$, as shown in Figure 3. The only change is that the c_2 component of the ciphertext–used for so-called *plaintext confirmation*–is derived as $c_2 \leftarrow H'(m, c_1)$ instead of as $c_2 \leftarrow H'(m)$. However, this seemingly minor change not only allows the $\mathsf{HFO}^{\perp'}$ transform to result in IND-CCA secure KEMs, but also strongly anonymous (ANO-CCA secure) and robust (SROB-CCA secure) KEMs in the QROM. In the full version [22], we formally state and prove the corresponding theorems.

4.2 KEMs With Implicit Rejection

Given a PKE scheme $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\mathcal{M}$ and hash functions G and H, the KEM $\mathsf{KEM}^{\not\perp} = \mathsf{FO}^{\not\perp}[\mathsf{PKE}, G, H]$ is shown in Figure 4. As described in [24], the $\mathsf{FO}^{\not\perp}$ transform "implicitly" uses a modular transformation T that converts a OW-CPA/IND-CPA secure PKE scheme PKE into a *deterministic* PKE scheme $\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G] = (\mathsf{KGen}, \mathsf{Enc}', \mathsf{Dec}')$ that is secure in the presence of so-called *plaintext-checking attacks*. The deterministic encryption $\mathsf{Enc}'(\mathsf{pk}, m)$ returns c where $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m; G(m))$. The decryption $\mathsf{Dec}'(\mathsf{sk}, c)$ first computes $m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$ and then returns m' if the *re-encryption* check "$\mathsf{Enc}(\mathsf{pk}, m'; G(m')) = c$" succeeds; otherwise, $\perp$ is returned.

It was proved in [26] that the $\mathsf{FO}^{\not\perp}$ transform lifts IND-CPA security of PKE to IND-CCA security of $\mathsf{KEM}^{\not\perp}$ in the QROM. We provide some further enhancement results for $\mathsf{FO}^{\not\perp}$. They demonstrate that, provided the starting PKE scheme PKE and the derived deterministic scheme PKE_1 satisfy some mild security assumptions on anonymity (wANO-CPA[3]) and collision-freeness (SCFR-CPA) respectively, then $\mathsf{FO}^{\not\perp}$ confers strong anonymity (ANO-CCA) and collision-freeness (SCFR-CCA) to the final $\mathsf{KEM}^{\not\perp}$ in the QROM.

[3] The wANO-CPA security notion for PKE is a weaker variant of ANO-CPA where, in the corresponding security game, the challenger encrypts a uniformly random *secret* message under either of the two honestly generated public-keys and *only* provides the resulting ciphertext to the adversary, along with the generated public-keys.

Theorem 5. *Suppose* $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is* δ*-correct and has message space* $\mathcal{M}$*. Then for any* ANO-CCA *adversary* $\mathcal{A}$ *against* $\mathsf{KEM}^{\not\perp} = \mathsf{FO}^{\not\perp}[\mathsf{PKE}, G, H]$ *issuing at most* q_G *(resp.* q_H*) queries*[4] *to the quantum random oracle* G *(resp.* H*) and at most* q_D *queries to the (classical) decapsulation oracles, there exist* wANO-CPA *adversary* $\mathcal{B}$ *and* OW-CPA *adversary* $\mathcal{C}$ *against* PKE*, and* SCFR-CPA *adversary* $\mathcal{D}$ *against* $\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G]$ *issuing at most* q_G *queries to* G*, such that:*

$$\mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{KEM}^{\not\perp}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{wANO-CPA}}_{\mathsf{PKE}}(\mathcal{B}) + 2(q_G + q_H)\sqrt{\mathbf{Adv}^{\text{OW-CPA}}_{\mathsf{PKE}}(\mathcal{C})} + q_D \cdot \mathbf{Adv}^{\text{SCFR-CPA}}_{\mathsf{PKE}_1}(\mathcal{D}) + \frac{4q_H}{\sqrt{|\mathcal{M}|}} + 2q_G(q_D+2)\sqrt{2\delta}\,.$$

Moreover, the running times of $\mathcal{B}$*,* $\mathcal{C}$ *and* $\mathcal{D}$ *are the same as that of* $\mathcal{A}$*.*

Proof (sketch). In a reduction from ANO-CCA security of $\mathsf{KEM}^{\not\perp}$ to wANO-CPA security of PKE, note that we need to simulate two different decapsulation oracles consistently without possessing the corresponding secret keys. Our approach is to generalize the simulation trick of [26, 34] in the QROM from a single-key setting (in the context of IND-CCA security) to a two-key setting (ANO-CCA). Namely, given two public-keys pk_0, pk_1, note that the encapsulation algorithm for both of them uses a common key-derivation function (KDF) "$k = H(m, c)$" (see Fig. 4). So we associate this KDF with two *secret* random functions H_0 and H_1 as follows: given an input (m, c), if $c = \mathsf{Enc}(\mathsf{pk}_i, m; G(m))$ (i.e., c results likely from $\mathsf{Encap}(\mathsf{pk}_i)$), then replace the KDF with "$k = H_i(c)$". Note that in this case, we can simply simulate the decapsulation oracles as $\mathsf{Decap}(\mathsf{sk}_i, c) = H_i(c)$ without requiring the secret keys. Now to argue that this replacement of KDF is indistinguishable w.r.t. an adversary, we require the functions $\mathsf{Enc}(\mathsf{pk}_i, \cdot\,; G(\cdot))$ to be injective. Thus, following [26], we first replace oracle G with G' where G' only returns "good" encryption randomness w.r.t. $(\mathsf{pk}_0, \mathsf{sk}_0)$ and $(\mathsf{pk}_1, \mathsf{sk}_1)$ – i.e., $\forall m$, $\mathsf{Dec}(sk_i, \mathsf{Enc}(\mathsf{pk}_i, m; G'(m))) = m$, for $i \in \{0, 1\}$. We again generalize the argument of [26] from a single-key setting to a two-key setting to show that this replacement of G is indistinguishable, relying on the δ-correctness of PKE.

However, note that we additionally have to account for pairs (m, c) which satisfy $\mathsf{Enc}(\mathsf{pk}_0, m; G'(m)) = \mathsf{Enc}(\mathsf{pk}_1, m; G'(m)) = c$; in this case, the reduction does not know which public-key was used to generate c during key-encapsulation. So we rely on SCFR-CPA security to argue that it is computationally hard for an adversary to ask for the (classical) decapsulation of such "peculiar" ciphertexts c. Such a c results in $\mathsf{Dec}(sk_0, c) = \mathsf{Dec}(sk_1, c) = m$, thereby breaking the SCFR-CPA security of $\mathsf{T}[\mathsf{PKE}, G']$, and hence, that of $\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G]$ (up to an additive loss). Complete details of the proof can be found in the full version [22]. Note that it is similar in structure to that of [26, Theorem 1] in terms of the

[4] Following [24, 26], we make the convention that the number q_O of queries made by an adversary $\mathcal{A}$ to a random oracle O counts the total number of times O is executed in the corresponding security experiment; i.e., the number of $\mathcal{A}$'s explicit queries to O plus the number of implicit queries to O made by the experiment.

sequence of game-hops. But for the sake of completeness, we provide a self-contained proof.

To establish strong collision-freeness of the implicit-rejection KEMs constructed using $\mathsf{FO}^{\not\perp}$, we require the following *claw-freeness* property of quantum random oracles.

Lemma 1 ([39, Lemma 2.3]). *There is a universal constant α (< 648) such that the following holds: Let $\mathcal{X}_0$, $\mathcal{X}_1$ and $\mathcal{Y}$ be finite sets. Let $N_0 = |\mathcal{X}_0|$ and $N_1 = |\mathcal{X}_1|$, with $N_0 \leq N_1$. Let $H_0 : \mathcal{X}_0 \to \mathcal{Y}$ and $H_1 : \mathcal{X}_1 \to \mathcal{Y}$ be two random oracles.*

If an unbounded time quantum adversary $\mathcal{A}$ makes a query to H_0 and H_1 at most q times, then we have

$$\Pr[H_0(x_0) = H_1(x_1) : (x_0, x_1) \leftarrow \mathcal{A}^{H_0,H_1}] \leq \frac{\alpha(q+1)^3}{|\mathcal{Y}|},$$

where all oracle accesses of $\mathcal{A}$ can be quantum.

For the following result, we in-fact need a weaker property than the one described in the above lemma; namely, it's hard for an adversary to return a value $x \in \mathcal{X}_0 \cap \mathcal{X}_1$ such that $H_0(x) = H_1(x)$. We leave the derivation of the corresponding upper-bound as an open problem.

Theorem 6. *Suppose $\mathsf{PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is δ-correct. Then for any* SCFR-CCA *adversary $\mathcal{A}$ against $\mathsf{KEM}^{\not\perp} = \mathsf{FO}^{\not\perp}[\mathsf{PKE}, G, H]$ issuing at most q_D queries to the (classical) decapsulation oracles, at most q_G (resp. q_H) queries to the quantum random oracle G (resp. H), there exists an* SCFR-CPA *adversary $\mathcal{B}$ against $\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G]$ issuing at most q_G queries to G such that*

$$\mathbf{Adv}^{\text{SCFR-CCA}}_{\mathsf{KEM}^{\not\perp}}(\mathcal{A}) \leq q_D \cdot \mathbf{Adv}^{\text{SCFR-CPA}}_{\mathsf{PKE}_1}(\mathcal{B}) + \frac{\alpha(q_H+1)^3}{|\mathcal{K}|} + \frac{4q_H}{\sqrt{|\mathcal{M}|}} + 2q_G(q_D+2)\sqrt{2\delta}\,.$$

Here $\mathcal{K}$ denotes the encapsulated key-space of $\mathsf{KEM}^{\not\perp}$ and α (< 648) is the constant from Lemma 1. The running time of $\mathcal{B}$ is the same as that of $\mathcal{A}$.

Proof (sketch). Here we reduce the SCFR-CCA security of $\mathsf{KEM}^{\not\perp}$ to the hardness of claw-finding w.r.t. QROs. The proof is similar in structure to that of Theorem 5. Namely, we start with an SCFR-CCA adversary $\mathcal{A}$ and do a similar sequence of game-hops until the point where the decapsulation oracles don't require the corresponding secret keys – namely, $\mathsf{Decap}(\mathsf{sk}_i, c) = H_i(c)$ for (secret) random functions $H_0, H_1 : \overline{\mathcal{C}} \to \mathcal{K}$, where $\overline{\mathcal{C}}$ denotes the ciphertext space of $\mathsf{PKE}/\mathsf{KEM}^{\not\perp}$. Now $\mathcal{A}$ wins this modified SCFR-CCA game if it returns c such that $\mathsf{Decap}(\mathsf{sk}_0, c) = \mathsf{Decap}(\mathsf{sk}_1, c)$, or equivalently, $H_0(c) = H_1(c)$. Note that (c, c) is then a *claw* w.r.t. the pair of QROs (H_0, H_1). Hence, we can bound $\mathcal{A}$'s winning probability using Lemma 1. A complete proof can be found in the full version [22].

From Theorems 5 and 6, we see that by applying the $\mathsf{FO}^{\not\perp}$ transformation to weakly secure (i.e., OW-CPA) and weakly anonymous (i.e., wANO-CPA) PKE schemes, with an additional assumption of strong collision-freeness (against chosen plaintext attacks) of the deterministic version of the underlying PKE scheme ($\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G]$), not only do we obtain strongly secure KEMs (i.e., IND-CCA security) but also KEMs that are strongly anonymous (i.e., ANO-CCA) and are strongly collision-free against chosen ciphertext attacks (SCFR-CCA) in the QROM.

At the same time, we showed a negative result in Theorem 4. It essentially shows that starting from a KEM that is IND-CCA, ANO-CCA and SCFR-CCA secure does not *generically* result in a strongly anonymous (ANO-CCA) hybrid PKE scheme via the KEM-DEM composition. Nonetheless, we are able to show the following positive result for KEMs obtained via the $\mathsf{FO}^{\not\perp}$ transform. We only need a weak additional property of the underlying PKE scheme, namely that it be γ-spread.

Theorem 7. *Let* $\mathsf{PKE}^{hy} = (\mathsf{KGen}', \mathsf{Enc}^{hy}, \mathsf{Dec}^{hy})$ *be a hybrid encryption scheme obtained by composing* $\mathsf{KEM}^{\not\perp} = \mathsf{FO}^{\not\perp}[\mathsf{PKE}, G, H]$ *with a one-time authenticated encryption scheme* $\mathsf{DEM} = (\mathsf{Enc}^{sym}, \mathsf{Dec}^{sym})$*. Suppose* PKE *is* δ*-correct and* γ*-spread (with message space* $\mathcal{M}$*). Then for any* ANO-CCA *adversary* $\mathcal{A}$ *against* PKE^{hy} *issuing at most* q_G *(resp.* q_H*) queries to the quantum random oracle* G *(resp.* H*), there exist* ANO-CCA *adversary* $\mathcal{B}$ *and* IND-CCA *adversary* $\mathcal{C}$ *against* $\mathsf{KEM}^{\not\perp}$*,* WCFR-CPA *adversary* $\mathcal{D}$ *against* $\mathsf{PKE}_1 = \mathsf{T}[\mathsf{PKE}, G]$*, and* INT-CTXT *adversary* $\mathcal{E}$ *against* DEM *such that:*

$$\mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{PKE}^{hy}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{KEM}^{\not\perp}}(\mathcal{B}) + 2\mathbf{Adv}^{\text{IND-CCA}}_{\mathsf{KEM}^{\not\perp}}(\mathcal{C}) + \mathbf{Adv}^{\text{WCFR-CPA}}_{\mathsf{PKE}_1}(\mathcal{D}) + 2\mathbf{Adv}^{\text{INT-CTXT}}_{\mathsf{DEM}}(\mathcal{E}) + \frac{4q_H}{\sqrt{|\mathcal{M}|}} + 4q_G\sqrt{\delta} + 2^{-\gamma}\,.$$

Moreover, the running times of $\mathcal{B}$*,* $\mathcal{C}$ *and* $\mathcal{E}$ *are the same as that of* $\mathcal{A}$*. The running time of* $\mathcal{D}$ *is independent (and less than that) of the running time of* $\mathcal{A}$*.*

Proof (sketch). We use the proof of Theorem 1. Let $(\mathsf{pk}_0, \mathsf{sk}'_0)$ and $(\mathsf{pk}_1, \mathsf{sk}'_1)$ be two key-pairs generated in the ANO-CCA security game w.r.t. PKE^{hy}, and $b \leftarrow_{\$} \{0,1\}$ be the challenge bit. Let $c^* = (c^*_1, c^*_2)$ be the challenge ciphertext given to an adversary $\mathcal{A}$; i.e., $(c^*_1, k^*) \leftarrow \mathsf{KEM}^{\not\perp}.\mathsf{Encap}(\mathsf{pk}_b)$ and $c^*_2 \leftarrow \mathsf{Enc}^{sym}(m)$ where m is chosen by $\mathcal{A}$ upon first receiving pk_0, pk_1. In the proof of Theorem 1, we make some initial game-hops to modify the $\mathsf{Dec}^{hy}(sk'_{1-b}, \cdot)$ oracle such that if the query is of the form (c^*_1, c_2), the oracle returns $\perp$. There we rely on the WROB-CPA security of the underlying KEM to justify this modification. However, $\mathsf{KEM}^{\not\perp}$ is trivially not WROB-CPA secure. Nevertheless, we show that by relying on γ-spreadness of PKE, WCFR-CPA security of PKE_1 and INT-CTXT security of DEM, we can still make the above modification of the $\mathsf{Dec}^{hy}(sk'_{1-b}, \cdot)$ oracle. From that point on, we essentially use the same game-hops as in the proof of Theorem 1 in our reduction to ANO-CCA security of $\mathsf{KEM}^{\not\perp}$. Complete details can be found in the full version [22].

5 Anonymity and Robustness of NIST PQC Candidates

After analyzing the anonymity and robustness enhancing properties of the "standard" FO transforms in Section 4, we extend our analysis to the specific instantiations of these transforms used by Classic McEliece, proto-Saber (the simplified version of Saber in [14]) and FrodoKEM. We conclude this section by discussing some limitations of our techniques w.r.t. analyzing Saber and Kyber.

5.1 Classic McEliece

KGen'		$\mathsf{Encap}(\mathsf{pk})$		$\mathsf{Decap}(\mathsf{sk}', (c, h))$	
1 :	$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}$	1 :	$m \leftarrow_{\$} \mathcal{M}$	1 :	Parse $\mathsf{sk}' = (\mathsf{sk}, \mathsf{pk}, s)$
2 :	$s \leftarrow_{\$} \mathcal{M}$	2 :	$c \leftarrow \mathsf{Enc}(\mathsf{pk}, m)$	2 :	$m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$
3 :	$\mathsf{sk}' \leftarrow (\mathsf{sk}, \mathsf{pk}, s)$	3 :	$h \leftarrow H_2(m)$	3 :	$c' \leftarrow \mathsf{Enc}(\mathsf{pk}, m')$
4 :	**return** $(\mathsf{pk}, \mathsf{sk}')$	4 :	$k \leftarrow H_1(m, (c, h))$	4 :	**if** $c' = c \wedge H_2(m') = h$ **then**
		5 :	**return** $((c, h), k)$	5 :	**return** $H_1(m', (c, h))$
				6 :	**else return** $H_0(s, (c, h))$

Fig. 5. Classic McEliece uses a slight variant of the $\mathsf{FO}^{\not\perp}$ transform that starts with deterministic PKE schemes. Here H_0 and H_1 are two different hash functions. The so-called "Dent hash" H_2 is used as an additional component in the KEM ciphertext [3].

Classic McEliece (CM) as defined in its third round NIST specification [3] applies a slight variant of the $\mathsf{FO}^{\not\perp}$ transform to its starting deterministic PKE scheme (see Fig. 5). It can easily be shown that our generic transformation results on $\mathsf{FO}^{\not\perp}$, namely Theorems 5 and 6, apply to the $\mathsf{FO}^{\not\perp}$-like transformation used by CM, while accounting for the additional "Dent hash". Hence, the only thing that would remain to be analyzed is whether the base PKE scheme used by CM satisfies the pre-requisite security properties of Theorems 5 and 6, namely wANO-CPA and SCFR-CPA. As we show next, the base PKE scheme used by CM fails to be collision-free in a striking way that rules out the application of these results. This failure also propagates to PKE schemes built from the CM KEM via the standard KEM-DEM construction.

The base CM scheme: The base CM scheme is deterministic. To encrypt a message m, first encode m as a binary column vector e of some fixed length n and fixed Hamming weight t. Then compute ciphertext $c = He \in \mathbb{F}_2^n$ where H is an $(n-k) \times n$ matrix of the form $H = (I_{n-k} \,|\, T)$, where T is some $(n-k) \times k$ matrix whose value is unimportant below. Matrix H is the parity check matrix of an error correcting code whose error correcting capacity is at least t. Decryption is done by using the private key to rewrite matrix H in such a way that efficient

decoding can be performed to recover e with perfect correctness. The base CM scheme is closely related to the Niederreiter variant of the McEliece PKE scheme.

Collision-freeness of the base CM scheme: Recall that we would require the base CM scheme to satisfy the SCFR-CPA property in order to make use of our generic results concerning the $\text{FO}^{\not\perp}$ transform. This property is crucial in the CPA $\rightarrow$ CCA security proofs where we have to simulate the decapsulation oracles under two different secret keys without access to the keys. As we will show now, the base CM scheme is not SCFR-CPA secure, nor even WCFR-CPA secure. In fact, we can go further and exhibit a strong robustness failure of the base CM scheme, and explain how it leads to robustness failures in the CM KEM and hybrid PKE schemes built from it.

Consider any weight t error vector e in which the t 1's in e are concentrated in the first $n-k$ bit positions of e (in all the parameter sets used in Classic McEliece, $n-k=mt \geq t$, for a positive integer m, so this is always possible). We call such an e *concentrated*. Note that any concentrated e can be written $e = \binom{e_{n-k}}{0_k}$ with e_{n-k} of length $n-k$ and 0_k being the vector of k zeros. Since encryption is done by computing $c = He$, and H is of the form $(I_{n-k} \,|\, T)$, it is easy to see that c is a fixed vector independent of the T part of H: namely, $He = e_{n-k}$ which depends only on the first $n-k$ bit positions of e.

Note that this property holds independent of the public key of the base CM scheme (which is effectively the matrix H). Thus there is a class of base CM messages (of size $\binom{n-k}{t}$) for which the resulting ciphertext c can be predicted as a function of the message *without even knowing the public key*. By correctness of the base CM scheme, such ciphertexts must decrypt to the selected message *under any base CM scheme private key*.

It is immediate that this property can be used to violate SCFR-CPA and WCFR-CPA security of the base CM scheme. This presents a significant barrier to the application of our general theorems for establishing robustness and anonymity of the full CM KEM.

Robustness of the CM KEM and Hybrid PKEs derived from it: The base CM scheme is used to construct the CM KEM according to procedure described in Figure 5. This means that the CM KEM encapsulations are also of the form $c = (He, H_2(e))$ where $H_2(\cdot)$ is a hash function; meanwhile the encapsulated keys are set as $H_1(e, c)$ where $H_1(\cdot)$ is another hash function. The CM KEM performs implicit rejection, so one cannot hope for robustness. However, one might hope for some form of collision-freeness. Our analysis above shows that the CM KEM does not provide even this, since when e is concentrated, $c = (He, H_2(e))$ decapsulates to $H_1(e, c)$ under any CM private key.

Finally, one might ask about the robustness of PKE scheme built by combining the CM KEM with a DEM in the standard way. Again, such a PKE cannot be strongly collision free (and therefore not strongly robust either), since it is trivial using our observations to construct a hybrid PKE ciphertext that decrypts correctly under *any* CM private key to *any* fixed choice of message m (without

even knowing the public key). To see this, simply consider hybrid ciphertexts of the form $(He, H_2(e), \text{AEAD.}\mathsf{Enc}(K, m; r))$ where e is concentrated, $K = H_1(e, c)$ is the symmetric key encapsulated by the KEM part $c = (He, H_2(e))$ of the hybrid ciphertext, and r is some fixed randomness for the AEAD scheme. Such ciphertexts decrypt to the freely chosen message m under any CM private key.

Robustness could plausibly be conferred on this hybrid PKE scheme by including a hash of the public key in the key derivation step. However CM keys are large, so this would have a negative effect on performance. Robustness is *not* conferred in general by replacing the DEM with an AEAD scheme and including the hash of the public key in the associated data to create a "labelled DEM". This is easy to see by adapting the counter-example construction used in the proof of Theorem 3.

Further remarks on CM: The analysis above shows that we cannot hope to establish anonymity or robustness of the CM KEM or PKEs built from it via the standard KEM-DEM construction using the sequence of results in this paper. But this does not rule out more direct approaches to proving anonymity. For example, Persichetti [33] has analysed the anonymity of a scheme called HN (for "hybrid Niederreiter") that is rather close to the natural hybrid scheme one would obtain from CM. However, the analysis is in the ROM rather than the QROM. We are not aware of any further analysis of the anonymity properties of schemes that are close to CM and that might be easily adapted to CM.

In the context of the NIST PQC process, it remains an important open problem to establish anonymity of the CM scheme.

5.2 proto-Saber

KGen'		$\mathsf{Encap}(\mathsf{pk})$		$\mathsf{Decap}(\mathsf{sk}', c)$	
1 :	$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}$	1 :	$m \leftarrow\!\!\$\, \mathcal{M}$	1 :	Parse $\mathsf{sk}' = (\mathsf{sk}, \mathsf{pk}, F(\mathsf{pk}), s)$
2 :	$s \leftarrow\!\!\$\, \mathcal{M}$	2 :	$h \leftarrow F(\mathsf{pk})$	2 :	$m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$
3 :	$\mathsf{pk}' \leftarrow (\mathsf{pk}, F(\mathsf{pk}))$	3 :	$(\hat{k}, r) \leftarrow G(h, m)$	3 :	$(\hat{k}', r') \leftarrow G(F(\mathsf{pk}), m')$
4 :	$\mathsf{sk}' \leftarrow (\mathsf{sk}, \mathsf{pk}', s)$	4 :	$c \leftarrow \mathsf{Enc}(\mathsf{pk}, m; r)$	4 :	$c' \leftarrow \mathsf{Enc}(\mathsf{pk}, m'; r')$
5 :	**return** $(\mathsf{pk}, \mathsf{sk}')$	5 :	$k \leftarrow H(\hat{k}, c)$	5 :	**if** $c' = c$ **then**
		6 :	**return** (c, k)	6 :	**return** $H(\hat{k}', c)$
				7 :	**else return** $H(s, c)$

Fig. 6. pSaber uses a variant of the $\mathsf{FO}^{\not\perp}$ transform. Here G, F and H are hash functions.

The scheme "proto-Saber" (pSaber for short) is a KEM that was introduced in [14] and which is included in the NIST third round specification document for

Saber [7]. Saber and pSaber use the same base PKE scheme but apply *different* FO-type transforms to obtain their respective KEMs. The QROM IND-CCA security claims for Saber [7, Theorem 6.5] seem to have been taken directly from those for pSaber [14, Theorem 6] without any modification. However, as we will explain below, there are issues with pSaber's IND-CCA security claims, and yet further issues for Saber's.

Now pSaber uses a transform that differs significantly from the standard $\mathsf{FO}^{\not\perp}$ one (see Fig. 6). These significant deviations act as an obstacle to applying our generic results on anonymity and SCFR enhancement of $\mathsf{FO}^{\not\perp}$ to pSaber. The nature of these deviations also led us to ask whether they also act as a barrier in applying the results of [26] to establish the IND-CCA security of pSaber, as claimed in [14]. We believe this to be the case, as we explain next.

IND-CCA security of pSaber in the QROM: We claim that the specific proof techniques used by [26], to obtain relatively tight IND-CCA security bounds for the standard $\mathsf{FO}^{\not\perp}$ transform in the QROM, do not directly apply to pSaber's variant of the FO transform. An important trick used by [26] in their security proofs of $\mathsf{FO}^{\not\perp}$ is to replace the computation of the key "$k \leftarrow H(m, c)$" with "$k \leftarrow H'(g(m))(= H'(c))$" for function $g(\cdot) = \mathsf{Enc}(\mathsf{pk}, \cdot\,; G(\cdot))$ and a secret random function $H'(\cdot)$; note that in this case, we simply have $\mathsf{Decap}(\mathsf{sk}, c) = H'(c)$ leading to an "efficient" simulation of the decapsulation oracle without using the secret key sk. To justify this replacement, the authors of [26] then argue about the injectivity of $g(\cdot)$, relying on the correctness of the underlying PKE scheme to establish this.

But in pSaber, the keys are computed as "$k \leftarrow H(\hat{k}, c)$" where the "pre-key" $\hat{k}$ is derived as a hash of the message m (to be specific, $(\hat{k}, r) \leftarrow G(F(\mathsf{pk}), m)$). So there is an extra *layer* of hashing between m and the computation of k. Hence, to use a similar trick as [26], we would require some additional injectivity arguments. Thus, strictly speaking, the proof techniques of [26] do not directly apply to pSaber.

Nevertheless, we are able to overcome the above barrier by adapting the analysis of $\mathsf{FO}^{\not\perp}$ in [26] to obtain an explicit IND-CCA security proof for pSaber in the QROM, with the *same* tightness as claimed in [14]. The formal proof can be found in the full version [22]. We give a high-level overview of our approach below.

First, note that we can replace the step "$(\hat{k}, r) \leftarrow G(F(\mathsf{pk}), m)$" in pSaber's encapsulation by "$\hat{k} \leftarrow G_{\hat{k}}(m)$" and "$r \leftarrow G_r(m)$" for two fresh random oracles $G_{\hat{k}}$, $G_r : \{0,1\}^{256} \to \{0,1\}^{256}$. Now our key observation is that the extra layer of hashing "$G_{\hat{k}}(\cdot)$" between m and k is actually *length-preserving*, i.e., the hash function has the same domain and range. So following [24, 37], we can replace the random oracle $G_{\hat{k}}(\cdot)$ with a random *polynomial* of degree $2q_G - 1$ over a finite field representation of $\{0,1\}^{256}$ (i.e., a $2q_G$-wise independent function). Here q_G is the number of queries made to oracle G in the IND-CCA security reduction for pSaber. Thanks to a result in [40], this change is perfectly indistinguishable to an adversary making at most q_G queries to $G_{\hat{k}}$. This will allow us to recover

m from a corresponding pre-key value $\hat{k}$ by computing roots of the polynomial $G_{\hat{k}}(x) - \hat{k}$. Hence we can *invert* this "nested" hashing of m in order to apply the trick of [26]. Namely, we can now replace the key derivation "$k \leftarrow H(\hat{k}, c)$" with "$k \leftarrow H'(g(m))(= H'(c))$" for function $g(\cdot) = \mathsf{Enc}(\mathsf{pk}, \cdot\,; G_r(\cdot))$, where in addition, m is a root of the polynomial $G_{\hat{k}}(x) - \hat{k}$.

Anonymity and Robustness of pSaber in the QROM: Our approach to repairing pSaber's IND-CCA proof also allows us to derive proofs of anonymity and SCFR enhancement for pSaber with similar tightness.

Now pSaber, and Saber, is a KEM whose claimed security relies on the hardness of the module learning-with-rounding problem, or mod-LWR for short (see [7, 14] for a precise description of the assumption). In the following, we prove the ANO-CPA security of the base PKE scheme Saber.PKE that is used by pSaber, and also currently used by Saber (as per [7]). The result relies on the hardness of mod-LWR. The proof can be found in the full version [22]. The proof adapts the proof of [14, Theorem 3] showing IND-CPA security of Saber.PKE.

Theorem 8. *For any* ANO-CPA *adversary* $\mathcal{A}$ *against* Saber.PKE*, there exists a distinguisher* $\mathcal{B}_1$ *(resp.,* $\mathcal{B}_2$*) between* l *(resp.* $l+1$*) samples from a mod-LWR distribution from that of a uniform distribution, with corresponding parameters* l, μ, q *and* p*, such that*

$$\mathbf{Adv}^{\text{ANO-CPA}}_{\mathsf{Saber.PKE}}(\mathcal{A}) \leq 2 \cdot \mathbf{Adv}^{\text{mod-lwr}}_{l,l,\mu,q,p}(\mathcal{B}_1) + \mathbf{Adv}^{\text{mod-lwr}}_{l+1,l,\mu,q,p}(\mathcal{B}_2).$$

Moreover, the running times of $\mathcal{B}_1$ *and* $\mathcal{B}_2$ *are the same as that of* $\mathcal{A}$*.*

Now we establish anonymity and strong collision-freeness of pSaber KEM, which we will denote as "pSaber.KEM" in the following to contrast the scheme with Saber.PKE. We use similar proof strategies that were used to show the same properties for $\mathsf{FO}^{\not\perp}$ in Section 4 (Theorems 5 and 6). A major difference is that instead of relying on the SCFR-CPA security property of Saber.PKE (specifically, its deterministic version), we again rely on hardness of the *claw-finding* problem in a quantum setting (see Lemma 1).

In our next results, we show that the stronger properties of ANO-CCA and SCFR-CCA hold for pSaber.KEM. Below we define $\mathbf{Coll}^{F}_{\mathsf{Saber.PKE}}$ as the probability of the event "$F(\mathsf{pk}_0) = F(\mathsf{pk}_1)$" where pk_0 and pk_1 are two honestly-generated Saber.PKE public-keys. Given the space of Saber's public-keys is sufficiently large (of size greater than 2^{256}), if the hash function F is sufficiently collision-resistant, then $\mathbf{Coll}^{F}_{\mathsf{Saber.PKE}}$ can be considered to be negligible. The proofs of Theorems 9 and 10 can be found in the full version [22].

Theorem 9. *Given* $\mathsf{Saber.PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is* δ*-correct, for any* ANO-CCA *adversary* $\mathcal{A}$ *against* $\mathsf{pSaber.KEM} = (\mathsf{KGen}', \mathsf{Encap}, \mathsf{Decap})$ *issuing at most* q_D *classical queries to the decapsulation oracles, at most* q_G *(resp.* q_H*) quantum queries to the random oracle* G *(resp.* H*), there exist* ANO-CPA *adversary* $\mathcal{B}$*,*

OW-CPA *adversary* $\mathcal{C}$ *against* $\mathsf{Saber.PKE}$ *and a distinguisher* $\mathcal{B}_1$ *between* l *samples from a mod-LWR distribution and a uniform distribution with corresponding parameters* l, μ, q *and* p*, such that*

$$\mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{pSaber.KEM}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{ANO-CPA}}_{\mathsf{Saber.PKE}}(\mathcal{B}) + 2(q_G + q_H)\sqrt{\mathbf{Adv}^{\text{OW-CPA}}_{\mathsf{Saber.PKE}}(\mathcal{C})}$$
$$+ \mathbf{Coll}^{F}_{\mathsf{Saber.PKE}} + \frac{\alpha(q_G+1)^3}{2^{256}} + \mathbf{Adv}^{mod\text{-}lwr}_{l,l,\mu,q,p}(\mathcal{B}_1) + \frac{2}{2^{256}} + \frac{4q_H}{2^{128}} + 8q_G\sqrt{\delta}$$

Here α (< 648) *is the constant from Lemma 1. The running times of* $\mathcal{B}$ *and* $\mathcal{C}$ *are the same as that of* $\mathcal{A}$*. The running time of* $\mathcal{B}_1$ *is independent (and less than that) of the running time of* $\mathcal{A}$*.*

Theorem 10. *Given* $\mathsf{Saber.PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is* δ*-correct, for any* SCFR-CCA *adversary* $\mathcal{A}$ *against* $\mathsf{pSaber.KEM} = (\mathsf{KGen}', \mathsf{Encap}, \mathsf{Decap})$ *issuing at most* q_D *queries to the (classical) decapsulation oracles, at most* q_G *(resp.* q_H*) queries to the quantum random oracle* G *(resp.* H*), we have*

$$\mathbf{Adv}^{\text{SCFR-CCA}}_{\mathsf{pSaber.KEM}}(\mathcal{A}) \leq \mathbf{Coll}^{F}_{\mathsf{Saber.PKE}} + \frac{\alpha(q_G+1)^3}{2^{256}} + \frac{\alpha(q_H+1)^3}{2^{256}} + \frac{4q_H}{2^{128}}$$

Here α (< 648) *is the constant from Lemma 1.*

Regarding hybrid PKE schemes obtained from $\mathsf{pSaber.KEM}$ via the KEM-DEM composition, we additionally show that such PKE schemes satisfy the stronger ANO-CCA notion of anonymity, in a similar vein to Theorem 7 w.r.t. $\mathsf{FO}^{\not\perp}$-based KEMs. The proof can be found in the full version [22].

Theorem 11. *Let* $\mathsf{pSaber.PKE}^{hy} = (\mathsf{KGen}', \mathsf{Enc}^{hy}, \mathsf{Dec}^{hy})$ *be a hybrid encryption scheme obtained by composing* $\mathsf{pSaber.KEM} = (\mathsf{KGen}', \mathsf{Encap}, \mathsf{Decap})$ *with a one-time authenticated encryption scheme* $\mathsf{DEM} = (\mathsf{Enc}^{sym}, \mathsf{Dec}^{sym})$*. Given* $\mathsf{Saber.PKE} = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is* δ*-correct, then for any* ANO-CCA *adversary* $\mathcal{A}$ *against* $\mathsf{pSaber.PKE}^{hy}$ *issuing at most* q_G *(resp.* q_H*) queries to the quantum random oracle* G *(resp.* H*), there exist* ANO-CCA *adversary* $\mathcal{B}$*,* IND-CCA *adversary* $\mathcal{C}$ *against* $\mathsf{pSaber.KEM}$*,* INT-CTXT *adversary* $\mathcal{E}$ *against* DEM *and distinguisher* $\mathcal{B}_1$ *between* l *samples from a mod-LWR distribution and a uniform distribution, with corresponding parameters* l, μ, q *and* p*, such that*

$$\mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{pSaber.PKE}^{hy}}(\mathcal{A}) \leq \mathbf{Adv}^{\text{ANO-CCA}}_{\mathsf{pSaber.KEM}}(\mathcal{B}) + 2\mathbf{Adv}^{\text{IND-CCA}}_{\mathsf{pSaber.KEM}}(\mathcal{C}) + \mathbf{Coll}^{F}_{\mathsf{Saber.PKE}}$$
$$+ 2\mathbf{Adv}^{\text{INT-CTXT}}_{\mathsf{DEM}}(\mathcal{E}) + \mathbf{Adv}^{mod\text{-}lwr}_{l,l,\mu,q,p}(\mathcal{B}_1) + \frac{4q_H}{2^{128}} + 8q_G\sqrt{\delta} + \frac{1}{2^{256}}$$

and the running times of $\mathcal{B}$*,* $\mathcal{C}$ *and* $\mathcal{E}$ *are the same as that of* $\mathcal{A}$*. The running time of* $\mathcal{B}_1$ *is independent (and less than that) of the running time of* $\mathcal{A}$*.*

At the same time, from Theorems 2 and 10, we note that if the DEM component is also FROB secure, then the corresponding hybrid PKE scheme will be strongly robust (i.e., SROB-CCA secure). Hence, our above results give a complete picture of anonymity and robustness properties of pSaber as well as the hybrid PKE schemes derived from it.

5.3 FrodoKEM

FrodoKEM uses an *identical* FO-type transform, described as "$\mathsf{FO}^{\not\perp\prime}$" in the specification document [4], as pSaber does (see Fig. 6) on its base PKE scheme "FrodoPKE". Hence, our positive results on tight IND-CCA security, anonymity and robustness of pSaber should also apply to FrodoKEM in a similar fashion; instead of relying on hardness of mod-LWR problem, we have to rely on hardness of the learning-with-errors (LWE) problem.

For example, when it comes to establishing anonymity of FrodoKEM, we only need to prove the ANO-CPA security of FrodoPKE and then rely on the "ANO-CPA $\rightarrow$ ANO-CCA" enhancement property of $\mathsf{FO}^{\not\perp\prime}$ (LWE variant of Theorem 9). The ANO-CPA security of FrodoPKE can be shown in a similar manner as that of $\mathsf{Saber.PKE}$ (Theorem 8): namely, by adapting the IND-CPA security proof of FrodoPKE. To be more precise, it is shown in [4, 30] w.r.t. FrodoPKE $= (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ that given $(\mathsf{pk}, \mathsf{sk}) \leftarrow_{\$} \mathsf{KGen}$ and *any* valid message m, the distribution $(\mathsf{pk}, \mathsf{Enc}(\mathsf{pk}, m))$ is computationally indistinguishable from (pk, c^*) where c^* is a uniformly random ciphertext, relying on the LWE hardness assumption. Hence, in the ANO-CPA security game w.r.t. FrodoPKE, given two honestly-generated public-keys $\mathsf{pk}_0, \mathsf{pk}_1$ and a message m chosen by an adversary, it cannot distinguish the encryption of m under pk_0 from a uniformly random ciphertext that is independent of pk_0. Similarly, the adversary also cannot distinguish the uniformly random ciphertext from the encryption of m under pk_1. It follows that the adversary cannot distinguish between the encryptions of m under pk_0 and pk_1, thereby establishing the ANO-CPA security of FrodoPKE.

5.4 Saber and Kyber

It turns out that Saber and Kyber implement a transform that deviates *even further* from the $\mathsf{FO}^{\not\perp}$ transform than pSaber does (see Fig. 7). Specifically, the keys in Saber are computed as "$k \leftarrow F(\hat{k}, F(c))$" where the "pre-key" $\hat{k}$ is derived as a hash of the message m (to be specific, $(\hat{k}, r) \leftarrow G(F(\mathsf{pk}), m)$). Again there is an extra hashing step between m and the computation of k, as we have seen for pSaber. But at the same time, there is also a "nested" hashing of ciphertext in the key-derivation (i.e., Saber uses "$F(c)$" in place of just "c") as opposed to the standard "single" hashing in $\mathsf{FO}^{\not\perp}$ and pSaber.

This "extra" hash of the ciphertext is a significant barrier to applying the techniques we used to prove anonymity of pSaber. It also acts as a barrier when trying to apply the generic proof techniques of [26] towards establishing the IND-CCA security of Saber in the QROM, with the *same* bounds as was claimed in its NIST third round specification [7]. At least for pSaber, as discussed above, we were able to account for the "nested" hashing of message because it was *length-preserving*. However, this is not the case for "$F(c)$" in Saber. We believe that an IND-CCA security reduction for Saber, along the lines of [26], in the QROM would need to rely on the collision-resistance of $F(\cdot)$ when modelled as a quantum random oracle. But a corresponding additive term is missing in the IND-CCA security bounds claimed in the Saber specification. We have shared

KGen'	$\mathsf{Encap}(\mathsf{pk})$	$\mathsf{Decap}(\mathsf{sk}', c)$
1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}$	1 : $m \leftarrow\!\!\$\ \mathcal{M}$	1 : Parse $\mathsf{sk}' = (\mathsf{sk}, \mathsf{pk}, F(\mathsf{pk}), s)$
2 : $s \leftarrow\!\!\$\ \mathcal{M}$	2 : $m \leftarrow F(m)$	2 : $m' \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$
3 : $\mathsf{pk}' \leftarrow (\mathsf{pk}, F(\mathsf{pk}))$	3 : $h \leftarrow F(pk)$	3 : $(\hat{k}', r') \leftarrow G(F(\mathsf{pk}), m')$
4 : $\mathsf{sk}' \leftarrow (\mathsf{sk}, \mathsf{pk}', s)$	4 : $(\hat{k}, r) \leftarrow G(h, m)$	4 : $c' \leftarrow \mathsf{Enc}(\mathsf{pk}, m'; r')$
5 : **return** $(\mathsf{pk}, \mathsf{sk}')$	5 : $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m; r)$	5 : **if** $c' = c$ **then**
	6 : $k \leftarrow F(\hat{k}, F(c))$	6 : **return** $F(\hat{k}', F(c))$
	7 : **return** (c, k)	7 : **else return** $F(s, F(c))$

Fig. 7. Saber uses a variant of the $\mathsf{FO}^{\not\perp}$ transform. Here G and F are hash functions [7].

these observations with the Saber team. A representative of the team [38] accepted our findings on the IND-CCA security of pSaber. Regarding Saber, they maintain that the nested hash of ciphertext $F(c)$ should not pose a security problem for Saber as c is "deterministically derived from limited entropy". However, they do not know if this allows a security proof to go through in the QROM [38].

When it comes to robustness however, the news is better. Namely, we can apply similar proof strategies used to establish strong collision-freeness of $\mathsf{FO}^{\not\perp}$-based KEMs (Theorem 6) and pSaber (Theorem 10) to show SCFR-CCA security of Saber in the QROM. The corresponding proof, presented in detail in the full version [22], on a high-level uses the fact that the hash of public-keys are included in Saber's key-derivation step (in contrast to Classic McEliece). This allows us to establish the SCFR-CCA security of Saber KEM by mainly relying on properties of quantum random oracles G and F, namely collision-resistance and claw-freeness.

Theorem 12. *For any* SCFR-CCA *adversary* $\mathcal{A}$ *against the scheme* $\mathsf{Saber.KEM} = (\mathsf{KGen}', \mathsf{Encap}, \mathsf{Decap})$ *issuing at most* q_G *(resp.* q_F*) queries to the quantum random oracle* G *(resp.* F*), we have*

$$\mathbf{Adv}^{\text{SCFR-CCA}}_{\mathsf{Saber.KEM}}(\mathcal{A}) \leq \mathbf{Coll}^{F}_{\mathsf{Saber.PKE}} + \frac{\alpha(q_G+1)^3}{2^{256}} + \frac{4\alpha(q_F+1)^3}{2^{256}} + \frac{4q_F}{2^{128}}$$

Here α (< 648) *is the constant from Lemma 1.*

Kyber uses an FO-type transform which is essentially the same as that of Saber (see Fig. 7). Hence, the issues we identified with Saber above w.r.t. IND-CCA security claims in the QROM as described in the specification document, as well as establishing anonymity of the scheme, apply to Kyber too. We have shared these observations with the Kyber team. At the 3rd NIST PQC Standardization Conference, a representative of the Kyber team [36] acknowledged that the nested hash of ciphertext $F(c)$ could make it "tricky" to prove the security of Kyber in the QROM, while removing this nested hash would overcome this issue.

But on the positive side, our result on strong collision-freeness (SCFR-CCA security) of Saber–namely, Theorem 12 above–also applies to Kyber in the same fashion, because of the similarity in their respective FO-type transforms. In other words, the current versions of Kyber and Saber also lead to strongly robust hybrid PKE schemes in the QROM.

In conclusion, we consider the *concrete* IND-CCA security–as claimed in [7, 5]–and anonymity (ANO-CCA security) of Saber and Kyber to remain open. We also suggest a modification to Saber and Kyber: namely, to apply the *same* FO-type transform as pSaber uses (as in Figure 6) to the relevant base PKE scheme, thus replacing the "nested" hashing of ciphertext in key-derivation with single hashing. In doing so, not only would the two NIST finalists then enjoy the same provable IND-CCA security guarantees of $\mathsf{FO}^{\not\perp}$-based KEMs in the QROM as established in the literature [26, 34], but this would also allow our techniques establishing anonymity of pSaber to be extended to Saber and Kyber.[5]

6 Conclusions and Future Work

In this work, we initiated the study of anonymous and robust KEMs and PKE schemes in the post-quantum setting. We resolved several core technical questions, and showed that proto-Saber, a simplified version of Saber, and FrodoKEM can be used to build anonymous, robust hybrid PKE schemes. We also pointed out gaps in the current IND-CCA security analyses of Saber and Kyber. Both NIST finalists do lead to robust hybrid PKE from our analysis. Finally, we highlighted a surprising property of Classic McEliece (CM) showing that it does not lead to robust PKE schemes via the standard KEM-DEM construction.

Important questions remain about the anonymity and robustness of the NIST finalists and alternate candidates. For example, it is plausible that the anonymity of CM could be proven by a direct approach; the same applies for Saber and Kyber. Notable among the alternate schemes is SIKE, which uses radically different algebraic problems to build a KEM; extending our work to SIKE would be interesting. One broader question about post-quantum PKE which has not been widely studied is multi-receiver hybrid PKE (with or without anonymity/robustness). Such schemes would have applications in group-oriented end-to-end secure messaging.

Acknowledgements. It is our pleasure to thank the Classic McEliece, Kyber, Saber and FrodoKEM teams, along with Kathrin Hövelmanns and Keita Xagawa, for helpful discussions. We also thank the anonymous reviewers of Eurocrypt 2022 for their constructive comments and suggestions. Paterson's research was supported in part by a gift from VMware.

[5] For Kyber's anonymity, we would rely on the hardness of module learning-with-errors (mod-LWE) problem instead of mod-LWR, akin to our discussion on FrodoKEM; see Subsection 5.3.

References

1. M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi. Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. In *CRYPTO 2005*, pages 205–222, 2005.
2. M. Abdalla, M. Bellare, and G. Neven. Robust encryption. In *TCC 2010*, pages 480–497, 2010.
3. M. R. Albrecht, D. J. Bernstein, T. Chou, C. Cid, J. Gilcher, T. Lange, V. Maram, I. von Maurich, R. Misoczki, R. Niederhagen, K. G. Paterson, E. Persichetti, C. Peters, P. Schwabe, N. Sendrier, J. Szefer, C. J. Tjhai, M. Tomlinson, and W. Wang. Classic McEliece: NIST round 3 submission, 2021.
4. E. Alkim, J. W. Bos, L. Ducas, P. Longa, I. Mironov, M. Naehrig, V. Nikolaenko, C. Peikert, A. Raghunathan, and D. Stebila. FrodoKEM: NIST round 3 submission, 2021.
5. R. Avanzi, J. Bos, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-Kyber: NIST round 3 submission, 2021.
6. A. Barth, D. Boneh, and B. Waters. Privacy in encrypted content distribution using private broadcast encryption. In *FC 2006*, pages 52–64, 2006.
7. A. Basso, J. M. B. Mera, J. D'Anvers, A. Karmakar, S. S. Roy, M. V. Beirendonck, and F. Vercauteren. Saber: NIST round 3 submission, 2021.
8. M. Bellare, A. Boldyreva, A. Desai, and D. Pointcheval. Key-privacy in public-key encryption. In *ASIACRYPT 2001*, pages 566–582, 2001.
9. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT 2006*, pages 409–426, 2006.
10. E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474, 2014.
11. D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *ASIACRYPT 2011*, pages 41–69, 2011.
12. J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT 2001*, pages 93–118, 2001.
13. C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, T. Saito, J. M. Schanck, P. Schwabe, W. Whyte, K. Xagawa, T. Yamakawa, and Z. Zhang. NTRU: NIST round 3 submission, 2021.
14. J.-P. D'Anvers, A. Karmakar, S. S. Roy, and F. Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In *AFRICACRYPT 18*, pages 282–305, 2018.
15. Y. Dodis, P. Grubbs, T. Ristenpart, and J. Woodage. Fast message franking: From invisible salamanders to encryptment. In *CRYPTO 2018, Part I*, pages 155–186, 2018.
16. P. Farshim, B. Libert, K. G. Paterson, and E. A. Quaglia. Robust encryption, revisited. In *PKC 2013*, pages 352–368, 2013.
17. P. Farshim, C. Orlandi, and R. Roşie. Security of symmetric primitives under incorrect usage of keys. *IACR Trans. Symm. Cryptol.*, 2017(1):449–473, 2017.
18. E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. In *PKC'99*, pages 53–68, 1999.

19. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *CRYPTO'99*, pages 537–554, 1999.
20. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, 2013.
21. P. Grubbs, J. Lu, and T. Ristenpart. Message franking via committing authenticated encryption. In *CRYPTO 2017, Part III*, pages 66–97, 2017.
22. P. Grubbs, V. Maram, and K. G. Paterson. Anonymous, robust post-quantum public key encryption. Cryptology ePrint Archive, Report 2021/708, 2021. https://ia.cr/2021/708.
23. R. Hayashi and K. Tanaka. PA in the two-key setting and a generic conversion for encryption with anonymity. In *ACISP 06*, pages 271–282, 2006.
24. D. Hofheinz, K. Hövelmanns, and E. Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In *TCC 2017, Part I*, pages 341–371, 2017.
25. K. Hövelmanns, E. Kiltz, S. Schäge, and D. Unruh. Generic authenticated key exchange in the quantum random oracle model. In *PKC 2020, Part II*, pages 389–422, 2020.
26. H. Jiang, Z. Zhang, L. Chen, H. Wang, and Z. Ma. IND-CCA-secure key encapsulation mechanism in the quantum random oracle model, revisited. In *CRYPTO 2018, Part III*, pages 96–125, 2018.
27. H. Jiang, Z. Zhang, and Z. Ma. Key encapsulation mechanism with explicit rejection in the quantum random oracle model. In *PKC 2019, Part II*, pages 618–645, 2019.
28. J. Len, P. Grubbs, and T. Ristenpart. Partitioning oracle attacks. In *USENIX Security*, 2021.
29. B. Libert, K. G. Paterson, and E. A. Quaglia. Anonymous broadcast encryption: Adaptive security and efficient constructions in the standard model. In *PKC 2012*, pages 206–224, 2012.
30. R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA 2011*, pages 319–339, 2011.
31. C. A. Melchor, N. Aragon, S. Bettaieb, L. Bidoux, O. Blazy, J. Bos, J. Deneuville, A. Dion, P. Gaborit, J. Lacan, E. Persichetti, J. Robert, P. Véron, and G. Zémor. HQC: NIST round 3 submission, 2021.
32. P. Mohassel. A closer look at anonymity and robustness in encryption schemes. In *ASIACRYPT 2010*, pages 501–518, 2010.
33. E. Persichetti. Secure and anonymous hybrid encryption from coding theory. In *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013*, pages 174–187, 2013.
34. T. Saito, K. Xagawa, and T. Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In *EUROCRYPT 2018, Part III*, pages 520–551, 2018.
35. K. Sako. An auction protocol which hides bids of losers. In *PKC 2000*, pages 422–432, 2000.
36. P. Schwabe. Crystals-kyber round 3 presentation. 3rd NIST PQC Standardization Conference, 2021.
37. E. E. Targhi and D. Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In *TCC 2016-B, Part II*, pages 192–216, 2016.
38. F. Vercauteren. Private communication, 2021.
39. K. Xagawa. Ntru leads to anonymous, robust public-key encryption. Cryptology ePrint Archive, Report 2021/741, 2021. https://ia.cr/2021/741.
40. M. Zhandry. Secure identity-based encryption in the quantum random oracle model. In *CRYPTO 2012*, pages 758–775, 2012.