

Robust Encryption, Revisited

Pooya Farshim¹, Benoît Libert², Kenneth G. Paterson³, and
Elizabeth A. Quaglia⁴

¹ Fachbereich Informatik, Technische Universität Darmstadt, Germany

² Technicolor, France

³ Information Security Group, Royal Holloway, University of London, UK

⁴ Département d’Informatique, École Normale Supérieure – Paris, France

Abstract. We revisit the notions of robustness introduced by Abdalla, Bellare, and Neven (TCC 2010). One of the main motivations for the introduction of strong robustness for public-key encryption (PKE) by Abdalla et al. is to prevent certain types of attack on Sako’s auction protocol. We show, perhaps surprisingly, that Sako’s protocol is still vulnerable to attacks exploiting robustness problems in the underlying PKE scheme, even when it is instantiated with a *strongly* robust scheme. This demonstrates that current notions of robustness are insufficient even for one of its most natural applications. To address this and other limitations in existing notions, we introduce a series of new robustness notions for PKE and explore their relationships. In particular, we introduce *complete* robustness, our strongest new notion of robustness, and give a number of constructions for completely robust PKE schemes.

Keywords. Robustness, Anonymity, Public-key encryption, Security proofs.

1 Introduction

A commonly pursued goal in cryptography is message privacy, which is typically achieved by means of encryption. In recent years, the privacy of users has become an equally relevant concern. It has led the research community to strive for anonymity properties when designing cryptographic primitives. In public-key encryption, in particular, *key-privacy* (a.k.a. receiver anonymity) was introduced in [4] to capture the idea that a ciphertext does not leak any information about the public key under which it was created, thereby making the communication anonymous. In this context, Abdalla, Bellare, and Neven [2] raised a fundamental question: how does a legitimate user know if an anonymous ciphertext is intended for him? Moreover, what happens if he uses his secret key on a ciphertext *not* created under his public key? To address this question, Abdalla et al. formalized a property called *robustness*, which (informally speaking) guarantees that decryption attempts fail with high probability if the “wrong” private key is used. They argued that, in all applications requiring anonymous public-key encryption, robustness is usually needed as well. These applications include auction protocols with bid privacy [23], consistency [1] in searchable encryption [7] and anonymous broadcast encryption [3,21]. As shown by Mohassel [22], robustness is also important in guaranteeing the anonymity of hybrid encryption schemes

resulting from the combination of anonymous asymmetric and symmetric components.

1.1 Robust public-key encryption

Robustness ensures that a ciphertext cannot correctly decrypt under two different secret keys. This notion has (often implicitly) been present in the literature (e.g., [23,7,9,19,3]), but formal definitions remained lacking until the recent foundational work of Abdalla et al. [2]. In particular, Abdalla et al. introduced two flavors of encryption robustness: *weak* and *strong* robustness.

Weak robustness is modeled as a game in which a winning adversary outputs a valid message M and two distinct public keys pk_0 and pk_1 such that the encryption of M under pk_0 decrypts to a valid message under sk_1 , the secret key corresponding to pk_1 . This notion is of interest since it precisely addresses the issue of *using the wrong key* that arises in anonymity contexts (such as anonymous broadcast encryption [3,21], for instance), but it is also useful in achieving the stronger notion of strong robustness.

Strong robustness—also called SROB-CCA when the adversary has access to a decryption oracle—allows for a more powerful adversary which chooses a ciphertext C (as opposed to a message which will be honestly encrypted) and two distinct public keys, and wins if C decrypts to a valid message under both corresponding secret keys. In [2] the need for this notion is motivated by scenarios where ciphertexts can be adversarially chosen. The authors of [2] give Sako’s auction protocol [23] as an example of such a situation, explaining that strong robustness is required in order to prevent an attack on the fairness of this protocol by a cheating bidder and a colluding auctioneer.

As pointed out by Abdalla et al. [2], merely appending the receiver’s public key to the ciphertext is not an option for providing robustness, since it destroys key-privacy properties. Abdalla et al. also showed that the seemingly natural solution of using an unkeyed redundancy function to modify the message before encryption does not achieve even weak robustness, thus demonstrating the non-triviality of the problem. They then gave several anonymity-preserving constructions to obtain both weak and strong robustness for public-key encryption. Using a simple tweak, they also showed how to render the Cramer-Shoup cryptosystem [12] strongly robust without introducing any overhead.

More recently, Mohassel [22] studied robustness in the context of hybrid encryption [13]. He showed that weak robustness (and not only anonymity) is needed in the asymmetric part of a hybrid encryption scheme to ensure anonymity of the overall scheme. Mohassel also considered relaxations, called *collision-freeness*, of both weak and strong robustness. He showed that many constructions in the literature are natively collision-free and showed how to generically turn any weakly (resp., strongly) collision-free scheme into a weakly (resp., strongly) robust one.

1.2 Our contributions

THE NEED FOR STRONGER DEFINITIONS. In this paper, we argue that some applications require even stronger forms of robustness than those considered in [2,22]. The first such application is, perhaps surprisingly, the construction of auction protocols with bid privacy, like that of Sako [23]. Recall that this was one of the initial motivations for analyzing robustness in [2]. Strong robustness actually turns out *not* to suffice for thwarting attacks against the fairness of Sako’s auction protocol [23]: strong robustness assumes honestly generated public keys whereas, if the auctioneer can collude with cheating bidders (as assumed in [2]), what really needs to be considered is an adversary who can maliciously generate ciphertexts *and* the public keys. To illustrate this, we show an attack on the fairness of Sako’s protocol when instantiated with \mathcal{CS}^* , a variant of the Cramer–Shoup encryption scheme which was proven to be key-private and strongly robust in [2]. This observation, then, motivates us to introduce notions of robustness where keys may be maliciously generated. We do not offer a full treatment of the delicate issue of fairness in auction protocols and its relation to robustness, since that is beyond the scope of this paper. Rather, as with [2], we use Sako’s protocol as a motivation for introducing and studying stronger robustness notions.

The limitations of existing robustness notions, and therefore the motivation for this work, are not solely restricted to Sako’s protocol. For instance, existing notions are not necessarily strong enough to provide robustness guarantees if the scheme is used to encrypt *key-dependent* messages [6] or messages encrypted under *related keys* [5]. This is because the adversary is denied access to the secret keys in these notions. The strongest of our new notions gives the adversary sufficient power and automatically provides robustness in these more challenging settings.

NEW NOTIONS OF ROBUSTNESS AND THEIR RELATIONS. Our strongest new notion is called *complete* robustness (CROB) and is obtained by progressively removing various restrictions on adversarial capabilities in the strong robustness security model. First, we give access to honestly generated secret keys and arrive at an intermediate notion which we term *unrestricted (strong) robustness* (US-ROB). Next, we also remove the honest key-generation requirement to get to the notion of *full robustness* (or FROB for short). We then view robustness in terms of the behavior of the encryption and decryption algorithms with respect to each other, and obtain our CROB notion. Roughly speaking, in CROB, the adversary should not be able to find “collisions” in the scheme beyond those which are already implied by the correctness property of the scheme. For example, he should not be able to “explain” a ciphertext C of his choice as an encryption under two different adversarially chosen public keys pk_0, pk_1 by revealing the plaintext and the encryption coins for pk_0 and the secret key sk_1 for pk_1 . As we will see, full robustness can be viewed as the “decryption-only part” of CROB. Another natural notion of robustness, which we call *key-less robustness* (KROB), arises as the dual notion corresponding to the “encryption-only part” of CROB, and is




Fig. 1. Relations among notions of robustness.

also implied by CROB. Finally, XROB is a “mixed” notion derived from FROB and KROB that has no natural interpretation but is a useful tool in establishing results about these notions.

We next study how these new notions of robustness relate to each other and to existing notions. Figure 1 summarizes the main relations that we prove between our new and existing robustness notions. In this figure, the lack of an implication between two notions should be interpreted as meaning that we prove a separation. Thus, for example, we will show that CROB is strictly stronger than FROB. It is apparent from the figure that we provide a complete account of the pairwise relations between the various robustness notions. In addition to these relations, we can prove several pairwise separations. For example, we will show that no two of the three notions from $\{\text{FROB}, \text{KROB}, \text{XROB}\}$ are sufficient to prove CROB, but that their combination is. Thus we obtain a characterization of CROB in terms of the three intermediate notions. These separations are not displayed in the figure for ease of visual presentation.

That robustness can come in so many flavors may be unsettling to some readers. Certainly, one should not seek to clutter the definitional landscape unnecessarily. Yet, with the exception of XROB, all of our notions arise as natural generalizations of the existing notions. Exploring their relations is then a natural endeavor. This is not so different from the situation for, say, confidentiality and anonymity notions for public-key encryption, where we now have many different security definitions and developing an understanding of their relations has taken several years.

CONSTRUCTIONS OF COMPLETELY ROBUST ENCRYPTION. Having defined CROB and its weaker relatives, we prove it to be achievable via a variety of efficient and natural constructions.

We first show that the generic construction for strong robustness presented in [2] is *already* powerful enough as to also achieve CROB. Further, we observe that a slight modification of this transformation allows dispensing with the weak robustness assumption—which was necessary in [2]—on the underlying PKE scheme. Moreover, we point out that the random-oracle-based generic transformation of Mohassel [22] also achieves CROB.

In the standard model, we also answer in a positive sense a question left open in [2] as to whether the Canetti–Halevi–Katz [11] (CHK) paradigm—which is known to provide chosen-ciphertext secure cryptosystems from weakly secure identity-based encryption (IBE) schemes—can be leveraged to construct systems that are simultaneously anonymous and offer message privacy under chosen-ciphertext attacks (AI-CCA security) *and* are robust in a strong sense. Answering this question is non-trivial: Abdalla et al. pinpointed that applying the one-time-signature-based CHK transformation to, say, the Boyen–Waters IBE [10] does not provide SROB-CCA or even SROB-CPA. Here, we show how to obtain AI-CCA-secure, completely robust PKE schemes from weakly secure IBE schemes. Our construction is a variant of the Boneh–Katz construction for chosen-ciphertext security [8], and it only requires the underlying IBE to satisfy a weak level of security under chosen-plaintext attacks. In comparison, the most powerful transformation of [2] must start from a scheme that is already AI-CCA-secure to achieve a comparable result. Because our technique simultaneously provides complete robustness *and* AI-CCA security, it enjoys better efficiency than applying the strongest robustness-conferring transformation of [2] to an AI-CCA-secure scheme obtained from the original Boneh–Katz transformation.

Finally, we also ask whether we can improve upon the efficiency of generic constructions with concrete schemes whose security rests on specific computational assumptions. By giving a concrete construction of a scheme that is CROB and AI-CCA-secure, we present a different and potentially more efficient way of directly achieving CROB for certain hybrid encryption schemes such as the Hofheinz–Kiltz [17] or Kurosawa–Desmedt [20] schemes. To do so, we take advantage of certain properties in the underlying symmetric components. Namely, we consider hybrid schemes that build on the *encrypt-then-MAC* paradigm in their symmetric part to obtain a suitably secure symmetric cipher. We show that, if the message authentication code (MAC) is what we call *committing*, then a simple modification in the hybrid scheme gives complete robustness without any significant computational overhead. The use of committing MACs readily extends as a tool to design AI-CCA-secure CROB hybrid constructions via the KEM/DEM framework [13]. Concretely, Mohassel [22] showed that the KEM/DEM framework gives an AI-CCA-secure hybrid encryption scheme when the KEM component is weakly robust and AI-CCA, and the DEM component is an authenticated symmetric encryption scheme. If the latter part is furthermore realized using the encrypt-then-MAC approach with a committing MAC, we easily obtain complete robustness as well. As we will see, the committing MAC technique can also offer certain advantages.

Taken altogether, our constructions achieving CROB rely on different building blocks and, when fully instantiated, allow us to rely on a variety of different hardness assumptions. They demonstrate that CROB, while providing strong guarantees, is attainable in an efficient and flexible manner.

ORGANIZATION. We start by highlighting the limitations of previous notions of robustness in Section 2. Section 3 presents our new notions. In Section 4,

we study the relations among notions of robustness. We describe our generic constructions in Section 5 and give an efficient construction in Section 6. We close by some concluding remarks in Section 7. Many details and all proofs are deferred to the full version [15].

2 Strong Robustness Does Not Suffice for Auction Protocols

Sako’s auction protocol [23] was the first practical protocol to ensure *bid privacy*, i.e., to hide the bids of losers. The basic idea is as follows. Let $V = \{v_1, \dots, v_N\}$ be the set of possible bid values. The auctioneer prepares N key-pairs $(sk_i, pk_i)_{i \in \{1, \dots, N\}}$ and publishes the N public keys. To *bid* for a value v_i a bidder encrypts a pre-determined message M under the public key pk_i . This is signed and posted by the bidder. To *open* a bid the auctioneer attempts to decrypt the encrypted bids one by one using sk_N . If at least one decrypts to M , the auctioneer publishes the winning bid v_N , a list of all the winning bidders and the secret key sk_N for the bidders to verify correctness of the result. If no decryption returns M , the auctioneer repeats the procedure using sk_{N-1} , and so on. For the auction to hide the bid values, the underlying public-key encryption scheme needs to be key-private, in the sense of [4].

In [23], Sako provided an example of an auction protocol scheme based on the ElGamal public-key encryption scheme, which is key-private. In [2], Abdalla et al. gave an attack which allows a cheating bidder and a colluding auctioneer to break the fairness of the protocol. This attack is based on the fact that the ElGamal scheme is not robust and therefore the auctioneer can open the cheating bidder’s bid to an arbitrary (winning) value. To prevent this attack, the authors of [2] suggest using any *strongly* robust scheme (strong robustness, instead of simply weak robustness, is required since the ciphertexts are generated adversarially; see [2,15] for the details).

We show that strong robustness is *not* sufficient to prevent an attack of the above type on Sako’s protocol. More precisely, in [15, Appendix C] we present an attack on the protocol when it is instantiated with a variant of the Cramer–Shoup encryption scheme, \mathcal{CS}^* , which is known to be key-private and strongly robust (the latter result was proved in [2]). Just as with the attack of Abdalla et al. [2], the attack we present assumes a cheating bidder and a colluding auctioneer. The key idea behind the attack is that an auctioneer can maliciously prepare the public keys so that the cheating bidder’s encryption decrypts to M under *any* secret key.

This attack shows that strong robustness is not enough to guarantee fairness in Sako’s auction protocol. Intuitively what is needed here is a form of robustness wherein all the public keys and ciphertexts in the system may be adversarially generated. In the coming sections we will formalize stronger notions of robustness which rule out such attacks.

3 New Notions of Robustness

3.1 A direct strengthening: full robustness

Recall that an SROB adversary has to output a ciphertext C and two public keys pk_0 and pk_1 such that C decrypts to a message M_0 under (sk_0, pk_0) and a message M_1 under (sk_1, pk_1) . The notion poses three restrictions on the adversary: (1) pk_0 and pk_1 have to be distinct; (2) The corresponding secret keys cannot have been queried by the adversary; (3) The public keys are honestly generated.

The first condition is *inherent* to modeling the behavior of an encryption scheme when used on different public keys, and removing it would make it trivial for an adversary to win.

We now look at the notion resulting from the removal of the second restriction, i.e., when the adversary is allowed to query secret keys even for the finally output public keys. We call this notion unrestricted strong robustness (USROB). This notion is powerful enough to model scenarios where keys are honestly generated, but an adversary may know the secret keys. This, for example, includes robustness for the encryption of key-dependent messages as discussed in the introduction.

However, as we have seen in the previous section, if an adversary can control the generation of keys, it may be unreasonable to assume that it can only generate the keys honestly. We therefore can strengthen USROB further by removing the third restriction on the adversary. We, however, ask the adversary to return secret keys for the public keys that it chooses. Two points deserve further attention at this point. First, returning the secret keys is to allow for a polynomial-time game definition which is not excessively strong. Second, we do not require the secret keys to be valid. Indeed, it is the responsibility of the decryption algorithm to check that the key-pair it receives is valid. Note that as a result of removing the two restrictions, the adversary has now full control over the keys, and we no longer need to provide the adversary with the oracles present in the SROB and USROB games. These modifications result in a simple, but strong, notion we call *full robustness* (FROB), and formalize in Figure 2.

<pre>proc Initialize $pars \leftarrow_s PG$ Return $pars$</pre>	<pre>proc Finalize($C, pk_0, pk_1, sk_0, sk_1$) // FROB If ($pk_0 = pk_1$) Then Return F $M_0 \leftarrow Dec(pars, pk_0, sk_0, C)$ $M_1 \leftarrow Dec(pars, pk_1, sk_1, C)$ Return ($M_0 \neq \perp$) \wedge ($M_1 \neq \perp$)</pre>
--	--

Fig. 2. Game defining full robustness.

3.2 A unified approach: complete robustness

At this point it can be asked if there are attacks which fall outside the FROB model. To answer this question, we take a somewhat different approach towards robustness and view it in terms of the behavior of the encryption and decryption routines of a scheme with respect to each other. In fact, this is the underlying intuition behind not only the original weak robustness notion,⁵ but also the standard correctness criterion for a PKE scheme (albeit for a single key). This leads us to a new notion which we term *complete robustness* (CROB). In this game the shared parameters of the system are passed to an adversary, which then arbitrarily interacts with the encryption and decryption routines on plaintexts, ciphertexts, keys, and even random coins of its choice. Its goal is to find an “unexpected collision” in the cryptosystem (i.e., one outside that imposed by the correctness criterion). We formalize the CROB game in Figure 3.

<pre>proc Initialize List ← [] pars ←_s PG Return pars</pre>	<pre>proc Enc(pk, M, r) C ← Enc(pars, pk, M; r) List ← (pk, M, C) ∪ List proc Dec(pk, sk, C) M ← Dec(pars, pk, sk, C) List ← (pk, M, C) ∪ List</pre>	<pre>proc Finalize() // CROB For (pk₀, M₀, C₀), (pk₁, M₁, C₁) ∈ List If (C₀ = C₁ ≠ ⊥) ∧ (pk₀ ≠ pk₁) ∧ (M₀ ≠ ⊥ ∧ M₁ ≠ ⊥) Return T Return F</pre>
---	---	--

Fig. 3. Game defining complete robustness.

KEY-LESS ROBUSTNESS. It can be seen through an easy inspection that full robustness is a sub-case of complete robustness where the adversary is restricted to querying the **Dec** oracle. One can also consider the dual case where the adversary only queries the **Enc** oracle. This results in a new notion which we call *key-less robustness* (KROB). Key-less robustness differs from full robustness in that an adversary no longer needs to return any secret keys, but instead “opens” a ciphertext by providing the random coins and the message used in the encryption. More precisely, the adversary outputs two messages, two distinct public keys and two sets of random coins, and its goal is to invoke a collision in the encryption algorithm. The game is shown in Figure 4.

In the next section we give a complete treatment of relations among different notions.

IDENTITY-BASED ENCRYPTION. In the IBE setting the identities (analogous to public keys in the PKE setting) are already chosen maliciously, while the natural extension of our notions would allow the adversary to also choose the IBE master keys maliciously. In particular, the identity-based analogue of FROB would be strong enough to guarantee *well-addressedness* according to the definition proposed by Hofheinz and Weinreb [18] (see also [15, Appendix B]), whereas

⁵ This then disappears in the SROB game as the adversary outputs ciphertexts.

<pre> proc Initialize $pars \leftarrow_s PG$ Return $pars$ </pre>	<pre> proc Finalize($M_0, M_1, pk_0, pk_1, r_0, r_1$) // KROB If ($pk_0 = pk_1$) Then Return F $C_0 \leftarrow \text{Enc}(pars, M_0, pk_0; r_0)$ $C_1 \leftarrow \text{Enc}(pars, M_1, pk_1; r_1)$ Return ($C_0 = C_1 \neq \perp$) </pre>
--	---

Fig. 4. Game defining key-less robustness.

SROB-CCA may not always do so. We leave the further development of the ID-based setting to future work.

4 Relations among Notions of Robustness

We now study how the various notions of robustness relate to each other. Starting with complete robustness, it may be asked if KROB and FROB are strong enough together to jointly imply CROB. We show that this is *not* the case. Indeed, there is a third “mixed” notion of robustness implicit in CROB, which we term XROB and formalize in Figure 5. As the next theorem shows, the XROB notion is necessary in the sense that it is not implied by KROB and FROB together.

In fact, no pair of the notions from {FROB, KROB, XROB} implies the third. Furthermore, the conjunction of all three notions is sufficient to imply CROB.

Theorem 1 (CROB characterization). *A PKE scheme is CROB if and only if it is simultaneously FROB, KROB, and XROB. Furthermore, no combination of at most two of FROB, KROB, and XROB is sufficient to provide the CROB guarantees.*

We prove the theorem via a sequence of propositions in [15, Appendix E].

<pre> proc Initialize $pars \leftarrow_s PG$ Return $pars$ </pre>	<pre> proc Finalize($M_0, pk_0, r_0, C_1, pk_1, sk_1$) // XROB If ($pk_0 = pk_1$) Then Return F $C_0 \leftarrow \text{Enc}(pars, M_0, pk_0; r_0)$ $M_1 \leftarrow \text{Dec}(pars, pk_1, sk_1, C_1)$ Return ($C_0 = C_1 \wedge (M_0 \neq \perp) \wedge (M_1 \neq \perp)$) </pre>
--	--

Fig. 5. Game defining mixed robustness.

As a next step we study how our new notions relate to the existing notions from Abdalla et al. [2]. Since USROB is a natural intermediate notion, for the sake of completeness, we also investigate where it stands in relation to existing notions. We start by observing that $FROB \implies USROB \implies SROB\text{-CCA}$ as the adversary becomes progressively more restricted in each game. Moreover, in the first part of the following theorem, we show that USROB is strictly stronger than SROB-CCA, and that FROB is strictly stronger than USROB. In the second part of the theorem we show that KROB does not even imply WROB-CPA,

separating this notion from all notions other than complete robustness. Finally, we show XROB implies WROB-CCA but *not* SROB-CPA. Hence XROB can be seen a strengthened version of weak robustness in a direction orthogonal to strong robustness.

Theorem 2 (Relation with WROB and SROB). *Let \mathcal{PKE} be a PKE scheme. We have the following.*

- **FROB:** *If \mathcal{PKE} is FROB, then it is also USROB. If \mathcal{PKE} is USROB then it is also SROB-CCA. Moreover, these implications are strict.*
- **KROB:** *KROB does not imply WROB-CPA and SROB-CCA does not imply KROB.*
- **XROB:** *If \mathcal{PKE} is XROB, then it is also WROB-CCA. Furthermore, XROB does not imply SROB-CPA and SROB-CCA does not imply XROB.*

We prove the theorem in [15, Appendix F]. The results of [2] together with Theorems 1 and 2 resolve all the relations between any pair of robustness notions as we have summarized in Figure 1. For example, to see that $\text{KROB} \not\Rightarrow \text{FROB}$, we use the facts that $\text{FROB} \Rightarrow \text{SROB-ATK}$ but $\text{KROB} \wedge \text{XROB} \not\Rightarrow \text{SROB-ATK}$. Moreover, although we do not formally prove it here, all our separating examples are designed such that they preserve the AI-ATK security of the underlying PKE schemes. Hence Figure 1 also applies in the presence of AI-ATK security.

5 Generic Constructions of Completely Robust Public-Key Encryption

5.1 Mohassel’s transformation

Mohassel [22] gives a generic transformation in the random-oracle model that converts an AI-ATK encryption scheme into one which is SROB-CCA without compromising its AI-ATK security. This construction also achieves complete robustness. In this construction, the hash value $H(pk, r, M)$, where r is the randomness used in the encryption, is attached to ciphertexts. This immediately rules out all forms of collisions between ciphertexts, as the hash values are unlikely to collide on two distinct public keys.

5.2 The ABN transformation

In [2, Theorem 4.2] the authors give a generic construction for a scheme $\overline{\mathcal{PKE}}$ which confers strong robustness and preserves the AI-ATK security of the starting scheme \mathcal{PKE} , provided that the latter scheme is additionally WROB. We briefly describe how the transformation works, and refer the reader to the original work for the details. At setup, include in *pars* for $\overline{\mathcal{PKE}}$ the parameters of a commitment scheme (see [15, Appendix G] for the definitions). When encrypting, commit to the public key, and encrypt the *de-commitment* along with the

message. Also include the commitment as a ciphertext component. Decryption checks the commitment/de-commitment pair for consistency and rejects if this is not the case. We strengthen the result of [2], showing that this construction achieves complete robustness:

Theorem 3 (The ABN transformation achieves CROB). *Let \mathcal{A} be a PPT CROB adversary against $\overline{\mathcal{PK}\mathcal{E}}$. Then there exist PPT adversaries \mathcal{B}_1 , \mathcal{B}_2 , and \mathcal{B}_3 against the binding property of \mathcal{CMT} such that*

$$\mathbf{Adv}_{\overline{\mathcal{PK}\mathcal{E}}}^{\text{croB}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{CMT}}^{\text{bind}}(\mathcal{B}_1) + \mathbf{Adv}_{\mathcal{CMT}}^{\text{bind}}(\mathcal{B}_2) + \mathbf{Adv}_{\mathcal{CMT}}^{\text{bind}}(\mathcal{B}_3).$$

The proof of this theorem is given in [15, Appendix H], where we show scheme $\overline{\mathcal{PK}\mathcal{E}}$ is simultaneously FROB, KROB, and XROB.

5.3 A modification of the ABN transformation

While the original transformation [2] *does* provide AI-ATK and CROB guarantees, the AI-ATK security of the transformed scheme $\overline{\mathcal{PK}\mathcal{E}}$ relies on the weak robustness of the underlying encryption scheme $\mathcal{PK}\mathcal{E}$ in the case of chosen-ciphertext adversaries (i.e., when $\text{ATK} = \text{CCA}$). We show that, if the underlying encryption scheme supports labels [24] (in which case the encryption and decryption algorithms both take an additional public string L as input; see [15, Appendix A]), this assumption can be eliminated and we only need $\mathcal{PK}\mathcal{E}$ to be AI-ATK-secure.

Although the weak robustness assumption is not too demanding in theory (since any encryption scheme can be made weakly robust by means of a keyed redundancy-based transformation [2]), our construction provides better efficiency in some settings since many AI-CCA encryption schemes, such as the Cramer–Shoup or the Kurosawa–Desmedt scheme, natively support labels.⁶

Our transformation, which relies on a commitment scheme \mathcal{CMT} consisting of algorithms (CPG, Com, Ver), is as follows.

$\overline{\text{PG}}(1^\lambda)$: Run $\text{pars} \leftarrow_{\$} \text{PG}(1^\lambda)$ to obtain public parameters pars for $\mathcal{PK}\mathcal{E}$. Then, generate $\text{cpars} \leftarrow_{\$} \text{CPG}(1^\lambda)$ for \mathcal{CMT} . Finally, return $(\text{pars}, \text{cpars})$.

$\overline{\text{KG}}(\text{pars}, \text{cpars})$: Compute and return $(sk, pk) \leftarrow_{\$} \text{KG}(\text{pars})$.

$\overline{\text{Enc}}((\text{pars}, \text{cpars}), pk, M)$: The algorithm proceeds in two steps.

1. Commit to pk by computing a pair $(com, dec) \leftarrow_{\$} \text{Com}(\text{cpars}, pk)$.
2. Encrypt $M || dec$ under the label $L = com$ by setting the ciphertext C to be $\text{Enc}(\text{pars}, pk, M || dec, L)$.

Return (C, com) as the final ciphertext.

$\overline{\text{Dec}}((\text{pars}, \text{cpars}), pk, sk, (C, com))$: The algorithm proceeds in two steps.

⁶ In the worst case, labeled public-key encryption schemes can always be obtained by appending the label to the encrypted plaintext and checking whether the correct label is recovered at decryption.

1. Compute $M' \leftarrow \text{Dec}(pars, pk, sk, (C, com), L)$, with $L = com$. Then, parse M' as $M||dec$ (and return \perp if M' cannot be parsed properly).
2. Return M if $\text{Ver}(cpars, pk, com, dec) = 1$. Else return \perp .

Theorem 4, whose proof is in [15, Appendix I], shows that thanks to the use of labels, we do not have to rely on any weaker form of robustness of \mathcal{PKE} when proving the AI-ATK security of $\overline{\mathcal{PKE}}$.

Theorem 4. *If \mathcal{PKE} is AI-ATK-secure and \mathcal{CMT} is a hiding commitment, then $\overline{\mathcal{PKE}}$ is AI-ATK-secure. More precisely, for any PPT AI-ATK adversary \mathcal{A} against $\overline{\mathcal{PKE}}$, there exists a PPT AI-ATK adversary \mathcal{B}_1 against \mathcal{PKE} and a PPT distinguisher \mathcal{B}_2 against \mathcal{CMT} such that*

$$\text{Adv}_{\overline{\mathcal{PKE}}}^{\text{ai-atk}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\mathcal{PKE}}^{\text{ai-atk}}(\mathcal{B}_1) + \text{Adv}_{\mathcal{CMT}}^{\text{hide}}(\mathcal{B}_2).$$

Furthermore, the above construction is CROB if \mathcal{CMT} is a binding commitment. More precisely, for any PPT CROB adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} against the binding property of the commitment scheme such that

$$\text{Adv}_{\overline{\mathcal{PKE}}}^{\text{croB}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{CMT}}^{\text{bind}}(\mathcal{B}).$$

5.4 Completely robust AI-CCA-secure PKE from selectively secure IBE

Next, we present a modification of the Boneh–Katz approach [8] which provides both CROB and AI-CCA security when applied to any IBE scheme that only provides TA anonymity in the multi-authority selective-ID setting (or sID-TAA-CPA security, as defined in [15, Appendix J]). In particular, this positively answers the question of whether CHK-like techniques can be used to achieve a strong flavor of robustness from weakly secure IBE.

Let \mathcal{IBE} be an sID-TAA-CPA secure IBE scheme. We obtain a completely robust AI-CCA-secure public-key encryption scheme $\overline{\mathcal{PKE}}$ by combining \mathcal{IBE} with a strongly secure message authentication code \mathcal{MAC} and a trapdoor commitment scheme \mathcal{TCMT} .

Recall that a trapdoor commitment scheme \mathcal{TCMT} consists of efficient algorithms (CPG, Com, Ver, Equiv) where (CPG, Com, Ver) function as in an ordinary commitment except that CPG outputs public parameters $cpars$ and a trapdoor td . In addition, Equiv allows equivocating a commitment using the trapdoor td : for any two messages m_1, m_2 and any tuple (com, dec_1) produced as $(com, dec_1) \leftarrow_{\$} \text{Com}(cpars, m_1)$, the trapdoor td allows computing the value $dec_2 \leftarrow_{\$} \text{Equiv}(td, com, m_1, dec_1, m_2)$ such that $\text{Ver}(cpars, com, m_2, dec_2) = 1$. Moreover, (com, dec_2) has the same distribution as $\text{Com}(cpars, m_2)$.

Our IBE-based construction $\overline{\mathcal{PKE}} = (\overline{\text{PG}}, \overline{\text{KG}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ is as follows.

$\overline{\text{PG}}(1^\lambda)$: Run $pars \leftarrow_{\$} \mathcal{IBE}.\text{PG}(1^\lambda)$ to obtain common public parameters $pars$. Also run $cpars \leftarrow_{\$} \text{CPG}(1^\lambda)$ to obtain public parameters for a trapdoor commitment scheme \mathcal{TCMT} . Then, choose a message authentication code \mathcal{MAC} with key length $\ell \in \text{poly}(\lambda)$. Finally, return $(pars, cpars, \mathcal{MAC})$.

$\overline{\text{KG}}(pars, cpars, \mathcal{MAC})$: Generate $(msk, mpk) \leftarrow_s \text{IBE.MPG}(pars)$ for IBE . Return the key pair $(sk, pk) := (msk, mpk)$.

$\overline{\text{Enc}}((pars, cpars, \mathcal{MAC}), pk, M)$: To encrypt M under $pk = mpk$, the algorithm proceeds as follows.

1. Choose a random MAC key $k \leftarrow_s \{0, 1\}^\ell$.
2. Commit to $mpk\|k$ by computing $(com, dec) \leftarrow_s \text{Com}(cpars, mpk\|k)$.
3. Encrypt $M\|k\|dec$ under the identity com by setting C to the output of $\text{IBE.Enc}(pars, mpk, com, M\|k\|dec)$.
4. Compute $tag = \text{MacGen}_k(C)$ and return (C, com, tag) as the final ciphertext.

$\overline{\text{Dec}}((pars, cpars, \mathcal{MAC}), pk, sk, (C, com, tag))$: Given $pk = mpk$ and $sk = msk$, conduct the following steps.

1. Compute $dk_{com} \leftarrow_s \text{IBE.KG}(pars, msk, com)$ and then set M' to be $\text{IBE.Dec}(pars, mpk, dk_{com}, com, C)$. Then, parse M' as $M\|k\|dec$ (and return \perp if $M' = \perp$ or if M' cannot be parsed properly).
2. If $\text{MacVer}_k(C, tag) = 1$ and $\text{Ver}(cpars, mpk\|k, com, dec) = 1$, return M . Otherwise, return \perp .

A difference with the original Boneh–Katz construction—which can use a weak form of commitment called *encapsulation*—is that our construction requires a full-fledged commitment scheme. This is because, in order to achieve complete robustness, we need to commit to the master public key of the scheme at the same time as the MAC key in the encryption algorithm. Moreover, the proof of AI-CCA security requires the commitment to be a *trapdoor* commitment: the trapdoor plays an essential role when we reduce the sID-TAA-CPA security of the IBE to the AI-CCA security of the encryption scheme.

The proof of the following theorem can be found in [15, Appendix J].

Theorem 5. *If IBE is sID-TAA-CPA-secure, \mathcal{MAC} is strongly unforgeable, and TCMT is a computationally binding trapdoor commitment scheme, then $\overline{\text{PKE}}$ is AI-CCA-secure. Moreover, the scheme $\overline{\text{PKE}}$ is CROB if TCMT is computationally binding.*

6 A Concrete CROB Scheme

In this section, we describe a simple way to achieve complete robustness using hybrid encryption where the symmetric component uses the encrypt-then-MAC approach. To this end, we require the MAC to satisfy a “MAC analogue” of the notion of committing symmetric encryption [16]. Informally this notion requires that a given MAC tag is valid for a single message regardless of the key used.

COMMITTING MAC. We say $\mathcal{MAC} = (\text{MacGen}, \text{MacVer})$ is *committing* if for any message m and any key k , there exists no message-key pair (m', k') such that $m' \neq m$ and $\text{MacVer}_{k'}(m', \text{MacGen}_k(m)) = 1$.

We also need the MAC to computationally hide the message. Note that the following definition is implied by the definition of message-hiding security used in [14, Definition 2.2].

INDISTINGUISHABLE MAC. We say a message authentication code $\mathcal{MAC} = (\text{MacGen}, \text{MacVer})$ with key space KSp provides *indistinguishability* if, for any two messages m_0, m_1 , it is computationally infeasible to distinguish the distributions $\mathcal{D}_b := \{\text{tag} \leftarrow_s \text{MacGen}_k(m_b) : k \leftarrow_s \text{KSp}\}$ for $b \in \{0, 1\}$.

For our purposes, the MAC only has to provide one-time strong unforgeability. Namely, the adversary is allowed to see one pair of the form (m, tag) , where $\text{tag} = \text{MacGen}_k(m)$, and should not be able to produce a pair (m', tag') such that $(m', \text{tag}') \neq (m, \text{tag})$ and $\text{MacVer}_k(m', \text{tag}') = 1$.

Using ideas from [16], it is easy to construct a MAC which is simultaneously committing, indistinguishable, and strongly unforgeable. The idea is to use a family of *injective* and *key-binding* pseudorandom functions: for any distinct keys k_1, k_2 , the functions $f_{k_1}(\cdot)$ and $f_{k_2}(\cdot)$ have disjoint ranges, i.e., there exist no two pairs $(k_1, x_1), (k_2, x_2)$ such that $k_1 \neq k_2$ and $f_{k_1}(x_1) = f_{k_2}(x_2)$. The key space of the MAC is that of the PRF. For any message $m \neq 1^\lambda$, the MAC generation computes and outputs the pair $(f_k(1^\lambda), f_k(m))$. The first component serves as a perfectly binding commitment to the key k while the injectivity of $f_k(\cdot)$ guarantees that the MAC is only valid for one message. In addition, its strong unforgeability and indistinguishability properties are both implied by the pseudorandomness of $\{f_k\}_k$ as long as the message space of the MAC, MSp^{mac} , does not include 1^λ (the proof is straightforward).

We show a simple variant of the Hofheinz–Kiltz (HK) hybrid encryption scheme [17] that provides CROB and AI-CCA security when the underlying authenticated symmetric encryption scheme uses a MAC with the aforementioned properties. Besides providing new ways to achieve robustness, our scheme comes with the advantage that its computational efficiency is the same as the original HK scheme and in particular it is more efficient than combining HK with a commitment using the ABN transformation.

PG(1^λ): Choose a group \mathbb{G} of prime order $p > 2^\lambda$ with $g \leftarrow_s \mathbb{G}$. Also, choose a symmetric encryption scheme (E, D) of key length ℓ_0 and a message authentication code $\mathcal{MAC} = (\text{MacGen}, \text{MacVer})$ of key length ℓ_1 . Finally, choose a key-derivation function $\text{KDF} : \mathbb{G} \rightarrow \{0, 1\}^{\ell_0 + \ell_1}$, a target collision-resistant hash function⁷ $\text{TCR} : \mathbb{G} \rightarrow \mathbb{Z}_p$, and a collision-resistant hash function $H : \{0, 1\}^* \rightarrow \text{MSp}^{\text{mac}}$, where MSp^{mac} is the message space of \mathcal{MAC} . The public parameters consist of $\text{pars} := (\mathbb{G}, p, g, (\text{E}, \text{D}), \mathcal{MAC}, \text{TCR}, \text{KDF}, H)$.

KG(pars): Choose $x, y, z \leftarrow_s \mathbb{Z}_p^*$ and compute $u = g^x$, $v = g^y$, and $h = g^z$. The public key is $pk = (u, v, h)$ and the private key is $sk = (x, y, z) \in (\mathbb{Z}_p^*)^3$.

Enc(pars, pk, M): Choose $s \leftarrow_s \mathbb{Z}_p^*$ and compute

$$C_1 = g^s, C_2 = (u^\tau \cdot v)^s, C_3 \leftarrow_s \text{E}_{K_0}(M), \text{tag} = \text{MacGen}_{K_1}(H(C_3, u, v, h))$$

where $\tau = \text{TCR}(C_1) \in \mathbb{Z}_p^*$ and $(K_0, K_1) = \text{KDF}(h^s) \in \{0, 1\}^{\ell_0 + \ell_1}$. Return $C = (C_1, C_2, C_3, \text{tag})$.

⁷ As in [17], this function can be replaced by an injective encoding from \mathbb{G} to \mathbb{Z}_p .

$\text{Dec}(\text{pars}, pk, sk, C)$: Given $C = (C_1, C_2, C_3, \text{tag})$, return \perp if $C_2 \neq C_1^{\tau \cdot x + y}$, where $\tau = \text{TCR}(C_1)$. Else, compute $(K_0, K_1) = \text{KDF}(C_1^z)$ and $M \leftarrow \text{D}_{K_0}(C_3)$. Return M if $\text{MacVer}_{K_1}(H(C_3, pk), \text{tag}) = 1$. Else, return \perp .

The scheme was known to be IND-CCA-secure. We are also able to prove that it provides AI-CCA security, essentially because the ciphertexts can be shown to be indistinguishable from dummy ciphertexts that are statistically independent of the public key, even in the presence of a decryption oracle. Proofs of the following results may be found in [15, Appendix K].

Theorem 6. *The scheme provides AI-CCA security assuming that: (1) The DDH assumption holds in \mathbb{G} ; (2) (E, D) is a semantically secure symmetric encryption scheme; (3) KDF is a secure key-derivation function;⁸ (4) MAC is one-time strongly unforgeable and provides indistinguishability; (5) H and TCR are collision-resistant and target collision-resistant, respectively. Furthermore, the scheme is CROB if H is collision-resistant and MAC is committing.*

Interestingly, if the construction of Section 5.4 is modified to use a committing MAC, it can be instantiated using any commitment scheme and in particular a perfectly binding commitment or even an encapsulation scheme (as in the original Boneh–Katz construction) also work. In this case, the sender no longer needs to commit to the master public key: (com, dec) is generated by committing to the MAC key only. Instead, the sender computes tag as $\text{tag} = \text{MacGen}_k(H(C, \text{mpk}))$ using a collision-resistant hash function H . If the MAC is committing, the resulting construction is easily seen to provide complete robustness. It also remains AI-CCA-secure provided the MAC satisfies the notion of indistinguishability.

7 Closing Remarks

Motivated in part by the shortcomings of existing definitions of robustness, we have made a thorough exploration of the landscape of robustness definitions and their relations, and given a suite of flexible and efficient methods for obtaining completely robust AI-CCA-secure public-key encryption schemes. In future work, one could explore the situation in the ID-based setting. Another open question, well beyond the remit of this paper, is to formalize the fairness of auctions and formally prove that our CROB notion is strong enough to ensure this property for Sako’s protocol or its variants.

Acknowledgments. The authors would like to thank Mihir Bellare for his valuable comments. Pooya Farshim is supported by grant Fi 940/4-1 of the German Research Foundation (DFG). Part of this work was done while Benoît Libert was an F.R.S.-F.N.R.S. scientific collaborator at the Université catholique de Louvain (Belgium). Kenneth G. Paterson was supported by an EPSRC Leadership Fellowship, EP/H005455/1.

⁸ The standard KDF security requires that no distinguisher can tell if it is given the output of the KDF for a random input or just a random element in the range of the KDF.

References

1. Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi. Searchable encryption revisited: consistency properties, relation to anonymous IBE, and extensions. *Journal of Cryptology*, 21(3):350–391, 2008.
2. Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. In Daniele Micciancio, editor, *TCC 2010*, vol. 5978 of *LNCS*, pp. 480–497. Springer, 2010.
3. Adam Barth, Dan Boneh, and Brent Waters. Privacy in encrypted content distribution using private broadcast encryption. In Giovanni Di Crescenzo and Aviel D. Rubin, editors, *Financial Cryptography 2006*, vol. 4107 of *LNCS*, pp. 52–64. Springer, 2006.
4. Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *ASIACRYPT 2001*, vol. 2248 of *LNCS*, pp. 566–582. Springer, 2001.
5. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham, editor, *EUROCRYPT 2003*, vol. 2656 of *LNCS*, pp. 491–506. Springer, 2003.
6. John Black, Phillip Rogaway, and Thomas Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In Kaisa Nyberg and Howard M. Heys, editors, *Selected Areas in Cryptography 2002*, vol. 2595 of *LNCS*, pp. 62–75. Springer, 2002.
7. Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. Public key encryption with keyword search. In Ronald Cramer, editor, *EUROCRYPT 2004*, vol. 3027 of *LNCS*, pp. 506–522. Springer, 2004.
8. Dan Boneh and Jonathan Katz. Improved efficiency for CCA-secure cryptosystems built using identity-based encryption. In Alfred Menezes, editor, *CT-RSA 2005*, vol. 3376 of *LNCS*, pp. 87–103. Springer, 2005.
9. Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In Serge Vaudenay, editor, *EUROCRYPT 2006*, vol. 4004 of *LNCS*, pp. 535–554. Springer, 2006.
10. Xavier Boyen and Brent Waters. Anonymous hierarchical identity-based encryption (without random oracles). In Cynthia Dwork, editor, *CRYPTO 2006*, vol. 4117 of *LNCS*, pp. 290–307. Springer, 2006.
11. Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, vol. 3027 of *LNCS*, pp. 207–222. Springer, 2004.
12. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *CRYPTO'98*, vol. 1462 of *LNCS*, pp. 13–25. Springer, 1998.
13. Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33:167–226, 2003.
14. Yevgeniy Dodis, Eike Kiltz, Krzysztof Pietrzak, and Daniel Wichs. Message authentication, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, vol. 7237 of *LNCS*, pp. 355–374. Springer, 2012.
15. Pooya Farshim, Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Robust encryption, revisited. Cryptology ePrint Archive, Report 2012/673, 2012. <http://eprint.iacr.org/>. Full version of this paper.

16. Marc Fischlin. Pseudorandom function tribe ensembles based on one-way permutations: Improvements and applications. In Jacques Stern, editor, *EUROCRYPT'99*, vol. 1592 of *LNCS*, pp. 432–445. Springer, 1999.
17. Dennis Hofheinz and Eike Kiltz. Secure hybrid encryption from weakened key encapsulation. In Alfred Menezes, editor, *CRYPTO 2007*, vol. 4622 of *LNCS*, pp. 553–571. Springer, 2007.
18. Dennis Hofheinz and Enav Weinreb. Searchable encryption with decryption in the standard model. Cryptology ePrint Archive, Report 2008/423, 2008. <http://eprint.iacr.org/>.
19. Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In Nigel P. Smart, editor, *EUROCRYPT 2008*, vol. 4965 of *LNCS*, pp. 146–162. Springer, 2008.
20. Kaoru Kurosawa and Yvo Desmedt. A new paradigm of hybrid encryption scheme. In Matthew K. Franklin, editor, *CRYPTO 2004*, vol. 3152 of *LNCS*, pp. 426–442. Springer, 2004.
21. Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Anonymous broadcast encryption: adaptive security and efficient constructions in the standard model. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *Public Key Cryptography 2012*, vol. 7293 of *LNCS*, pp. 206–224. Springer, 2012.
22. Payman Mohassel. A closer look at anonymity and robustness in encryption schemes. In Masayuki Abe, editor, *ASIACRYPT 2010*, vol. 6477 of *LNCS*, pp. 501–518. Springer, 2010.
23. Kazue Sako. An auction protocol which hides bids of losers. In Hideki Imai and Yuliang Zheng, editors, *Public Key Cryptography 2000*, vol. 1751 of *LNCS*, pp. 422–432. Springer, 2000.
24. Victor Shoup. A proposal for an ISO standard for public key encryption (version 2.1). Cryptology ePrint Archive, Report 2001/112, 2001. <http://eprint.iacr.org/>.