

Key Rotation for Authenticated Encryption

Adam Everspaugh¹, Kenneth Paterson², Thomas Ristenpart³, Sam Scott²

¹University of Wisconsin–Madison, ²Royal Holloway, University of London,
³Cornell Tech

Abstract. A common requirement in practice is to periodically rotate the keys used to encrypt stored data. Systems used by Amazon and Google do so using a hybrid encryption technique which is eminently practical but has questionable security in the face of key compromises and does not provide full key rotation. Meanwhile, symmetric updatable encryption schemes (introduced by Boneh *et al.* CRYPTO 2013) support full key rotation without performing decryption: ciphertexts created under one key can be rotated to ciphertexts created under a different key with the help of a re-encryption token. By design, the tokens do not leak information about keys or plaintexts and so can be given to storage providers without compromising security. But the prior work of Boneh *et al.* addresses relatively weak confidentiality goals and does not consider integrity at all. Moreover, as we show, a subtle issue with their concrete scheme obviates a security proof even for confidentiality against passive attacks.

This paper presents a systematic study of *updatable Authenticated Encryption (AE)*. We provide a set of security notions that strengthen those in prior work. These notions enable us to tease out real-world security requirements of different strengths and build schemes that satisfy them efficiently. We show that the hybrid approach currently used in industry achieves relatively weak forms of confidentiality and integrity, but can be modified at low cost to meet our stronger confidentiality and integrity goals. This leads to a practical scheme that has negligible overhead beyond conventional AE. We then introduce *re-encryption indistinguishability*, a security notion that formally captures the idea of fully refreshing keys upon rotation. We show how to repair the scheme of Boneh *et al.*, attaining our stronger confidentiality notion. We also show how to extend the scheme to provide integrity, and we prove that it meets our re-encryption indistinguishability notion. Finally, we discuss how to instantiate our scheme efficiently using off-the-shelf cryptographic components (AE, hashing, elliptic curves). We report on the performance of a prototype implementation, showing that fully secure key rotations can be performed at a throughput of approximately 139 kB/s.

1 Introduction

To cryptographically protect data while stored, systems use authenticated encryption (AE) schemes that provide strong message confidentiality as well as ciphertext integrity. The latter allows detection of active attackers who manipulate ciphertexts. When data is stored for long periods of time, good key management practice dictates that systems must support key rotation: moving encrypted data from an old key to a fresh one. Indeed, key rotation is mandated by regulation in some contexts, such as the payment card industry data security standard (PCI DSS) that dictates how credit card data must be secured [PCI16]. Key rotation can also be used to revoke old keys that are comprised, or to effect data access revocation.

Deployed approaches to key rotation. Systems used in practice typically support a type of key rotation using a symmetric key hierarchy. Amazon’s Key Management Service [AWS], for example, enables users to encrypt a plaintext M under a fresh data encapsulation key via $C_{dem} = \text{Enc}(K_d, M)$ and then wrap K_d via $C_{kem} = \text{Enc}(K, K_d)$ under a long-term key K owned by the client. Here Enc is an authenticated encryption (AE) scheme. By analogy with the use of hybrid encryption in the asymmetric setting, we refer to such a scheme as a KEM/DEM construction, with KEM and DEM standing for key and data encapsulation mechanisms, respectively; we refer to the specific scheme as AE-hybrid.

The AE-hybrid scheme then allows a simple form of key rotation: the client picks a fresh K' and re-encrypts K_d as $C'_{kem} = \text{Enc}(K', \text{Dec}(K, C_{kem}))$. Note that the DEM key K_d does not change during key rotation. When deployed in a remote storage system, a client can perform key rotation just by fetching from the server the small, constant-sized ciphertext C_{kem} , operating locally on it to produce C'_{kem} , and then sending C'_{kem} back to the server. Performance is independent of the actual message length. The Google Cloud Platform [Goo] uses a similar approach to enable key rotation.

To our knowledge, the level of security provided by this widely deployed AE-hybrid scheme has never been investigated, let alone formally defined in a security model motivated by real-world security considerations. It is even arguable whether AE-hybrid truly rotates keys, since the DEM key does not change. Certainly it is unclear what security is provided if key compromises occur, one of the main motivations for using such an approach in the first place. On the other hand, the scheme is fast and requires only limited data transfer between the client and the data store, and appears to be sufficient to meet current regulatory requirements.

Updatable encryption. Boneh, Lewi, Montgomery, and Raghunathan (BLMR) [BLMR15] (the full version of [BLMR13]) introduced another approach to enabling key rotation that they call *updatable encryption*. An updatable encryption scheme is a symmetric encryption scheme that, in addition to the usual triple of (KeyGen, Enc, Dec) algorithms, comes with a pair of algorithms ReKeyGen and ReEnc. The first, ReKeyGen, generates a compact rekey token given the old and new secret keys and a target ciphertext, while the second, ReEnc, uses a rekey token output by the first to rotate the ciphertext without performing decryption. For example, AE-hybrid can be seen as an instance of an updatable encryption scheme in which the rekey token output by ReKeyGen is C'_{kem} and where ReEnc simply replaces C_{kem} with C'_{kem} . BLMR introduced an IND-CPA-style security notion in which adversaries can additionally obtain some rekey tokens. Their definition is inspired by, but different from, those used for CCA-secure proxy re-encryption schemes [CH07]. Given its obvious limitations when it comes to key rotation, it is perhaps surprising that the AE-hybrid construction provably meets the BLMR confidentiality notion for updatable encryption schemes.

BLMR also introduced and targeted a second security notion for updatable encryption, called ciphertext independence. It demands that a ciphertext and its rotation to another key are identically distributed to a ciphertext and a rotation of another ciphertext (for the same message). The intuition is that this captures the idea that true key rotation should refresh all randomness used during encryption. This definition is *not* met by the AE-hybrid construction above. But it is both unclear what attacks meeting their definition would prevent, and, relatedly, whether more intuitive definitions exist.

BLMR gave a construction for an updatable encryption scheme and claimed that it provably meets their two security definitions. Their construction cleverly combines an IND-CPA KEM with a DEM that uses a key-homomorphic PRF [NPR99, BLMR15] to realize a stream cipher. This enables rotation of both the KEM and the DEM keys, though the latter requires a number of operations that is linear in the plaintext length. Looking ahead, their proof sketch has a bug and we provide strong evidence that it is unlikely to be fixable. Moreover, BLMR do not yet target or achieve any kind of authenticated encryption goal, a must for practical use.

Our contributions. We provide a systematic treatment of AE schemes that support key rotation without decryption, a.k.a. updatable AE.

Specifically, we provide a new security notion for confidentiality, UP-IND, that is strictly stronger than that of BLMR [BLMR15], a corresponding notion for integrity, UP-INT (missing entirely from BLMR but essential for practice), and a new notion called re-encryption indistinguishability (UP-REENC) that is strictly stronger and more natural in capturing the spirit of “true key rotation” than the ciphertext indistinguishability notion of BLMR.

Achieving our UP-REENC notion means that an attacker, having access to both a ciphertext and the secret key used to generate it, should not be able to derive any information that helps it attack a rotation of that ciphertext. Thus, for example, an insider with access to the encryption keys at some point in time but who is then excluded from the system cannot make use of the old keys to learn anything useful once key rotation has been carried out on the AE ciphertexts. Teasing out the correct form of this notion turns out to be a significant challenge in our work.

Armed with this set of security notions, we go on to make better sense of the landscape of constructions for updatable AE schemes. Table 1 summarises the security properties of the different schemes that we consider. Referring to this table, our security notions highlight the limitations of the AE-hybrid scheme: while it meets the confidentiality notion of BLMR, it only satisfies our UP-IND and UP-INT notions when considering a severely weakened adversary who has no access to any compromised keys. We propose an improved construction, KSS, that satisfies both notions for any number of compromised keys and which is easily deployable via small adjustments to AE-hybrid. KSS uses a form of secret sharing to embed key shares in the KEM and DEM components to avoid the issue of leaking the DEM key in the updating process, and adds a cryptographic hash binding the KEM and DEM components to prevent mauling attacks. These

Scheme	Section	Tokens	CT dep.	UP-IND	UP-INT	UP-REENC
AE-hybrid [†]	4.1	uni-dir.	depend.	✗	✗	✗
KSS*	4.3	uni-dir.	depend.	✓	✓	✗
XOR-KEM*	A.1	bi-dir.	indep.	✓	✗	✗
BLMR	6	uni-dir.	depend.	✗	✗	✗
ReCrypt*	7	uni-dir.	depend.	✓	✓	✓

Table 1. Summary of schemes studied. [†] In-use by practitioners today. * Introduced in this work.

changes could easily be adopted by practitioners with virtually no impact on performance, while concretely improving security.

However, the improved scheme KSS cannot satisfy our UP-REENC notion, because it still uses a KEM/DEM-style approach in which the DEM key is never rotated. The BLMR scheme might provide UP-REENC security, but, as noted above, its security proof contains a bug which we consider unlikely to be fixable. Indeed, we show that proving the BLMR scheme confidential would imply that one could also prove circular security [BRS03,CL01] for a particular type of hybrid encryption scheme assuming only the key encapsulation is IND-CPA secure. Existing counter-examples of IND-CPA secure, but circular insecure, schemes [ABBC10,CGH12] do not quite rule out such a result. But the link to the very strong notion of circular security casts doubt on the security of this scheme. One can easily modify the BLMR scheme to avoid this issue, but even having done so the resulting encryption scheme is still trivially malleable and so cannot meet our UP-INT integrity notion.

We therefore provide another new scheme, ReCrypt, meeting all three of our security notions: UP-IND, UP-INT and UP-REENC. We take inspiration from the previous constructions, especially that of BLMR: key-homomorphic PRFs provide the ability to fully rotate encryption keys; the KEM/DEM approach with secret sharing avoids the issue of leaking the DEM key in the updating process; and finally, adding a cryptographic hash to the KEM tightly binds the KEM and DEM portions and prevents ciphertext manipulation. We go on to instantiate the scheme using the Random Oracle Model (ROM) key-homomorphic PRF from [NPR99], having the form $H(M)^k$, where H is a hash function into a group in which DDH is hard. This yields a construction of an updatable AE scheme meeting all three of our security notions in the ROM under the DDH assumption. We report on the performance of an implementation of ReCrypt using elliptic curve groups, concluding that it is performant enough for practical use with short plaintexts. However, because of its reliance on exponentiation, ReCrypt is still orders of magnitude slower than our KSS scheme (achieving only UP-IND and UP-INT security). This, currently, is the price that must be paid for true key rotation in updatable encryption.

Summary. In summary, the main contributions of this paper are:

- To provide the first definitions of security for AE supporting key rotation without exposing plaintext.
- To explain the gap between existing, deployed schemes using the KEM/DEM approach and “full” refreshing of ciphertexts.
- To provide the first proofs of security for AE schemes using the KEM/DEM approach, namely AE-hybrid and KSS.
- To detail the first updatable AE scheme, ReCrypt, that fully and securely refreshes ciphertexts by way of key rotations without exposing plaintext. We implement a prototype and report on microbenchmarks, showing that rotations can be performed in less than 10 μ s per byte.

2 Updatable AE

We turn to formalizing the syntax and semantics of AE schemes supporting key rotation. Our approach extends that of Boneh et al. [BLMR15] (BLMR), the main syntactical difference being that we allow rekey token generation, re-encryption, and decryption to all return a distinguished error symbol \perp . This is required to enable us to later cater for integrity notions. We also modify the syntax so that ciphertexts include two portions, a header and a body. In our formulation, only the former is used during generation of rekey tokens (while in BLMR the full ciphertext is formally required).




Fig. 1. Interaction between client and cloud during a ciphertext-dependent update. Client retrieves a small ciphertext header, and runs ReKeyGen to produce a compact rekey token Δ . The cloud uses this token to re-encrypt the data. At the end of the update, the data is encrypted using k_2 , and cannot be recovered using only k_1 .

Definition 1 (Updatable AE). An updatable AE scheme is a tuple of algorithms $\Pi = (\text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec})$ with the following properties:

- $\text{KeyGen}() \rightarrow k$. Outputs a secret key k .
- $\text{Enc}(k, m) \rightarrow C$. On input a secret key k and message m , outputs a ciphertext $C = (\tilde{C}, \bar{C})$ consisting of a ciphertext header \tilde{C} and ciphertext body \bar{C} .
- $\text{ReKeyGen}(k_1, k_2, \tilde{C}) \rightarrow \Delta_{1,2,\tilde{C}}$. On input two secret keys, k_1 and k_2 , and a ciphertext header \tilde{C} , outputs a rekey token or \perp .
- $\text{ReEnc}(\Delta_{1,2,\tilde{C}}, (\tilde{C}, \bar{C})) \rightarrow C_2$. On input a rekey token and ciphertext, outputs a new ciphertext or \perp . We require that ReEnc is deterministic.
- $\text{Dec}(k, C) \rightarrow m$. On input a secret key k and ciphertext C outputs either a message or \perp .

Of course we require that all algorithms are efficiently computable. Note that, in common with [BLMR15], our definition is *not* in the nonce-based setting that is widely used for AE. Rather, we will assume that Enc is randomised. We consider this sufficient for a first treatment of updatable AE; it also reflects common industry practice as per the schemes currently used by Amazon [AWS] and Google [Goo]. We relegate the important problem of developing a parallel formulation in the nonce-based setting to future work. Similarly, we assume that all our AE schemes have single decryption errors, cf. [BDPS14], and we do not consider issues such as release of unverified plaintext, cf. [ABL⁺14], tidiness, cf. [NRS14] and length-hiding, cf. [PRS11].

Correctness. An updatable AE scheme is *correct* if decrypting a legitimately generated ciphertext reproduces the original message. Of course, legitimate ciphertexts may be rotated through many keys, complicating the formalization of this notion.

Definition 2 (Correctness). Fix an updatable AE scheme Π . For any message m and any sequence of secret keys k_1, \dots, k_T output by running KeyGen T times, let $C_1 = (\tilde{C}_1, \bar{C}_1) = \text{Enc}(k_1, m)$ and recursively define for $1 \leq t < T$

$$C_{t+1} = \text{ReEnc}(\text{ReKeyGen}(k_t, k_{t+1}, \tilde{C}_t), C_t).$$

Then Π is correct if $\text{Dec}(k_T, C_T) = m$ with probability 1.

Compactness. We say that an updatable AE scheme is compact if the size of both ciphertext headers and rekeying tokens are independent of the length of the plaintext. In practice the sizes should be as small as possible, and for the constructions we consider these are typically a small constant multiple of the key length.

Compactness is important for efficiency of key rotation. Considering the abstract architecture in Figure 1, header values must be available to the key server when rekey tokens are generated. Typically this will mean having to fetch them from storage. Likewise, the rekey token must be sent back to the storage system. Note that there are simple constructions that are not compact, such as the one that sets \tilde{C} to be a standard authenticated encryption of the message and in which ReKeyGen decrypts \tilde{C} , re-encrypts it, and outputs a “rekeying token” as the new ciphertext.

Ciphertext-dependence. As formulated above, updatable AE schemes require part of the ciphertext, the ciphertext header \tilde{C} , in order to generate a rekey token. We will also consider schemes for which \tilde{C} is the empty string, denoted ε . We will restrict attention to schemes for which encryption either always outputs $\tilde{C} = \varepsilon$ or never does. In the former case we call the scheme ciphertext-independent and, in the latter case, ciphertext-dependent. When discussing ciphertext-independent schemes, we will drop \tilde{C} from notation, e.g., writing $\Delta_{i,j}$ instead of $\Delta_{i,j,\tilde{C}}$.

However, we primarily focus on ciphertext-dependent schemes which appear to offer more flexibility and achieve stronger security guarantees (though it is an open question whether a ciphertext-independent scheme can achieve our strongest security notion). We do propose a very lightweight ciphertext-independent scheme included in Appendix A.1, but we show it achieves strictly weaker confidentiality and integrity notions. One can generically convert a ciphertext-independent scheme into a ciphertext-dependent one, simply by deriving a ciphertext-specific key using some unique identifier for the ciphertext. We omit the formal treatment of this trivial approach.

Directionality of rotations. Some updatable AE schemes are bidirectional, meaning rekey tokens can be used to go forwards or backwards.

We only consider bi-directionality to be a feature of ciphertext-independent schemes. Formally, we say that a scheme is *bidirectional* if there exists an efficient algorithm $\text{Invert}(\cdot)$ that produces a valid rekey token $\Delta_{j,i}$ when given $\Delta_{i,j}$ as input.

Schemes that are not bidirectional might be able to ensure that an adversary cannot use rekey tokens to “undo” a rotation of a ciphertext. We will see that ciphertext-dependence can help in building such unidirectional schemes, whereas ciphertext-independent schemes seem harder to make unidirectional. This latter difficulty is related to the long-standing problem of constructing unidirectional proxy re-encryption schemes in the public key setting.

Relationship to proxy re-encryption. Proxy re-encryption targets a different setting than updatable encryption (or AE): the functional ability to allow a ciphertext encrypted under one key to be converted to a ciphertext decryptable by another key. The conversion should not leak plaintext data, but, unlike key rotation, it is not necessarily a goal of proxy re-encryption to remove all dependency on the original key, formalised as indistinguishability of re-encryptions (UP-REENC security) in our work. For example, previous work [CK05, ID03] suggests twice encrypting plaintexts under different keys. To rotate, the previous outer key and a freshly generated outer key is sent to the proxy to perform conversion, but the inner key is never modified. Such an approach does not satisfy the goals of key rotation.

That said, any bidirectional, ciphertext-independent updatable AE ends up also being usable as a symmetric proxy re-encryption scheme (at least as formalized by [BLMR15]).

3 Confidentiality and Integrity for Updatable Encryption

Updatable AE should provide confidentiality for messages as well as integrity of ciphertexts, even in the face of adversaries that obtain rekey tokens and re-encryptions, and that can corrupt some number of secret keys. Finding definitions that rule out trivial wins — e.g., rotating a challenge ciphertext to a compromised key, or obtaining sequences of rekey tokens that allow such rotations — is delicate. We provide a framework for doing so.

Our starting point will be a confidentiality notion which improves significantly upon the previous notion of BLMR by including additional attack vectors, and strengthening existing ones.

For ciphertext integrity, we develop a new definition, building on the usual INT-CTXT notion for standard AE [BN00]. Looking ahead, we will target unidirectional schemes that simultaneously achieve both UP-IND and UP-INT security.

We will follow a concrete security approach in which we do not strictly define security, but rather measure advantage as a function of the resources (running time and number of queries) made by an adversary. Informally, schemes are secure if no adversary with reasonable resources can achieve advantage far from zero.

3.1 Message Confidentiality

The confidentiality game UP-IND is shown in the leftmost column of Figure 2. The adversary's goal is to guess the bit b . Success implies that a scheme leaks partial information about plaintexts. We parameterise the game by two values t and κ . The game initialises $t + \kappa$ secret keys, κ of which are given to the adversary, and t are kept secret for use in the oracles. We label the keys by k_1, \dots, k_t for the uncompromised keys, and by $k_{t+1}, \dots, k_{t+\kappa}$ for the compromised keys. We require at least one uncompromised key, but do not necessarily require any compromised keys, i.e. $t \geq 1$ and $\kappa \geq 0$. We leave consideration of equivalences between models with many keys and few keys and between models with active and static key compromises as interesting problems for future work.

<p>UP-IND</p> <hr/> $b \leftarrow_{\$} \{0, 1\}$ $k_1, \dots, k_{t+\kappa} \leftarrow_{\$} \text{KeyGen}()$ $b' \leftarrow_{\$} \mathcal{A}^O(k_{t+1}, \dots, k_{t+\kappa})$ return $(b' = b)$ <hr/> $\text{Enc}(i, m)$ return $\text{Enc}(k_i, m)$ <hr/> $\text{ReKeyGen}(i, j, \tilde{C})$ if $\text{Invalid}_{\text{RK}}(i, j, \tilde{C})$ then return \perp $\Delta_{i,j,\tilde{C}} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C})$ return $\Delta_{i,j,\tilde{C}}$ <hr/> $\text{ReEnc}(i, j, (\tilde{C}, \bar{C}))$ $\Delta_{i,j,\tilde{C}} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C})$ $C' = (\tilde{C}', \bar{C}') \leftarrow \text{ReEnc}(\Delta_{i,j,\tilde{C}}, (\tilde{C}, \bar{C}))$ if $\text{Invalid}_{\text{RE}}(i, j, \tilde{C})$ then return \tilde{C}' else return C' <hr/> $\text{LR}(i, m_0, m_1)$ if $i > t$ then return \perp $C \leftarrow_{\$} \text{Enc}(k_i, m_b)$ return C	<p>UP-INT</p> <hr/> $\text{win} \leftarrow \text{false}$ $k_1, \dots, k_{t+\kappa} \leftarrow_{\$} \text{KeyGen}()$ $\mathcal{A}^O(k_{t+1}, \dots, k_{t+\kappa})$ return win <hr/> $\text{Enc}(i, m)$ return $\text{Enc}(k_i, m)$ <hr/> $\text{ReKeyGen}(i, j, \tilde{C})$ return $\text{ReKeyGen}(k_i, k_j, \tilde{C})$ <hr/> $\text{ReEnc}(i, j, (\tilde{C}, \bar{C}))$ $\Delta_{i,j,\tilde{C}} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C})$ $C' \leftarrow \text{ReEnc}(\Delta_{i,j,\tilde{C}}, (\tilde{C}, \bar{C}))$ return C' <hr/> $\text{Try}(i, C)$ if $\text{Invalid}_{\text{CTXT}}(i, C)$ then return \perp $M \leftarrow \text{Dec}(k_i, C)$ if $M = \perp$ then return \perp $\text{win} \leftarrow \text{true}$ return M
---	---

Fig. 2. Confidentiality and integrity games for updatable encryption security.

The game relies on two subroutines $\text{Invalid}_{\text{RK}}$ and $\text{Invalid}_{\text{RE}}$ to determine if a re-keygen and re-encryption query, respectively, should be allowed. These procedures are efficiently computed by the game as a function of the adversarial queries and responses. This reliance on the transcript we leave implicit in the notation to avoid clutter. Different choices of invalidity procedures gives rise to distinct definitions of security, and we explain two interesting ones in turn. Note that an invalid query (as determined by $\text{Invalid}_{\text{RE}}$) still results in

the adversary learning the ciphertext header, giving greater power to the adversary. We believe this to be an important improvement both in practice and theoretically over previous models, which consider only a partial compromise. The full compromise of a client results in the adversary playing the role of the client in the key update procedure, during which the server will return the ciphertext header. In practice, it is likely that an adversary who has initially breached the client would use this access to query related services.

Invalidity procedures. For the invalidity constraints used in UP-IND, we target a strong definition, while preventing the adversary from trivially receiving a challenge ciphertext re-encrypted to a compromised key.

We use the ciphertext headers to determine whether a ciphertext has been derived from a challenge ciphertext. It is natural to use only the headers since these will be processed by the client when performing an update. We define a procedure $\text{Derived}_{\text{LR}}(i, \tilde{C})$ that outputs `true` should \tilde{C} have been derived from the ciphertext header returned by an LR query.

Definition 3 (LR-derived headers). *We recursively define function $\text{Derived}_{\text{LR}}(i, \tilde{C})$ to output `true` iff any of the following conditions hold:*

- \tilde{C} was the ciphertext header output in response to a query $\text{LR}(i, m_0, m_1)$
- \tilde{C} was the ciphertext header output in response to a query $\text{ReEnc}(j, i, C')$ and $\text{Derived}_{\text{LR}}(j, \tilde{C}') = \text{true}$
- \tilde{C} is the ciphertext header output by running $\text{ReEnc}(\Delta_{j,i,\tilde{C}'}, C')$ where $\Delta_{j,i,\tilde{C}'}$ is the result of a query $\text{ReKeyGen}(j, i, \tilde{C}')$ for which $\text{Derived}_{\text{LR}}(j, \tilde{C}') = \text{true}$.

The predicate $\text{Derived}_{\text{LR}}(i, \tilde{C})$ is efficient to compute and can be computed locally by the adversary. The most efficient way to implement it is to grow a look-up table T indexed by a key identifier and a ciphertext header and whose entries are sets of ciphertexts. Any query to $\text{LR}(i, m_0, m_1)$ updates the table by adding the returned ciphertext to the set $T[i, \tilde{C}]$ where \tilde{C} is the oracle's returned ciphertext header value. For a query $\text{ReEnc}(j, i, C')$, if $T[j, \tilde{C}']$ is not empty, then it adds the returned ciphertext to the set $T[i, \tilde{C}^*]$ for \tilde{C}^* the returned ciphertext header. For a query $\text{ReKeyGen}(j, i, \tilde{C}')$ with return value $\Delta_{j,i,\tilde{C}'}$, apply $\text{ReEnc}(\Delta_{j,i,\tilde{C}'}, C)$ for all ciphertexts C found in entry $T[j, \tilde{C}']$ and add appropriate new entries to the table. In this way, one can maintain the table in worst-case time that is quadratic in the number of queries and compute in constant time $\text{Derived}_{\text{LR}}(i, \tilde{C})$ by simply checking if $T[i, \tilde{C}]$ is non-empty. If any call to ReKeyGen or ReEnc in $\text{Derived}_{\text{LR}}$ or the main oracle procedure returns \perp , then the entire procedure returns \perp .

Note that $\text{Derived}_{\text{LR}}$ relies on ReEnc being deterministic, a restriction we made in Section 2. To complete the definition, we specify the invalidity procedures that use $\text{Derived}_{\text{LR}}$ as a subroutine:

- $\text{Invalid}_{\text{RK}}(i, j, \tilde{C})$ outputs `true` if $j > t$ and $\text{Derived}_{\text{LR}}(i, \tilde{C}) = \text{true}$. In words, the target key is compromised and i, \tilde{C} derives from an LR query.
- $\text{Invalid}_{\text{RE}}(i, j, \tilde{C})$ outputs `true` if $j > t$ and $\text{Derived}_{\text{LR}}(i, \tilde{C}) = \text{true}$. In words, the target key is compromised and i, \tilde{C} derives from an LR query.

We denote the game defined by using these invalidity procedures by UP-IND. We associate to an UP-IND adversary \mathcal{A} and scheme Π the advantage measure:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) = 2 \cdot \Pr[\text{UP-IND}_{\Pi, \kappa, t}^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

This notion is very strong and bidirectional schemes cannot meet it.

Theorem 1. *Let Π be a bidirectional updatable encryption scheme. Then there exists an UP-IND adversary \mathcal{A} that makes 2 queries and for which*

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) = 1$$

for any $\kappa \geq 1$ and $t \geq 1$.

Proof. We explicitly define the adversary \mathcal{A} . It makes a query to $C_1 = \text{LR}(1, m_0, m_1)$ for arbitrary messages $m_0 \neq m_1$ and computes locally $C_{t+1} = \text{Enc}(k_{t+1}, m_1)$. It then makes a query $\Delta_{t+1,1,\tilde{C}_{t+1}} = \text{ReKeyGen}(t+1, 1, \tilde{C}_{t+1})$. It runs $C' = \text{ReEnc}(\text{Invert}(\Delta_{t+1,1,\tilde{C}_{t+1}}, C_{t+1}, C_1), C_1)$ locally and then decrypts C' using k_{t+1} . It checks whether the result is m_0 or m_1 and returns the appropriate bit. \square

BLMR confidentiality. In comparison, we define invalidity procedures corresponding to those in BLMR's security notion.

- $\text{InvalidBLMR}_{\text{RK}}(i, j, \tilde{C})$ outputs **true** if $i \leq t < j$ or $j \leq t < i$ and outputs **false** otherwise. In words, the query is not allowed if exactly one of the two keys is compromised.
- $\text{InvalidBLMR}_{\text{RE}}(i, j, \tilde{C})$ outputs **true** if $j > t$ and **false** otherwise. In words, the query is not allowed if the target key k_j is compromised.

We denote the game defined by using these invalidity procedures by UP-IND-BI (the naming will become clear presently). We associate to an UP-IND-BI adversary \mathcal{A} , scheme Π , and parameters κ, t the advantage measure:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind-bi}}(\mathcal{A}) = 2 \cdot \Pr[\text{UP-IND-BI}_{\Pi, \kappa, t}^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

A few observations are in order. First, it is apparent that the invalidity procedures for the BLMR notion are significantly stronger than ours, leading to a weaker security notion: the BLMR procedures are not ciphertext-specific but instead depend only on the compromise status of keys. We will show that this difference is significant. In addition, the corresponding BLMR definition did not consider leakage of the ciphertext header when $\text{InvalidBLMR}_{\text{RE}}$ returns **true**. Second, for ciphertext-independent schemes in which $\tilde{C} = \varepsilon$ always, the BLMR definition coincides with symmetric proxy re-encryption security (as also introduced in their paper [BLMR15]). Third, the BLMR confidentiality notion does not require unidirectional security of rekey tokens because it has the strong restriction of disallowing attackers from obtaining rekey tokens $\Delta_{i, j, \tilde{C}}$ with $i > t$ (so the corresponding key is compromised), but with $j < t$ (for an uncompromised key). Thus, in principle, bidirectional schemes could meet this notion, explaining our naming convention for the notion. Finally, the BLMR notion does not require ciphertext-specific rekey tokens because the invalidity conditions are based only on keys and not on the target ciphertext.

Detailed in Appendix A.1 is a bidirectional scheme that is secure in the sense of UP-IND-BI. This result and the negative result that no bidirectional scheme can achieve UP-IND given above (Theorem 1) yields as a corollary that UP-IND-BI security is strictly weaker than UP-IND security. This illustrates the enhanced strength of our UP-IND security notion compared to the corresponding BLMR notion, UP-IND-BI.

Given that bidirectional, ciphertext-independent schemes have certain advantages in terms of performance and deployment simplicity, practitioners may prefer them in some cases. For that flexibility, one trades off control over the specificity of rekey tokens, which could be dangerous to confidentiality in some compromise scenarios.

3.2 Ciphertext Integrity

We now turn to a notion of integrity captured by the game UP-INT shown in Figure 2. The adversary’s goal is to submit a ciphertext to the Try oracle that decrypts properly. Of course, we must exclude the adversary from simply resubmitting valid ciphertexts produced by the encryption oracle, or derived from such an encryption by way of re-encryption queries or rekey tokens.

In a bit more detail, in the Try oracle, we define a predicate InvalidCTXT which captures whether the adversary has produced a trivial derivation of a ciphertext obtained from the encryption oracle. This fulfills a similar role to that of the $\text{Invalid}_{\text{RE}}$ and $\text{Invalid}_{\text{RK}}$ subroutines in the UP-IND game.

For the unidirectional security game UP-INT, we define $\text{InvalidCTXT}(i, C = (\tilde{C}, \bar{C}))$ inductively, outputting **true** if any of the following conditions hold:

- $i > t$, i.e. k_i is known to the adversary
- (\tilde{C}, \bar{C}) was output in response to a query $\text{Enc}(i, m)$
- (\tilde{C}, \bar{C}) was output in response to a query $\text{ReEnc}(j, i, C')$ and $\text{InvalidCTXT}(j, C') = \text{true}$
- (\tilde{C}, \bar{C}) is the ciphertext output by running $\text{ReEnc}(\Delta_{j, i, \tilde{C}'}, C')$ for $C' = (\tilde{C}', \bar{C}')$ where $\Delta_{j, i, \tilde{C}'}$ was the result of a query $\text{ReKeyGen}(j, i, \tilde{C}')$ and $\text{InvalidCTXT}(j, C') = \text{true}$.

This predicate requires the transcript of queries thus far; to avoid clutter we leave the required transcript implicit in our notation. The definition of InvalidCTXT is quite permissive: it defines invalid ciphertexts as narrowly as possible, making our security notion stronger. Notably, the adversary can produce any ciphertext (valid or otherwise) using a corrupted key k_i , and use the ReKeyGen oracle to learn a token to update this ciphertext to a non-compromised key. Only the direct re-encryption of the submitted ciphertext is forbidden.

We associate to an updatable encryption scheme Π , an UP-INT adversary \mathcal{A} , and parameters κ, t the advantage measure:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-int}}(\mathcal{A}) = \Pr[\text{UP-INT}_{\Pi, \kappa, t}^{\mathcal{A}} \Rightarrow \text{true}] .$$

4 Practical Updatable AE Schemes

We first investigate the security of updatable AE schemes built using the KEM/DEM approach sketched in the introduction. Such schemes are in widespread use at present, for example in AWS’s and Google’s cloud storage systems [AWS,Goo], yet have received no formal analysis to date. We produce the AE-hybrid construction as a formalism of this common practice.

Using the confidentiality and integrity definitions from the previous section, we discover that this construction offers very weak security against an adversary capable of compromising keys. Indeed, we are only able to prove security when the number of compromised keys κ is equal to 0. Given the intention of key rotation this is a somewhat troubling result.

On a positive note, we show a couple of simple tweaks to the AE-hybrid which fix these issues. The resultant scheme, named KSS, offers improved security at little additional cost.

We leave to the appendix our bidirectional, ciphertext-independent scheme XOR-KEM which does not offer strong integrity guarantees but may be of interest for other applications.

4.1 Authenticated Encryption

In the following constructions we make use of authenticated encryption (AE) schemes which we define here.

Definition 4 (Authenticated encryption). *An authenticated encryption scheme π is a tuple of algorithms $(\mathcal{K}, \mathcal{E}, \mathcal{D})$. \mathcal{K} is a randomised algorithm outputting keys. We denote by $\mathcal{E}_k(\cdot)$ the randomised algorithm for encryption by key k and by $\mathcal{D}_k(\cdot)$ decryption. Decryption is a deterministic algorithm and outputs the distinguished symbol \perp to denote a failed decryption.*

In keeping with our definitional choices for updatable AE, we consider randomised AE schemes rather than nonce-based ones.

We use the all-in-one authenticated encryption security definition from [RS06].

Definition 5 (Authenticated Encryption Security). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an authenticated encryption scheme. Let Enc, Dec be oracles whose behaviors depends on hidden values $b \in \{0, 1\}$ and key $k \leftarrow_s \mathcal{K}$. Enc takes as input a bit string m and produces $\mathcal{E}_k(m)$ when $b = 0$, and produces a random string of the same length otherwise. Dec takes as input a bit string C and produces $\mathcal{D}_k(C)$ when $b = 0$, and produces \perp otherwise.*

Let $\text{AE-ROR}_{\pi}^{\mathcal{A}}$ be the game in which an adversary \mathcal{A} has access to the Enc and Dec oracles and must output a bit b' . The game outputs true when $b = b'$. We require that the adversary not submit outputs from the Enc oracle to the Dec oracle.

We define the advantage of \mathcal{A} in the AE-ROR security game for π as:

$$\text{Adv}_{\pi}^{\text{ae}}(\mathcal{A}) = 2 \cdot \Pr[\text{AE-ROR}_{\pi}^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

Unless otherwise stated, our AE schemes will be length-regular, so that the lengths of ciphertexts depend only on the lengths of plaintexts. This ensures that the above definition also implies a standard “left-or-right” security definition.

4.2 (In-)Security of AE-hybrid Construction

Figure 3 defines an updatable AE scheme, AE-hybrid, for any AE scheme $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. This is a natural key-wrapping scheme that one might create in the absence of security definitions. It is preferred by practitioners because key rotation is straightforward and performant. Using this scheme means re-keying requires constant time and communication, independent of the length of the plaintext. In fact, we note that this scheme sees

<u>Enc(k, m)</u>	<u>ReKeyGen(k_1, k_2, \tilde{C})</u>	<u>Dec($k, (\tilde{C}, \bar{C})$)</u>
$x \leftarrow_{\mathcal{S}} \mathcal{K}$	$x = \mathcal{D}(k_1, \tilde{C})$	$x = \mathcal{D}(k, \tilde{C})$
$\tilde{C} \leftarrow_{\mathcal{S}} \mathcal{E}(k, x)$	if $x = \perp$ return \perp	if $x = \perp$ return \perp
$\bar{C} \leftarrow_{\mathcal{S}} \mathcal{E}(x, m)$	$\Delta_{1,2,\tilde{C}} \leftarrow_{\mathcal{S}} \mathcal{E}(k_2, x)$	$m = \mathcal{D}(x, \bar{C})$
return (\tilde{C}, \bar{C})	return $\Delta_{1,2,\tilde{C}}$	return m
KeyGen : return $k \leftarrow_{\mathcal{S}} \mathcal{K}$		
ReEnc($\Delta_{1,2,\tilde{C}}, (\tilde{C}, \bar{C})$) : return $C' = (\Delta_{1,2,\tilde{C}}, \bar{C})$		

Fig. 3. Algorithms for the AE-hybrid updatable AE scheme.

widening deployment for encrypted cloud storage services. Both Amazon Web Services [AWS] and Google Cloud Platform [Goo] use AE-hybrid to perform key rotations over encrypted customer data.

We demonstrate severe limits of AE-hybrid: when keys are compromised confidentiality and integrity cannot be recovered through re-encryption. Later we will demonstrate straightforward modifications to AE-hybrid that allow it to recover both confidentiality and integrity without impacting performance.

Theorem 2 (UP-IND Insecurity of AE-hybrid). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and Π be the updatable AE scheme AE-hybrid using π as defined in Figure 3.*

Then there exists an adversary \mathcal{A} making 2 queries such that $\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) = 1$ for all $\kappa \geq 1$ and $t \geq 1$.

Proof. We construct a concrete adversary \mathcal{A} satisfying the theorem statement.

\mathcal{A} makes an initial query to LR($1, m_0, m_1$) for distinct messages $m_0 \neq m_1$ and receives challenge ciphertext $C^* = (\mathcal{E}(k_1, x), \mathcal{E}(x, m_b))$. \mathcal{A} subsequently calls ReKeyGen($1, t+1, C^*$). k_{t+1} is corrupted and thus Invalid_{RK} returns true, so the adversary receives the re-encrypted ciphertext header $\tilde{C}' = \mathcal{E}(k_{t+1}, x)$.

The adversary decrypts $x = \mathcal{D}(k_{t+1}, \tilde{C}')$, computes $m_b = \mathcal{D}(x, \bar{C}^*)$ and checks whether $m_b = m_0$ or m_1 . \square

The best one can achieve with this scheme is to prove security when $\kappa = 0$, that is, security is not degraded beyond the underlying AE scheme when the adversary does not obtain any compromised keys. However, such a weak security notion is not particularly interesting, since the intention of key rotation is to recover security in the face of key compromises.

Theorem 3 (UP-IND security of AE-hybrid with no key compromise). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and Π be the updatable AE scheme AE-hybrid using π as defined above. Then for any adversary \mathcal{A} for the game UP-IND, there exists an adversary \mathcal{B} for the AE security game where:*

$$\text{Adv}_{\Pi, 0, t}^{\text{up-ind}}(\mathcal{A}) \leq 2(t+1) \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B})$$

for all $t \geq 1$.

Proof. We argue using a series of games that the advantage of any adversary in the UP-IND game for AE-hybrid is bounded by the advantage in the AE security game for the underlying AE scheme π . The first game, G_0 , is the UP-IND game for AE-hybrid using the underlying scheme π . For $1 \leq i \leq t$, in game G_i we modify the LR oracle by dropping one key k_i , and instead replacing \tilde{C} with a random string of the same length (and store the \tilde{C} and x value) used to answer the query. The ReKeyGen and ReEnc oracles use these stored values to respond.

Let S_i be the event that G_i outputs true. We claim that for $0 < i \leq t$, there exists \mathcal{B} such that:

$$|\Pr[S_{i-1}] - \Pr[S_i]| \leq \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}).$$

We construct this \mathcal{B} by replacing the random string selection for oracle queries with key index i with calls to the Enc oracle in the AE security game for π . When the hidden bit \hat{b} in the AE game is 0, then \mathcal{B} perfectly simulates game G_{i-1} for key k_i and when $\hat{b} = 1$, then \mathcal{B} perfectly simulates game G_i .

Similarly, in game G_{t+1} the ciphertext body \bar{C} is replaced by a random string of the same length and again we have:

$$|\Pr[S_t] - \Pr[S_{t+1}]| \leq \text{Adv}_\pi^{\text{ae}}(\mathcal{B}).$$

And we construct \mathcal{B} in the same way as before, replacing \bar{C} in the LR oracle with a call to the AE encryption oracle. Depending on the value of the hidden bit in the AE security game, either G_t or G_{t+1} is perfectly simulated.

Combining our two claims above we see that:

$$|\Pr[S_0] - \Pr[S_{t+1}]| \leq (t+1) \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B}).$$

Finally, observe that the adversary in game G_{t+1} learns nothing about the encrypted message m_b , since it is given only random strings. Hence $\Pr[S_{t+1}] = \frac{1}{2}$ and we conclude:

$$\begin{aligned} \text{Adv}_{\Pi,0,t}^{\text{up-ind}}(\mathcal{A}) &= 2 \cdot \Pr[S_0] - 1 \\ &= 2 \cdot (\Pr[S_0] - \Pr[S_{t+1}] + \Pr[S_{t+1}]) - 1 \\ &\leq 2(t+1) \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B}). \end{aligned}$$

□

Similarly, AE-hybrid is trivially insecure in the UP-INT sense when $\kappa \geq 1$.

Theorem 4 (UP-INT insecurity of AE-hybrid). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and Π be the updatable AE scheme AE-hybrid using π as defined in Figure 3.*

Then there exists an adversary \mathcal{A} making 2 queries and one Try query such that $\text{Adv}_{\Pi,\kappa,t}^{\text{up-int}}(\mathcal{A}) = 1$ for all $\kappa \geq 1$ and $t \geq 1$.

Proof. We construct a concrete adversary \mathcal{A} satisfying the theorem statement.

\mathcal{A} first queries $\text{Enc}(1, m)$ to obtain an encryption $C = (\mathcal{E}(k_1, x), \mathcal{E}(x, m))$, and subsequently queries $\text{ReEnc}(1, t+1, C)$, receiving the re-encryption $C' = (\mathcal{E}(k_{t+1}, x), \mathcal{E}(x, m))$. Since \mathcal{A} has key k_{t+1} , \mathcal{A} recovers $x = \mathcal{D}(k_{t+1}, \tilde{C}')$ by performing the decryption locally.

Finally, \mathcal{A} constructs the ciphertext $C^* = (\tilde{C}, \mathcal{E}(x, m'))$ for some $m' \neq m$ and queries $\text{Try}(1, C^*)$. Since C^* is not derived from C and k_1 is not compromised, UP-INT outputs true. □

As before, we can only prove UP-INT security of AE-hybrid when $\kappa = 0$.

Theorem 5 (UP-INT Security of AE-hybrid with no key compromise). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and Π be the AE-hybrid scheme using π as defined above. Then for any adversary \mathcal{A} for the game UP-INT there exists an adversary \mathcal{B} for the AE security game where:*

$$\text{Adv}_{\Pi,0,t}^{\text{up-int}}(\mathcal{A}) \leq (t+1) \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B})$$

for all $t \geq 1$.

Proof. We use the same sequence of games as in the previous proof, leveraging the integrity properties of authenticated encryption.

In game G_i for $1 \leq i \leq t$, for any query (j, C) to the Try oracle where $j \leq i$, the decryption of the ciphertext header \tilde{C} will always decrypt to \perp , unless \tilde{C} was a previously generated random string output by the encryption oracle.

Hence in game G_t , the adversary can only win by re-using a previously seen header \tilde{C} which is associated to a DEK x , and constructing a distinct message body \tilde{C}' .

Between games G_t and G_{t+1} , we again replace the ciphertext body by a call to the AE encryption oracle. Game G_t corresponding to the real world, and G_{t+1} to the random world.

As stated previously, the adversary can only win (i.e. produce a distinct, unique ciphertext) by submitting a distinct, valid ciphertext body. However, in G_{t+1} , the decryption oracle always returns \perp and hence the adversary wins with probability 0. Finally, we conclude:

$$\text{Adv}_{\Pi,0,t}^{\text{up-int}}(\mathcal{A}) \leq (t+1) \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B}).$$

□

<u>Enc(k, m)</u>	<u>ReKeyGen(k_1, k_2, \tilde{C})</u>	<u>Dec($k, (\tilde{C}, \bar{C})$)</u>
$x, y \leftarrow_{\mathcal{S}} \mathcal{K}$	$(\chi \parallel \tau) = \mathcal{D}(k_1, \tilde{C})$	$(\chi \parallel \tau) = \mathcal{D}(k, \tilde{C})$
$\chi = x \oplus y$	if $(\chi \parallel \tau) = \perp$ return \perp	if $(\chi \parallel \tau) = \perp$
$\bar{C}^1 \leftarrow_{\mathcal{S}} \mathcal{E}(x, m)$	$y' \leftarrow_{\mathcal{S}} \mathcal{K}$	return \perp
$\tau = \mathcal{E}(x, h(m))$	return $(y', \mathcal{E}(k_2, (\chi \oplus y') \parallel \tau))$	$x = \chi \oplus \bar{C}^0$
$\tilde{C} \leftarrow_{\mathcal{S}} \mathcal{E}(k, \chi \parallel \tau)$		$m = \mathcal{D}(x, \bar{C}^1)$
return $(\tilde{C}, (y, \bar{C}^1))$		if $\mathcal{D}(x, \tau) \neq h(m)$
		return \perp
		return m
KeyGen(): return $k \leftarrow \mathcal{K}$		
ReEnc($\Delta_{1,2,\tilde{c}}, (\tilde{C}, \bar{C})$): return $(\Delta_{1,2,\tilde{c}}^1, (\bar{C}^0 \oplus \Delta_{1,2,\tilde{c}}^0, \bar{C}^1))$		

Fig. 4. Algorithms for the KSS updatable AE scheme.

4.3 Improving AE-hybrid

We make small modifications to AE-hybrid and show that the resulting construction has both UP-IND and UP-INT security. These modifications include masking the DEM key stored inside the ciphertext header (to gain UP-IND security), and including an encrypted hash of the message (for UP-INT). We note that these modifications are straightforward to implement on top of the AE-hybrid scheme and have only minimal impact on the scheme's performance.

Let $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme and h a hash function with ℓ_h output bits. Then we define KSS (KEM/DEM with Secret Sharing) in Figure 4.

Theorem 6 (UP-IND Security of KSS). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and Π be the updatable AE scheme KSS using π as defined in Figure 4. Then for any adversary \mathcal{A} for the game UP-IND, making at most q queries to the LR oracle, there exists an adversary \mathcal{B} for the AE security game where:*

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) \leq 2(t + q) \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B})$$

for all $\kappa \geq 0, t \geq 1$.

Proof. We argue using a series of games that the advantage of any adversary in the UP-IND game for KSS is bounded by the advantage in the AE security game for the underlying AE scheme π . The first game, G_0 , is the UP-IND game for KSS using the underlying scheme π . For $1 \leq i \leq t$, in game G_i we replace all ciphertext headers encrypted under key k_i and instead return a random string of the same length used to answer the query. The value of the returned ciphertext header \tilde{C} is stored along with the encrypted value $(\chi \parallel \tau)$ in order to simulate later calls to ReEnc and ReKeyGen.

Let S_i be the event that G_i outputs true. We claim that for $0 < i \leq t$, there exists \mathcal{B} such that:

$$|\Pr[S_{i-1}] - \Pr[S_i]| \leq \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}).$$

We construct \mathcal{B} by using the Enc oracle in the AE security game for π to encrypt ciphertext headers. When the hidden bit \hat{b} in the AE game is 0, then \mathcal{B} perfectly simulates game G_{i-1} for key k_i and when $\hat{b} = 1$, then \mathcal{B} perfectly simulates game G_i . \mathcal{B} simply returns as a guess $\hat{b} = 1$ if the adversary is correct. Any difference between the success probabilities in G_{i-1} and G_i results in an advantage for the adversary.

Now, we observe that the adversary in game G_t cannot learn anything about the DEM key x used to encrypt a challenge. Even through a re-encryption to a corrupted key, the most the adversary can learn is the value $\chi' = x \oplus y'$, where y' is in the ciphertext body and unobtainable by the adversary.

Hence consider another set of hybrids G_j for $t < j \leq t + q$, where the adversary makes at most q queries to the left-or-right oracle LR. In the same spirit as the previous hybrids, we construct an AE adversary such that $|\Pr[S_{j-1}] - \Pr[S_j]| \leq \text{Adv}_\pi^{\text{ae}}(\mathcal{B})$, this time replacing the encryption of the DEM during a challenge query with the output of the AE encryption oracle.

Finally, observe that in game G_{t+q} , all outputs from LR oracle are of the form (r, χ, z) where r, z are random strings. Since the outputs are unrelated to message inputs, and indeed the hidden bit b , the adversary can learn nothing from oracle queries and so $\Pr[S_{t+q}] = \frac{1}{2}$ and we conclude:

$$\begin{aligned} \text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) &= 2 \cdot \Pr[S_0] - 1 \\ &= 2 \cdot (\Pr[S_0] - \Pr[S_1] + \Pr[S_1] + \dots - \Pr[S_{t+q}] + \Pr[S_{t+q}]) - 1 \\ &\leq 2(t + q) \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B}). \end{aligned}$$

□

Our modification to include an encrypted hash of the ciphertext is in order to provide a measure of integrity protection. As we will see in the following theorem, collision resistance of the hash function is sufficient to provide UP-INT security, since the hash is integrity-protected by the AE encryption of the KEM. The hash is encrypted to avoid compromise of the ciphertext header being sufficient to distinguish messages.

We achieve collision resistance by assuming h to be a random oracle. However, this assumption could be avoided by either re-using the DEM key x to additionally key the hash function.

We note that this combination of hash function and AE encryption is used to provide an additional integrity mechanism that works for any AE scheme. However, some schemes may be able to avoid this additional computation by re-using components of the AE encryption. For example, if an encrypt-then-MAC scheme is used such that the encryption and MAC keys are both uniquely derived from the DEM key x , then we conjecture that the MAC itself can be used in place of the encrypted hash.

Theorem 7 (UP-INT Security of KSS). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme, h be a cryptographic hash function modelled as a random oracle with output length ℓ_h , and Π be the updatable AE scheme KSS using π and h as defined in Figure 4. Then for any polynomial-time adversary \mathcal{A} , making at most q_h queries to the random oracle h , there exists an adversary \mathcal{B} for the AE security game where:*

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-int}}(\mathcal{A}) \leq t \cdot \text{Adv}_\pi^{\text{ae}}(\mathcal{B}) + \frac{q_h^2}{2^{\ell_h}}$$

for all $\kappa \geq 0, t \geq 1$.

Proof. We use the same sequence of games as in the previous proof, leveraging the integrity properties of authenticated encryption.

In game G_i for $1 \leq i \leq t$, for any query (j, C) to the Try oracle where $j \leq i$, the decryption of the ciphertext header \tilde{C} will always decrypt to \perp , unless \tilde{C} was a previously generated random string output by the encryption oracle.

Hence in game G_t , the adversary can only win by re-using a previously seen header \tilde{C} .

Therefore, the adversary can only win by submitting a ciphertext of the form: $(\tilde{C}, \overline{C}')$, where $\overline{C}' \neq \overline{C}$, the ciphertext body returned in the same query as \tilde{C} .

Let $\overline{C} = (y, \overline{C}^1)$ and thus we require $\overline{C}' = (y', \mathcal{E}(\chi \oplus y', m'))$ for either $y' \neq y$ or $\mathcal{E}(\chi \oplus y', m') \neq \overline{C}^1$. Note that the adversary can learn the value of x, y used in \overline{C} by querying $\text{ReKeyGen}(i, t + 1, \tilde{C})$ and receiving $\mathcal{E}(k_{t+1}, \chi \parallel h(\overline{C}^1)), y$. From this, they can simply decrypt and learn $x = \chi \oplus y$.

To succeed, the adversary needs to find y', m' such that $h(m') = \mathcal{D}(\chi \oplus y', \mathcal{E}(x, h(m)))$. Modelling h as a random oracle, for any value of $y' \neq y$, the probability that the adversary finds an input z to h , such that $h(z) = \mathcal{D}(\chi \oplus y', \mathcal{E}(x, h(m)))$ is $\frac{1}{2^{\ell_h}}$. If such a value z is found, the adversary can compute $m' = \mathcal{D}(\chi \oplus y, z)$. However, the probability that the adversary, making at most q_h oracle queries to h finds such a value is given by $\frac{q_h}{2^{\ell_h}}$.

Alternatively, for $y = y'$, the adversary needs to find a message $m' \neq m$ such that $h(m') = h(m)$. Again, modelling h as a random oracle, the probability that the adversary, making at most q_h oracle queries to h , finds a collision is given by $\frac{q_h^2}{2^{\ell_h}}$.

Since this probability is greater than the probability of finding a pre-image when $y' \neq y$, the probability that the adversary wins in G_{t+1} is bounded by this probability of finding a collision, and we conclude that:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-int}}(\mathcal{A}) \leq t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + \frac{q_h^2}{2^{\ell_h}}.$$

□

5 Indistinguishability of Re-encryptions

The KSS scheme in the previous section achieves message confidentiality and ciphertext integrity, even though the actual DEM key is not modified in the course of performing a rotation. Modifying the scheme to ensure the DEM key is also rotated is non-trivial, requiring either significant communication complexity (linear in the length of the encrypted message) between the key server and storage, or the introduction of more advanced primitives such as key-homomorphic PRFs. The question that arises is whether or not changing DEM keys leaves KSS vulnerable to attacks not captured by the definitions introduced thus far.

BLMR’s brief treatment of updatable encryption attempts to speak to this issue by requiring that all randomness be refreshed during a rotation. Intuitively this would seem to improve security, but the goal they formalize for this, detailed below, is effectively a correctness condition (i.e., it does not seem to account for adversarial behaviors). It doesn’t help clarify what attacks would be ruled out by changing DEM keys.

Exfiltration attacks. We identify an issue with our KSS scheme (and the other schemes in the preceding section) in the form of an attack that is not captured by the confidentiality definitions introduced so far. Consider our simple KSS scheme in the context of our motivating key server and storage service application (described in Section 2). Suppose an attacker compromises for some limited time both the key server and the storage service. Then for each ciphertext (\tilde{C}, \bar{C}) encrypted under a key k_1 , the attacker can compute the DEM key $y \oplus \chi = x$ and exfiltrate it.

Suppose the compromise is cleaned up, and the service immediately generates new keys and rotates all ciphertexts to new secret keys. For the KSS scheme, the resulting ciphertexts will still be later decryptable using the previously exfiltrated DEM keys.

Although a confidentiality issue — the attacker later obtains access to plaintext data they should not have — our UP-IND security notion (and, by implication, the weaker BLMR confidentiality notion) do not capture these attacks. Technically this is because the security game does not allow a challenge ciphertext to be encrypted to a compromised key (or rotated to one). Intuitively, the UP-IND notion gives up on protecting the plaintexts underlying such ciphertexts, as the attacker in the above scenario already had access to the plaintext in the first phase of the attack.

One might therefore argue that this attack is not very important. All of the plaintext data eventually at risk of later decryption was already exposed to the adversary in the first time period because she had access to both the key and ciphertexts. But quantitatively there is a difference: for a given ciphertext an adversary in the first time period can exfiltrate just $|x|$ bits per ciphertext to later recover as much plaintext as she likes, whereas the trivial attack may require exfiltrating the entire plaintext.

The chosen-message attack game of UP-IND does not capture different time periods in which the adversary knows plaintexts in the first time period but “forgets them” in the next. One could explicitly model this, perhaps via a two-stage game with distinct adversaries in each stage, but such games are complex and often difficult to reason about (cf., [RSS11]). We instead develop what we believe is a more intuitive route that asks that the re-encryption of a ciphertext should leak nothing about the *ciphertext* that was re-encrypted. We use an indistinguishability-style definition to model this. The interpretation of our definition is that any information derivable from a ciphertext (and its secret key) before a re-encryption isn’t helpful in attacking the re-encrypted version.

Re-encryption indistinguishability. We formalize this idea via the game shown in Figure 5. The adversary is provided with a left-or-right *re-encryption* oracle, ReLR, instead of the usual left-or-right encryption

UP-REENC	$\text{Enc}(i, m)$	$\text{ReKeyGen}(i, j, \tilde{C})$
$b \leftarrow_{\$} \{0, 1\}$ $k_1, \dots, k_{t+\kappa} \leftarrow_{\$} \text{KeyGen}()$ $b' \leftarrow_{\$} \mathcal{A}^O(k_{t+1}, \dots, k_{t+\kappa})$ return $(b' = b)$	return $\text{Enc}(k_i, m)$	if $\text{Invalid}_{\text{RK}}(i, j, \tilde{C})$ then return \perp $\Delta_{i,j,\tilde{C}} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C})$ return $\Delta_{i,j,\tilde{C}}$
$\text{ReEnc}(i, j, (\tilde{C}, \bar{C}))$	$\text{ReLR}(i, j, C_0, C_1)$	
$\Delta_{i,j,\tilde{C}} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C})$ $C' = (\tilde{C}', \bar{C}') \leftarrow \text{ReEnc}(\Delta_{i,j,\tilde{C}}, (\tilde{C}, \bar{C}))$ if $\text{Invalid}_{\text{RE}}(i, j, \tilde{C})$ then return \tilde{C}' else return C'	if $j > t$ or $ C_0 \neq C_1 $ then return \perp for $\beta \in \{0, 1\}$ do $\Delta_{i,j,\tilde{C}_\beta} \leftarrow_{\$} \text{ReKeyGen}(k_i, k_j, \tilde{C}_\beta)$ $C'_\beta \leftarrow \text{ReEnc}(\Delta_{i,j,\tilde{C}_\beta}, C_\beta)$ if $C'_\beta = \perp$ then return \perp return C'_b	

Fig. 5. The game used to define re-encryption indistinguishability.

oracle, in addition to the usual collection of compromised keys, a re-encryption oracle, encryption oracle, and rekey token generation oracle. We assume that the adversary always submits ciphertext pairs such that $|C_0| = |C_1|$.

To avoid trivial wins, the game must disallow the adversary from simply re-encrypting the challenge to a corrupted key. Hence we define a $\text{Derived}_{\text{ReLR}}$ predicate, which is identical to the $\text{Derived}_{\text{LR}}$ predicated defined in Section 3 for UP-IND security, except that it uses the ReLR challenge oracle. We give it in full detail in the next definition.

Definition 6 (ReLR-derived headers). We recursively define the function $\text{Derived}_{\text{ReLR}}(i, \tilde{C})$ to output true iff $\tilde{C} \neq \varepsilon$ and any of the following conditions hold:

- \tilde{C} was the ciphertext header output in response to a query $\text{ReLR}(i, C_0, C_1)$.
- \tilde{C} was the ciphertext header output in response to a query $\text{ReEnc}(j, i, C')$ and $\text{Derived}_{\text{ReLR}}(j, \tilde{C}') = \text{true}$.
- \tilde{C} is the ciphertext header output by running $\text{ReEnc}(\Delta_{j,i,\tilde{C}'}, C')$ where $\Delta_{j,i,\tilde{C}'}$ is the result of a query $\text{ReKeyGen}(j, i, C')$ for which $\text{Derived}_{\text{ReLR}}(j, \tilde{C}') = \text{true}$.

Then the subroutines $\text{Invalid}_{\text{RK}}$, $\text{Invalid}_{\text{RE}}$ used in the game output true if $\text{Derived}_{\text{ReLR}}(i, \tilde{C})$ outputs true and $j > t$. We associate to an updatable encryption scheme Π , UP-REENC adversary \mathcal{A} , and parameters κ, t the advantage measure:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-reenc}}(\mathcal{A}) = 2 \cdot \Pr[\text{UP-REENC}_{\Pi, \kappa, t}^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

Informally, an updatable encryption scheme is UP-REENC secure if no adversary can achieve advantage far from zero given reasonable resources (run time, queries, and number of target keys).

Notice that exfiltration attacks as discussed informally above would not apply to a scheme that meets UP-REENC security. Suppose otherwise, that the exfiltration still worked. Then one could build an UP-REENC adversary that worked as follows. It obtains two encryptions of different messages under a compromised key, calculates the DEM key (or whatever other information is useful for later decryption) and then submits the ciphertexts to the ReLR oracle, choosing as target a non-compromised key ($j \leq t$). Upon retrieving the ciphertext, it uses the DEM key to decrypt, and checks which message was encrypted. Of course our notion covers many other kinds of attacks, ruling out even re-encryption that allows a single bit of information to leak.

UP-REENC01 _m	UP-REENC00 _m
$k_1, k_2 \leftarrow \text{KeyGen}()$	$k_1, k_2 \leftarrow \text{KeyGen}()$
$C_0 \leftarrow \text{Enc}(k_1, m), C_1 \leftarrow \text{Enc}(k_1, m)$	$C_0 \leftarrow \text{Enc}(k_1, m), C_1 \leftarrow \text{Enc}(k_1, m)$
$\Delta_{1,2,\tilde{C}_1} \leftarrow \text{ReKeyGen}(k_1, k_2, \tilde{C}_1)$	$\Delta_{1,2,\tilde{C}_0} \leftarrow \text{ReKeyGen}(k_1, k_2, \tilde{C}_0)$
$C'_1 \leftarrow \text{ReEnc}(\Delta_{1,2,\tilde{C}_1}, C_1)$	$C'_0 \leftarrow \text{ReEnc}(\Delta_{1,2,\tilde{C}_0}, C_0)$
return (C_0, C'_1)	return (C_0, C'_0)

Fig. 6. Re-encryption indistinguishability experiments from [BLMR15].

BLMR re-encryption security. BLMR introduced a security goal that we will call basic re-encryption indistinguishability.¹ In words, it asks that the distribution of a ciphertext and its re-encryption should be identical to the distribution of a ciphertext and a re-encryption of a distinct ciphertext of the same message.

More formally we have a pair of experiments shown in Figure 6, each parameterized by a message m . Then BLMR require that for all m and all ciphertext pairs (C, C') :

$$|\Pr[\text{UP-REENC00}_m \Rightarrow (C, C')] - \Pr[\text{UP-REENC01}_m \Rightarrow (C, C')]| = 0$$

where the probabilities are over the coins used in the experiments.

This goal misses a number of subtleties which are captured by our definition. Our definition permits the adversary, for example, to submit *any* pair of ciphertexts to the ReLR oracle. This includes ciphertexts which are encryptions of distinct messages, and even maliciously formed ciphertexts which may not even decrypt correctly. It is simple to exhibit a scheme that meets the BLMR notion but trivially is insecure under ours.²

On the other hand, suppose a distinguisher exists that can with some probability ϵ distinguish between the outputs UP-REENC00_m and UP-REENC01_m for some m . Then there exists an adversary against our UP-REENC notion which achieves advantage ϵ . This can be seen by the following simple argument. The adversary gets $C \leftarrow \text{Enc}(1, m), C' \leftarrow \text{Enc}(1, m)$ and submits the tuple $(1, 2, C, C')$ to its ReLR oracle and receives a re-encryption of one of the ciphertexts, C^* . The adversary then runs the distinguisher on (C, C^*) and outputs whatever the distinguisher guesses. If the distinguisher is computationally efficient, then so too is the UP-REENC adversary. Thus our UP-REENC notion would be stronger than a computational version of the BLMR notion.

6 Revisiting the BLMR Scheme

The fact that the simple KEM/DEM schemes of Section 4 fail to meet re-encryption security begs the question of finding new schemes that achieve it, as well as UP-IND and UP-INT security. Our starting point is the BLMR construction of an updatable encryption from key-homomorphic PRFs. Their scheme does not (nor did it attempt to) provide integrity guarantees, and so trivially does not meet UP-INT. But before seeing how to adapt it to become suitable as an updatable AE scheme, including whether it meets our stronger notions of UP-IND and UP-REENC security, we first revisit the claims of UP-IND-BI security from [BLMR15].

As mentioned in the introduction, BLMR claim that the scheme can be shown secure, and sketch a proof of UP-IND-BI security. Unfortunately the proof sketch contains a bug, as we explain below. Interestingly revelation of this bug does not lead to a direct attack on the scheme, and at the same time we could not determine if the proof could be easily repaired. Instead we are able to show that a proof is unlikely to exist.

¹ BLMR called this ciphertext independence, but we reserve that terminology for schemes that do not require ciphertexts during token generation as per Section 2.

² Such a scheme can be constructed by starting with a scheme that satisfies both security notions and adding a “counter” component to ciphertexts that records how many re-encryptions have been performed to obtain that ciphertext; one now exploits the property that any pair of ciphertexts can be input to the ReLR oracle in our UP-REENC game, while only fresh ciphertexts C_0, C_1 are rotated in the BLMR notion.

KeyGen()	Enc(k, m)	ReKeyGen(k_i, k_j, \tilde{C})
$k \leftarrow_s \mathcal{KG}()$	$x \leftarrow_s \mathcal{K}$	$x = \mathcal{D}(k_i, \tilde{C})$
return k	$\tilde{C} \leftarrow_s \mathcal{E}(k, x)$	$x' \leftarrow_s \mathcal{K}$
	$\bar{C} = (m_1 + F(x, 1), \dots, m_\ell + F(x, \ell))$	$\tilde{C}' = \mathcal{E}(k_j, x')$
	return $C = (\tilde{C}, \bar{C})$	$\Delta_{i,j,\tilde{C}} = (\tilde{C}', x' - x)$
		return $\Delta_{i,j,\tilde{C}}$
ReEnc($\Delta_{i,j,\tilde{C}}, C$)	Dec(k, C)	
$(\tilde{C}, \bar{C}) = C$	$(\tilde{C}, \bar{C}) = C$	
$(\tilde{C}', y) = \Delta_{i,j,\tilde{C}}$	$x = \mathcal{D}(k, \tilde{C})$	
$\bar{C}' = (\bar{C}_1 + F(y, 1), \dots, \bar{C}_\ell + F(y, \ell))$	$m = (\bar{C}_1 - F(x, 1), \dots, \bar{C}_\ell - F(x, \ell))$	
return $C = (\tilde{C}', \bar{C}')$	return m	

Fig. 7. The BLMR scheme.

Our main result of this section is the following: giving a proof showing the BLMR UP-IND-BI security would imply the existence of a reduction showing that (standard) IND-CPA security implies circular security [BRS03,CL01] for a simple KEM/DEM style symmetric encryption scheme. The latter seems quite unlikely given the known negative results about circular security [ABBC10,CGH12], suggesting that the BLMR scheme is not likely to be provably secure.

First we recall some basic tools that BLMR use to build their scheme.

Definition 7 (Key-homomorphic PRF [BLMR15]). Consider an efficiently computable function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that (\mathcal{K}, \oplus) and (\mathcal{Y}, \otimes) are both groups. We say that the F is key-homomorphic if for every $k_1, k_2 \in \mathcal{K}$ and every $x \in \mathcal{X}$, $F(k_1, x) \otimes F(k_2, x) = F(k_1 \oplus k_2, x)$.

To define security, let game $\text{PRF1}_F^{\mathcal{A}}$ be the game that selects a key $k \leftarrow_s \mathcal{K}$ and then runs an adversary \mathcal{A} that can adaptively query to an oracle that returns F_k applied to the queried message. The adversary outputs a bit. Let game $\text{PRF0}_F^{\mathcal{A}}$ be the game in which an adversary \mathcal{A} can adaptively query an oracle that returns a random draw from \mathcal{Y} . The adversary outputs a bit. We assume that \mathcal{A} , in either game, never queries the same value twice to its oracle. We define the PRF advantage of \mathcal{A} as

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) = |\Pr[\text{PRF1}_F^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{PRF0}_F^{\mathcal{A}} \Rightarrow 1]|.$$

A simple example of a secure key-homomorphic PRF in the ROM is the function $F(k, x) = k \cdot H(x)$ where $\mathcal{Y} = \mathbb{G}$ is an additive group in which the decisional Diffie–Hellman assumption holds. This construction is originally due to Naor, Pinkas, and Reingold [NPR99].

As an application of key-homomorphic PRFs, BLMR proposed the following construction. The construction follows a similar approach to the AE-hybrid scheme, but by using a key-homomorphic PRF in place of regular encryption the data encryption key can also be rotated.

Definition 8 (BLMR scheme). Let π be a symmetric-key IND-CPA encryption scheme $\pi = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$. Furthermore, let $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ be a key-homomorphic PRF where $(\mathcal{K}, +)$ and $(\mathcal{Y}, +)$ are groups, and where the elements of \mathcal{X} can be represented as integers. The BLMR scheme is the tuple of algorithms $(\text{KeyGen}, \text{Enc}, \text{ReKeyGen}, \text{ReEnc}, \text{Dec})$ depicted in Figure 7.

Note that encryption in the BLMR scheme is a key wrap followed by CTR mode encryption using the wrapped key x and PRF F . For simplicity of presentation we have assumed that a message m can be represented by a sequence m_1, \dots, m_ℓ of elements of \mathcal{Y} .

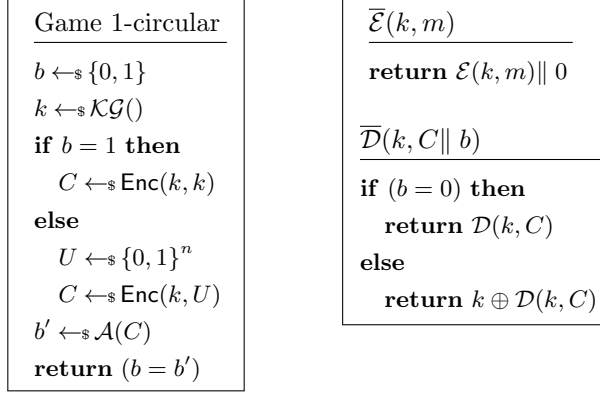


Fig. 8. Left: The 1-circular security game. Right: Definition of $\bar{\mathcal{E}}, \bar{\mathcal{D}}$ used in the proof of Theorem 8.




Fig. 9. Adversary \mathcal{B} for UP-IND using as a subroutine the adversary \mathcal{A} attacking 1-circular security of EncBad .

6.1 Negative Result about Provable UP-IND Security of BLMR

BLMR sketch a proof for the security of this construction in the UP-IND-BI model (as we refer to it). However, the proof misses a subtle point: the interaction with the ReKeyGen oracle behaves similarly to a decryption oracle and the informal argument given that the IND-CPA security of the KEM is sufficient to argue security is wrong. In fact, the BLMR scheme seems unlikely to be provably secure even in our basic security model. To argue this, we show that proving security of the BLMR scheme implies the 1-circular security of a specific KEM/DEM construction. Figure 8 depicts the security game capturing a simple form of 1-circular security for an encryption scheme $\pi = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$.

While our main result here (Theorem 8), can be stated for the BLMR scheme as described earlier, for the sake of simplicity we instead give the result for the special case of using a simple one-time pad DEM instead of the key-homomorphic PRF. This is a trivial example of what BLMR call a key-homomorphic PRG, and their theorem statement covers this construction as well. We will show that proving security for this special case is already problematic, and this therefore suffices to call into question their (more general) theorem. Thus encryption becomes $\text{Enc}(k, m) = (\mathcal{E}(k, r), r \oplus m)$ where \mathcal{E} is an IND-CPA secure KEM. We assume $|m| = n$. We then have $\text{ReKeyGen}(k_1, k_2, \tilde{C}) = (\mathcal{E}(k_2, r'), r' \oplus \mathcal{D}(k_1, \tilde{C}))$. We have the following theorem:

Theorem 8. *Let \mathcal{E} be an IND-CPA-secure symmetric encryption scheme. If there exists a reduction from BLMR UP-IND-BI-security to the IND-CPA security of \mathcal{E} , then there exists a reduction that shows that Enc is 1-circular secure.*

Proof. We start by introducing a slight variant of \mathcal{E} , denoted $\bar{\mathcal{E}}$, shown in Figure 8. It adds a bit to the ciphertext³ that is read during decryption: if the bit is 1 then decryption outputs the secret key xor'd with the plaintext. Let EncBad be the same as Enc above but using $\bar{\mathcal{E}}$, i.e., $\text{EncBad}(k, m) = \bar{\mathcal{E}}(k, r), r \oplus m$ and $\text{ReKeyGenBad}(k_1, k_2, C) = \bar{\mathcal{E}}(k_2, r'), r' \oplus \bar{\mathcal{D}}(k_1, C)$.

If \mathcal{E} is IND-CPA then $\bar{\mathcal{E}}$ is as well. Thus if $\bar{\mathcal{E}}$ is IND-CPA, then the security claim of BLMR implies that EncBad is UP-IND-BI. We will now show that UP-IND-BI security of EncBad implies the 1-circular security of EncBad . In turn it's easy to see that if EncBad is 1-circular secure then so too is Enc , and, putting it all together, the claim of BLMR implies a proof that IND-CPA of \mathcal{E} gives 1-circular security of Enc .

³ Notice that this scheme is not tidy in the sense of [NRS14]. While that doesn't affect the implications of our analysis — BLMR make no assumptions about tidiness — finding a tidy counter-example is an interesting open question.

It remains to show that UP-IND-BI security implies EncBad 1-circular security. Let \mathcal{A} be a 1-circular adversary against EncBad. Then we build an adversary \mathcal{B} against the UP-IND security of EncBad. It is shown in Figure 9. The adversary makes an LR query on a uniform message and the message 0^n . If the UP-IND-BI challenge bit is 1 then it gets back a ciphertext $C_1 = (\bar{\mathcal{E}}(k_1, r) \| 0, r \oplus U)$ and if it is 0 then $C_0 = (\bar{\mathcal{E}}(k_1, r) \| 0, r)$. Next it queries ReKeyGen oracle on the first component of the returned ciphertext but with the trailing bit switched to 1. It asks for a rekey token for rotating from k_1 back to k_1 . The value returned by this query is equal to $\bar{\mathcal{E}}(k_1, r') \| 0, r' \oplus k_1 \oplus r$. By XOR'ing the second component with the second component returned from the LR query the adversary gets finally a ciphertext that is, in the left world, the encryption of k_1 under itself and, in the right world, the encryption of a uniform point under k_1 . Adversary \mathcal{B} runs a 1-circular adversary \mathcal{A} on the final ciphertext and outputs whatever \mathcal{A} outputs. \square

The above result uses 1-circular security for simplicity of presentation, but one can generalize the result to longer cycles with additional queries.

The result is relative, only showing that a proof of BLMR's claim implies another reduction between circular security and IND-CPA security for the particular KEM/DEM scheme Enc above. It is possible that this reduction exists, however it seems unlikely. Existing counter-examples show IND-CPA schemes that are not circular-secure [KRW15]. While these counter-examples do not have the same form as the specific scheme under consideration, it may be that one can build a suitable counter-example with additional effort.

7 An Updatable AE Scheme with Re-encryption Indistinguishability

We first point out that one can avoid the issues raised in Section 6 by replacing the IND-CPA KEM with a proper AE scheme. This does not yet, however, address integrity of the full encryption scheme. To provide integrity overall, we can include a hash of the message in the ciphertext header. However, to prevent this from compromising confidentiality during re-keying, we further mask the hash by an extra PRF output.

This amended construction — which we refer to as ReCrypt — is detailed in Figure 10. It uses an AE scheme $\pi = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$, a key-homomorphic PRF $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$, and a hash function $h : \{0, 1\}^* \rightarrow \mathcal{Y}$. Here $(\mathcal{K}, +)$ and $(\mathcal{Y}, +)$ are groups. A message m is assumed to be a sequence m_1, \dots, m_ℓ of elements in \mathcal{Y} . When we come to discuss specific instantiations in Section 7.2, we will explain how to modify the scheme to handle messages that are binary strings. We also assume that elements of \mathcal{X} can be represented by integers (as in the BLMR scheme).

In the remainder of this section we show that the new scheme meets our strongest security notions for updatable encryption. We then assess the viability of using this scheme in practice, discussing how to instantiate F for high performance and reporting on performance of the full scheme.

7.1 Security of ReCrypt

We state three security theorems for ReCrypt: UP-IND, UP-INT, and UP-REENC notions. The proof of UP-INT relies on the collision resistance of the hash h , while the other two proofs do not. For simplicity, and because we will later instantiate the PRF F in the Random Oracle Model (ROM), we model h as a random oracle throughout our analysis. This modelling of h could be avoided using the approach of Rogaway [Rog06], since concrete collision-producing adversaries can be extracted from our proofs. Note also that the *almost* key-homomorphic PRF construction in the standard model presented by BLMR would not achieve UP-REENC since the number of re-encryptions is leaked by the ciphertext, allowing an adversary to distinguish two re-encryptions.

Theorem 9 (UP-IND security of ReCrypt). *Let $\pi = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$ be an AE scheme, $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ be a key-homomorphic PRF, $h : \{0, 1\}^* \rightarrow \mathcal{Y}$ be a hash function, and let Π be the ReCrypt scheme as depicted in Figure 10.*

Then for any adversary \mathcal{A} against Π , there exist adversaries \mathcal{B}, \mathcal{C} such that

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) \leq 2t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{C})$$

for all $\kappa \geq 0, t \geq 1$.

<u>KeyGen()</u>	<u>Enc(k, m)</u>	<u>ReKeyGen(k_i, k_j, \tilde{C})</u>
$k \leftarrow \mathcal{K}\mathcal{G}()$	$x, y \leftarrow \mathcal{K}$	$(\chi, \tau) = \mathcal{D}(k_i, \tilde{C})$
return k	$\chi = x + y$	if $(\chi, \tau) = \perp$ return \perp
	$\tau = h(m) + F(x, 0)$	$x', y' \leftarrow \mathcal{K}$
	$\tilde{C} = \mathcal{E}(k, (\chi, \tau))$	$\chi' = \chi + x' + y'$
	for $1 \leq l \leq \ell$	$\tau' = \tau + F(x', 0)$
	$\bar{C}_l = m_l + F(x, l)$	$\tilde{C}' \leftarrow \mathcal{E}(k_j, (\chi', \tau'))$
	return $(\tilde{C}, \bar{C} = (y, \bar{C}_1, \dots, \bar{C}_\ell))$	return $\Delta_{i,j,\tilde{C}} = (\tilde{C}', x', y')$
<u>ReEnc($\Delta_{i,j,\tilde{C}}, (\tilde{C}, \bar{C})$)</u>		<u>Dec($k, (\tilde{C}, \bar{C})$)</u>
$(\tilde{C}', x', y') = \Delta_{i,j,\tilde{C}}$		$(\chi, \tau) \leftarrow \mathcal{D}(k, \tilde{C})$
$y = \bar{C}_0$		if $(\chi, \tau) = \perp$ return \perp
for $1 \leq l \leq \ell$		$y = \bar{C}_0$
$\bar{C}'_l = \bar{C}_l + F(x', l)$		for $1 \leq l \leq \ell$
return $(\tilde{C}', \bar{C}' = (y + y', \bar{C}'_1, \dots, \bar{C}'_\ell))$		$m_l = \bar{C}_l - F(\chi - y, l)$
		if $h(m) + F(\chi - y, 0) = \tau$ then
		return $m = (m_1, \dots, m_\ell)$
		else
		return \perp

Fig. 10. The ReCrypt scheme.

Proof. We split the proof into two parts. The first part uses the AE security of π to show that the value of x used to key the key-homomorphic PRF is hidden from the adversary. The second part uses the fact that F is a PRF to show indistinguishability holds, given that the adversary cannot learn x .

Let G_0 be the original UP-IND game. For $1 \leq i \leq t$, Game G_i is the same as G_{i-1} , except we replace the encryption $\mathcal{E}(k_i, (\chi, h(m) + F(x, 0)))$ with a random string r and store the tuple (x, y, m) as a lookup in the table \mathbb{C} indexed by the random string r . For queries to $\text{ReKeyGen}(i, j, \tilde{C})$, if $\tilde{C} = r$ for some previously returned r , then compute x', y' as usual and either:

- if $j \leq i$, return (r', x', y') for some random r' , storing the value $(x + x', y + y', m)$, or
- if $j > i$, return $(\mathcal{E}(k_j, (\chi', h(m) + F(x + x', 0))), x', y')$.

If \tilde{C} has not been previously seen, return \perp . These modifications are shown in Figure 11.

Let S_i be the event that \mathcal{A} outputs the correct bit in game G_i . Then we can construct an adversary \mathcal{B} such that $|\Pr[S_i] - \Pr[S_{i-1}]| \leq \text{Adv}_\pi^{\text{ae}}(\mathcal{B})$. To see this, notice that by replacing the $\mathcal{E}(k_i, \cdot)$ and $\mathcal{D}(k_i, \cdot)$ functions with calls to an instance of the AE game for Π , in the real world, we get G_{i-1} , and the random world corresponds to G_i .

In game G_t , suppose the attacker queries the LR oracle to receive a challenge ciphertext C . Then the header \tilde{C} will be equal to some random string r , while the body \bar{C} consists of the key mask y and the message encrypted by the PRF F keyed by some value x . Note that the use of AE forces the adversary to submit only previously seen ciphertext headers to oracles with $i \leq t$.

From re-keygen queries, the attacker can learn the fresh values of x', y' , and r' . However, these are insufficient to learn anything about the value of x used to encrypt the message.

Similarly, re-encryption queries where the target key j is uncompromised, i.e., $j \leq t$, the adversary receives a fresh ciphertext encrypted under x' with no way of learning x' .

On the other hand, if $j > t$, the invalidity condition $\text{Invalid}_{\text{RE}}$ will be true, so the adversary learns the ciphertext header. This is of the form $\mathcal{E}(k_j, (x + y + x' + y', h(m) + F(x + x', 0)))$. Since the adversary possesses k_j , they may decrypt the ciphertext header. However, x is still masked by y .

From here, we proceed similarly to the original proof outline given by Boneh et al. We construct a game G_{t+1} such that $|\Pr[S_{t+1}] - \Pr[S_t]| \leq \text{Adv}_F^{\text{prf}}(\mathcal{C})$.

<u>Enc(i, m)</u>	<u>ReKeyGen(i, j, \tilde{C})</u>	<u>Dec($i, (\tilde{C}, \bar{C})$)</u>
$x, y \leftarrow_{\$} \mathcal{K}$	$(x, y, m) \leftarrow \mathbb{C}[\tilde{C}]$	$(x, y, m) \leftarrow \mathbb{C}_i[\tilde{C}]$
$r \leftarrow_{\$} \{0, 1\}^{ \tilde{C} }$	if $(x, y, m) = \perp$ return \perp	if $(x, y, m) = \perp$ return \perp
$\mathbb{C}_i[r] = (x, y, m)$	$x', y' \leftarrow_{\$} \mathcal{K}$	$y' = \bar{C}_0$
$\bar{C}_0 = y$	$\chi' = x + y + x' + y'$	$x' = x + y - y'$
for $1 \leq l \leq \ell$	if $j \leq i$ then	for $1 \leq l \leq \ell$
$\bar{C}_l = m_l + F(x, l)$	$r \leftarrow_{\$} \{0, 1\}^{ \tilde{C} }$	$m'_l = \bar{C}_l - F(x', l)$
return $(\tilde{C} = r, \bar{C})$	$\mathbb{C}_j[r] = (x + x', y + y', m)$	$\tau = h(m) + F(x, 0)$
	$\tilde{C}' = r$	if $\tau - F(x', 0) = h(m')$ then
	else	return $m = (m_1, \dots, m_\ell)$
	$\tau = h(m) + F(x + x', 0)$	else
	$\tilde{C}' \leftarrow_{\$} \mathcal{E}(k_j, (\chi', \tau))$	return \perp
	return $\Delta_{i,j,\tilde{C}} = (\tilde{C}', x', y')$	

Fig. 11. The replacement algorithms used in game G_i for the proofs of security for the ReCrypt construction. For the i -th key, encryption/decryption by Π is replaced by random strings r of the same length and a lookup.

When $i \leq t$, we replace the usage of $F(x, l)$ in the LR oracle by using the PRF challenge oracle, denoted $f(l)$. Suppose K is the key used by the PRF challenger. When the PRF challenge is from the real world, the computed encryption is of the form $m_l + F(x, l) + f(l) = m_l + F(K + x, l)$. While the ciphertext header includes the value $\tau = h(m) + F(K + x, 0)$.

Since we showed that the value of x is unknown to the adversary for challenge ciphertexts, this change perfectly models the game G_t .

Furthermore, the computation of re-keying tokens and re-encryptions does not remove the PRF mask f from the ciphertext: if j is uncorrupted, then we compute $\bar{C}_l + F(x', l) = m_l + F(x + x' + K, l)$; otherwise, either $\text{Invalid}_{\text{RE}}$ or $\text{Invalid}_{\text{RK}}$ will be false, and the oracle returns \perp for the ciphertext body. The adversary can also obtain the ciphertext header containing an encryption of $(x + x', \tau = h(m) + F(K + x + x', 0))$.

Now, let G_{t+1} correspond to the event that the PRF challenge returns a random value. Then the message is always perfectly masked by the output of $f(l)$ – we showed earlier that this value is present even after a re-encryption is computed. Therefore both the message and the hash are perfectly masked by the PRF values, and the adversary cannot win with a probability greater than $\frac{1}{2}$. Hence, we conclude that:

$$\begin{aligned} \text{Adv}_{\Pi, \kappa, t}^{\text{up-ind}}(\mathcal{A}) &= |2 \cdot \Pr[S_0] - 1| \\ &\leq 2t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{C}) \end{aligned}$$

for all $\kappa \geq 0, t \geq 1$. □

Theorem 10 (UP-INT security of ReCrypt). *Let $\pi = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$ be an AE scheme, $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ be a key-homomorphic PRF, h be a cryptographic hash function modelled as a random oracle with outputs in \mathcal{Y} , and let Π be the ReCrypt scheme as depicted in Figure 10.*

Then for any adversary \mathcal{A} against Π , there exists an adversary \mathcal{B} such that

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-int}}(\mathcal{A}) \leq 2t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + \frac{q^2 + q_h^2}{|\mathcal{Y}|} + \frac{q^2}{|\mathcal{X}| \cdot |\mathcal{Y}|}$$

for all $\kappa \geq 0, t \geq 1$, where the adversary makes q_h queries to h , and q oracle queries.

Proof. Apply the same t transformations shown in Figure 11, which substitutes encryption by π with random strings for uncorrupted keys. Decryption is then simulated by using a lookup of previously returned values. Let S_i be the event that \mathcal{A} outputs the correct bit in game G_i .

In game G_{t+1} , we replace the Dec verification check $\tau - F(x', 0) = h(m')$ and instead directly verify that $m' = m$. These two games are identical, unless the adversary submits for verification a ciphertext $C = (\tilde{C}, \bar{C})$ such that \tilde{C} has previously been generated for the tuple (x, y, m) and \bar{C} , decrypts using $\chi - y$ to a message m' such that $m' \neq m$ but $h(m) + F(x, 0) - F(x + y - y') = h(m')$. That is, $h(m) + F(y - y', 0) = h(m')$. For all adversarially chosen y' , the adversary needs to construct a ciphertext \bar{C}' such that $h(\bar{C}'_1 - F(x + y - y', 1), \dots, \bar{C}'_\ell - F(x + y - y', \ell))$ is equal to some fixed value $h(m) - F(y - y', 0)$. This is equivalent to the pre-image resistance of h . Alternatively, if the adversary does not modify y , then this reduces to finding a value of m' such that $h(m) = h(m')$, which is just the collision resistance of h .

By modelling h as a random oracle, we can bound the probability that the adversary making q_h queries to the random oracle, succeeds in finding suitable values of y', m , and m' by $q_h^2/|\mathcal{Y}|$. Therefore $|\Pr[S_{t+1}] - \Pr[S_t]| \leq q_h^2/|\mathcal{Y}|$.

Let (i, C) be a tuple submitted by the adversary to the Try oracle. We show that with high probability, either $\text{InvalidCTXT}(i, C) = 1$ or $\text{Dec}(k_i, C) = \perp$. By the first t game hops, we know that $\mathcal{D}(k_i, \tilde{C}) = \perp$ unless \tilde{C} was previously output by an oracle query (whether Enc or ReKeyGen). Furthermore, for all \tilde{C} , with high probability, there exists precisely one \bar{C} . For $1 \leq l \leq \ell$, decryption is computed as $\bar{C}_l + F(x, l)$. Therefore, any modification to a ciphertext block will result in a distinct message, which will result in a failed decryption due to the equality check. The exception to this is when the same random value \tilde{C} was generated for distinct messages. If the adversary makes a total of q queries, this happens with probability $q^2/2^{|\tilde{C}|} \leq q^2/2^{h(x)+|x|} \leq q^2/2^{\log|\mathcal{X}|+\log|\mathcal{Y}|} \leq q^2/(|\mathcal{X}| \cdot |\mathcal{Y}|)$.

This shows that for each \tilde{C} , there can only be one \bar{C} . But what about the converse? For a given ciphertext body \bar{C} , suppose there exists (x', m') such that $\bar{C}_l = m_l + F(x, l) = m'_l + F(x', l)$. Therefore $F(x' - x, l) = m_l - m'_l$. However, m' was submitted to the (re-)encryption oracle before x' was generated. Therefore, for any pair $(x, m), (x', m')$ the probability that $F(x' - x, l) = m_l - m'_l$ is $1/|\mathcal{Y}|$, and thus if the adversary makes q queries, the probability is bounded by $q^2/|\mathcal{Y}|$.

Together with the probability that two \bar{C} are equal, this becomes the probability the adversary wins in game G_{t+1} , otherwise we have established that all valid ciphertext headers have precisely one matching ciphertext body, in which case $\text{InvalidCTXT}(i, C)$ is true.

Therefore, the advantage of \mathcal{A} is:

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-int}}(\mathcal{A}) \leq t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + \frac{q^2 + q_h^2}{|\mathcal{Y}|} + \frac{q^2}{|\mathcal{X}| \cdot |\mathcal{Y}|}.$$

□

Finally, we prove that ReCrypt meets our re-encryption indistinguishability notion.

Theorem 11 (UP-REENC security of ReCrypt). *Let $\pi = (\text{KG}, \mathcal{E}, \mathcal{D})$ be an AE scheme, $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ be a key-homomorphic PRF, and let Π be the ReCrypt scheme as depicted in Figure 10.*

Then for any adversary \mathcal{A} against Π , there exist adversaries \mathcal{B}, \mathcal{C} such that

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-reenc}}(\mathcal{A}) \leq 2t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{C})$$

for all $\kappa \geq 0, t \geq 1$.

Proof. Apply the same t steps as used in the previous proof, shown in Figure 11. Then we have the ciphertext headers replaced with random strings, and the re-keygen process is simulated by recording inputs previously seen.

Hence, in game G_t , the adversary cannot learn anything about the PRF key x used in the challenge, nor does the ciphertext header reveal anything for compromised keys, since all values are masked by values stored in the ciphertext body.

As before, we construct a PRF adversary \mathcal{C} to finish the game. Suppose in G_t we replace the re-encryption routine from challenges to return: $(r, \bar{C}_b + f(l) + F(x' + y', l))$. That is, the ciphertext header \tilde{C} is a random string, by the previous hybrids, and we re-encrypt the ciphertext body \bar{C}_b to use (implicitly) key $\chi + K + x'$, where K is the key used by the PRF challenger. The re-keying token would be $(r, K + x', y')$, however, since we do not return this to the adversary, we do not need to compute it, nor does the adversary learn $K + x'$. Any further re-encryption or re-keygen queries by the adversary preserves the usage of $f(l)$ as a mask.

Now we let game G_{t+1} be the scenario in which the PRF challenger returns random values, and we conclude that the adversary cannot learn any information about the bit b .

The same computation as before results in:

$$\text{Adv}_{H,\kappa,t}^{\text{up-reenc}}(\mathcal{A}) \leq 2t \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B}) + 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{C}).$$

□

7.2 Implementation and Performance

We dedicate the remainder of this section to analysis of the viability of ReCrypt for use in practice.

We first must instantiate the key-homomorphic PRF used by ReCrypt. While BLMR suggest using a standard model construction, a more efficient route is to use the classic ROM construction due originally to Naor, Pinkas, and Reingold [NPR99]. Here one uses $F(k, x) = k \cdot H(x)$ where H is modelled as a random oracle $H : \mathcal{X} \rightarrow \mathbb{G}$, \mathcal{X} is any suitable set (bit-strings of some fixed length, for example), and $(\mathbb{G}, +)$ is a group in which the decisional Diffie–Hellman (DDH) assumption holds. We will use as group \mathbb{G} of subset of the points on an elliptic curve $E(\mathbb{F}_p)$ defined over a prime field. In our implementation we use the specific curve Curve25519 [Ber06], which has $p = 2^{255} - 19$. We use SHA256 for the hash function h and AES128-GCM for the AE scheme π .

Message encoding. Recall that encryption is done block-wise as $\bar{C}_l = m_l + F(x, l)$, where it was assumed that messages m are already represented as sequences over the group \mathbb{G} . To make a practical scheme for encrypting data represented as bitstrings, we therefore require an encoding function $\sigma_m : \{0, 1\}^n \rightarrow \mathbb{G}$ for some block length n . Padding can be used to ensure messages are of length a multiple of n .

For a suitable message encoding function, we require the function and its inverse to be efficiently computable. In order to avoid side-channel leakage on plaintexts, we require the function to be amenable to constant-time implementation. In addition, security mandates that σ_m^{-1} can be implemented in such a way that it rejects on input any element that it is outside the range of σ_m . For otherwise, σ_m^{-1} may not act as a bijection on its inputs; the existence of two points P, P' on the curve such that $\sigma_m^{-1}(P) = \sigma_m^{-1}(P')$ might then allow an adversary to construct a ciphertext forgery.

We use the Elligator function [BHK13] to realize σ_m . For elliptic curve groups including Curve25519 [Ber06], the Elligator function produces an injective mapping from $\{0, 1, \dots, \frac{p-1}{2}\}$ to points on the curve. The inverse mapping returns elements of the form $r, p - r$; we obtain a bijection by inverting the map and taking as output the value that is less than $\frac{p-1}{2}$.

Using Elligator, the natural block size for messages would be $\lfloor \log(p-1) - 1 \rfloor = 254$ bits, which is not an integral number of bytes. However, splitting bytes would have a deleterious impact on performance, and so we actually work with a 31 byte (248 bit) block size. Hence, on recovering the message using σ_m^{-1} , we must additionally check that the top byte is zero.

Hashing to the curve. In order to instantiate the key-homomorphic PRF, we require a hash function H outputting elements in the group \mathbb{G} . Here we are less concerned about side-channel attacks because the inputs to H in our scheme are not secret. The natural way to build H would be to take a regular hash function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$, modelled as a random oracle, and apply any encoding function σ_m , resulting in $H(x) := \sigma_m(h(x))$. However, for the function H to additionally act like a random oracle, we would require at least that σ_m maps uniformly to elements in \mathbb{G} , imposing additional requirements on σ_m . We therefore require a second encoding function σ_h for constructing the key-homomorphic PRF, and require that σ_h composed with a random oracle h is indifferentiable [MRH04] from a random oracle.

An in-depth treatment of indifferentiable hashing onto elliptic curves appears in [BCI⁺10]. For simplicity, we opted to use their general approach: $\sigma_h(m) = f(h_1(m)) + h_2(m) \cdot G$. Here h_1 and h_2 are cryptographic hash functions suitable to be modeled as random oracles, instantiated in our implementation by SHA-256 with appropriate domain separation. The function f needs to be what [BCI⁺10] refer to as a weak encoding function. Elligator satisfies the necessary properties, as it is polynomial time computable, the probability of any point being output is either 0 or $2/\#\mathbb{G}$, and it is samplable by simply picking from $\pm f^{-1}(P)$ for any $P \in E(\mathbb{F}_p)$.

Point compression. Curve25519 supports an x -coordinate-only Montgomery ladder for scalar multiplication. Since we want to unambiguously add points rather than just computing scalar multiples, we need to work

ReCrypt Operation	1 block	Time per CPU				cycles/byte
		1 KB	1 MB	1 GB		
Encrypt	353 μ s	8.5 ms	8.9 s	2.5 hours	30.7 K	
ReEnc	239 μ s	7.3 ms	7.2 s	2.0 hours	26.3 K	
Decrypt	328 μ s	7.7 ms	7.5 s	2.2 hours	27.9 K	
ReKeyGen (total)		178 μ s			2.3 M	

Fig. 12. Processing times for ReCrypt operations measured on a 3.8GHz CPU. 1 block represents any plaintext \leq 31 bytes. Number of iterations: 1000 (for 1 block, 1 KB), 100 (for 1 MB) and 1 (for 1 GB). Cycles per byte given for 1MB ciphertexts.

with both x and y coordinates. Therefore, we used the twisted Edwards form of Curve25519, implemented in [Lov], which in turn is based on [Lan]. We also used the latter to guide our Elligator implementation.

We save bandwidth using point compression. When a curve point is serialized, only the y coordinate and the sign of the x coordinate (1-bit) needs to be recorded. Since the y coordinate requires less than the full 32 bytes, we are able to serialize points as 32 byte values. Each 32 byte serialized value represents 31 bytes of plaintext, giving a ciphertext expansion of 3%. Upon deserialization, the x coordinate must be recomputed. This requires computing a square root, taking approximately 20 μ s. Of course this cost could be avoided by instead serializing both x and y coordinates. This would create a 64 byte ciphertext for each 31 bytes of plaintext, an expansion of 106%. We consider that to be unacceptable.

Microbenchmarks. We built our reference implementation using the Rust [MKI14] programming language. Our implementation is single-threaded and we measured performance on an Intel CPU (Haswell), running at 3.8GHz in turbo mode.

Figure 12 shows wall clock times for ReCrypt operations over various plaintext sizes. As might be expected given the nature of the cryptographic operations involved, performance is far from competitive with conventional AE schemes. For comparison, AES-GCM on the same hardware platform encrypts 1 block, 1 KB, 1 MB and 1 GB of plaintext in 15 μ s, 24 μ s, 9 ms, and 11 s, respectively. KSS has performance determined by that of AES-GCM, while the performance of the ReCrypt scheme is largely determined by the scalar multiplications required to evaluate the PRF. Across all block sizes there is a 1000x performance cost to achieve our strongest notion of security.

Discussion. Given the performance difference, ReCrypt is best suited to very small or very valuable plaintexts (ideally, both). If the plaintext corpus is moderately or very large, cost and performance may prohibit practitioners from using ReCrypt over more performant schemes like KSS that give strictly weaker security. To make this discussion concrete, we consider two examples in more detail.

For privacy and security reasons, regulation mandates that credit card numbers must be stored in encrypted form [PCI16]. These include the recommendation that a mechanism must be in place to rotate keys on a regular basis and in the face of known or suspected compromise. Given the sensitive nature of payment information, the naive solution of decrypting and encrypting the data to rotate keys exposes it to some risk, since it makes the data available in plaintext form for at least a period of time. Similarly, NIST guidelines [Bar16] recommend balancing the risk introduced by the re-encryption process with the benefits offered by key rotation. On the other hand, our updatable AE schemes enable secure rotation of keys using an untrusted storage service.

Consider a payment system with, say, 1 billion credit card entries. The storage required for encrypted credit card numbers using ReCrypt is about 30 GB (assuming one block per entry). Projecting from performance measurements in Figure 12, a full key rotation across the this dataset requires 60 CPU-hours: a significant amount of computation, but potentially not prohibitive given infrequent rotations and many available CPUs. The time to decrypt a single entry is sub-millisecond; this is small compared to processing time for a credit card transaction.

As a second example we consider long-term storage of static data, commonly known as “deep” or “cold” storage. Such data is accessed infrequently, yet data owners may still desire (or be required to) periodically rotate the encryption keys used for protecting the data. In such cases it may be more convenient for the data owner to allow the storage provider to rotate the encryption keys using a system local to the data, as opposed to the data owner retrieving the data and performing the re-encryption.

For a rough estimation of costs, we compute the cost of performing re-encryption using ReCrypt on Amazon Web Service’s Elastic Compute Cloud (AWS EC2). Using the price per CPU-hour of an AWS EC2 instance of 0.05 \$/hour (USD)⁴, we compute the cost to perform updates using ReCrypt as 0.10 \$/GB. For small- to medium-sized data sets or for data sets that are particularly valuable (e.g. financial information), this cost may be justified. However, for moderately-sized to large data sets, the cost may be prohibitive and clients may favor basic security schemes like KSS.

8 Conclusion and Open Problems

We have given a systematic study of updatable AE, providing a hierarchy of security notions meeting different real-world security requirements and schemes that satisfy them efficiently. Along the way, we showed the limitations of currently deployed approach, as represented by AE-hybrid, improved it at low cost to obtain the KSS scheme meeting our UP-IND and UP-INT notions, identified a flaw in the BLMR scheme, repaired it, and showed how to instantiate the repaired scheme in the ROM. Through this, we arrived at ReCrypt, a scheme that is secure in our strongest security models (UP-IND, UP-INT and UP-REENC). We implemented ReCrypt and presented basic performance benchmarks for our prototype. The scheme is slower than the hybrid approaches but offers true key rotation.

Our work puts updatable AE on a firm theoretical foundation and brings schemes with improved security closer to industrial application. While there is a rich array of different security models for practitioners to choose from, it is clear that achieving strong security (currently) comes at a substantial price. Meanwhile, weaker but still useful security notions can be achieved at almost zero cost over conventional AE. It is an important challenge to find constructions which lower the cost compared to ReCrypt without reducing security. But it seems that fundamentally new ideas are needed here, since what are essentially public key operations are intrinsic to our construction.

From a more theoretical perspective, it would also be of interest to study the exact relations between our security notions, in particular whether UP-REENC is strong enough to imply UP-IND and UP-INT. There is also the question of whether a scheme that is UP-REENC is necessarily ciphertext-dependent. Finally, we reiterate the possibility of formulating updatable AE in the nonce-based setting.

Acknowledgements

We thank the anonymous Crypto reviewers for their feedback and discussion relating to a bug in our initial KSS construction, and for their additional comments and suggestions. This work was supported in part by NSF grant CNS-1330308, EPSRC grants EP/K035584/1 and EP/M013472/1, by a research programme funded by Huawei Technologies and delivered through the Institute for Cyber Security Innovation at Royal Holloway, University of London, and a gift from Microsoft.

References

- ABBC10. Tolga Acar, Mira Belenkiy, Mihir Bellare, and David Cash. Cryptographic agility and its relation to circular encryption. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 403–422. Springer, Heidelberg, May 2010.
- ABL⁺14. Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. How to securely release unverified plaintext in authenticated encryption. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 105–125. Springer, Heidelberg, December 2014.
- AWS. AWS. Protecting data using client-side encryption. <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html>.
- Bar16. Elaine Barker. Recommendation for key management Part 1: General. In *NIST special publication 800-57 Part 1 Revision 4*, 2016.

⁴ This rate is based on the lowest per-CPU cost on AWS as of June 2017. Namely, an m4.16xlarge instance available in the AWS US-East (Ohio) region which provides 64 CPUs and is available on-demand for 3.20 \$/hour.

- BCI⁺10. Eric Brier, Jean-Sébastien Coron, Thomas Icart, David Madore, Hugues Randriam, and Mehdi Tibouchi. Efficient indifferentiable hashing into ordinary elliptic curves. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 237–254. Springer, Heidelberg, August 2010.
- BDPS14. Alexandra Boldyreva, Jean Paul Degabriele, Kenneth G. Paterson, and Martijn Stam. On symmetric encryption with distinguishable decryption failures. In Shiho Moriai, editor, *FSE 2013*, volume 8424 of *LNCS*, pages 367–390. Springer, Heidelberg, March 2014.
- Ber06. Daniel J. Bernstein. Curve25519: New Diffie-Hellman speed records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 207–228. Springer, Heidelberg, April 2006.
- BHKL13. Daniel J. Bernstein, Mike Hamburg, Anna Krasnova, and Tanja Lange. Elligator: elliptic-curve points indistinguishable from uniform random strings. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 967–980. ACM Press, November 2013.
- BK03. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 491–506. Springer, Heidelberg, May 2003.
- BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Heidelberg, August 2013.
- BLMR15. Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. Cryptology ePrint Archive, Report 2015/220, 2015. <http://eprint.iacr.org/2015/220>.
- BN00. Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, Heidelberg, December 2000.
- BRS03. 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, *SAC 2002*, volume 2595 of *LNCS*, pages 62–75. Springer, Heidelberg, August 2003.
- CGH12. David Cash, Matthew Green, and Susan Hohenberger. New definitions and separations for circular security. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 540–557. Springer, Heidelberg, May 2012.
- CH07. Ran Canetti and Susan Hohenberger. Chosen-ciphertext secure proxy re-encryption. In Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors, *ACM CCS 07*, pages 185–194. ACM Press, October 2007.
- CK05. DL Cool and Angelos D Keromytis. Conversion and proxy functions for symmetric key ciphers. In *Information Technology: Coding and Computing, 2005. ITCC 2005. International Conference on*, volume 1, pages 662–667. IEEE, 2005.
- CL01. Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 93–118. Springer, Heidelberg, May 2001.
- Goo. Google. Managing data encryption. <https://cloud.google.com/storage/docs/encryption>.
- ID03. Anca Ivan and Yevgeniy Dodis. Proxy cryptography revisited. In *NDSS 2003*. The Internet Society, February 2003.
- KRW15. Venkata Koppula, Kim Ramchen, and Brent Waters. Separations in circular security for arbitrary length key cycles. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 378–400. Springer, Heidelberg, March 2015.
- Lan. Adam Langley. ed25519 for go. <http://godoc.org/github.com/agl/ed25519>.
- Lov. Isis Agora Lovecruft. curve25519-dalek. <https://github.com/isislovecruft/curve25519-dalek/>.
- MKI14. Nicholas D Matsakis and Felix S Klock II. The rust language. *ACM SIGAda Ada Letters*, 34(3):103–104, 2014.
- MRH04. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, February 2004.
- NPR99. Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and KDCs. In Jacques Stern, editor, *EUROCRYPT’99*, volume 1592 of *LNCS*, pages 327–346. Springer, Heidelberg, May 1999.
- NRS14. Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering generic composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, Heidelberg, May 2014.
- PCI16. PCI Security Standards Council. Requirements and security assessment procedures. In *PCI DSS v3.2*, 2016.

- PRS11. Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 372–389. Springer, Heidelberg, December 2011.
- Rog06. Phillip Rogaway. Formalizing human ignorance. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 2006, First International Conference on Cryptology in Vietnam, Hanoi, Vietnam, September 25–28, 2006, Revised Selected Papers*, volume 4341 of *Lecture Notes in Computer Science*, pages 211–228. Springer, 2006.
- RS06. Phillip Rogaway and Thomas Shrimpton. Deterministic authenticated-encryption: A provable-security treatment of the key-wrap problem. *Cryptology ePrint Archive*, Report 2006/221, 2006. <http://eprint.iacr.org/2006/221>.
- RSS11. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indistinguishability framework. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, Heidelberg, May 2011.

A Bidirectional Updatable AE

A.1 XOR-KEM: A Bidirectional Updatable AE Scheme

The AE-hybrid and KSS schemes are unidirectional and ciphertext-dependent. This means that in practice the client must fetch, from storage, ciphertext headers prior to computing rekey tokens for updating ciphertexts. It would be simpler to utilize a ciphertext-independent scheme that has rekey tokens that work for any ciphertext encrypted with a particular key. This would make the re-encryption process “non-interactive”, requiring that the key holder only push a single rekey token to the place where ciphertexts are stored. Given the obvious performance benefits that such a scheme would have, we also provide such a scheme, called XOR-KEM. This scheme is exceptionally fast, and is built from a (non-updatable) AE scheme that is assumed to be secure against a restricted form of related-key attack (RKA). This latter notion adapts the Bellare-Kohno RKA-security notions for block ciphers [BK03] to the setting of AE schemes. To the best of our knowledge, this definition is novel, and RKA secure AE may itself be of independent interest as a primitive. However, the XOR-KEM scheme cannot meet our integrity notions against an attacker in possession of compromised keys. (And because of its bidirectionality, XOR-KEM also provides the counter-example that we used to separate UP-IND-BI and UP-IND security in Section 3.1.)

Let $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme. Then we define the ciphertext-independent scheme, XOR-KEM, as follows:

- **KeyGen()**: **return** $k \leftarrow \mathcal{K}$
- **Enc** (k, m) : $x \leftarrow \mathcal{K}$; $C \leftarrow (x \oplus k, \mathcal{E}(x, M))$; **return** C
- **ReKeyGen** (k_1, k_2) : **return** $\Delta_{1,2} = k_1 \oplus k_2$
- **ReEnc** $(\Delta_{1,2}, C = (C_0, C_1))$: $C' \leftarrow (\Delta_{1,2} \oplus C_0, C_1)$; **return** C'
- **Dec** $(k, C = (C_0, C_1))$: **return** $\mathcal{D}(C_0 \oplus k, C_1)$

The XOR-KEM scheme has a similar format to the AE-hybrid scheme above. However, instead of protecting the DEM key x by encrypting it, we instead XOR it with the secret key k . The resulting scheme becomes a bidirectional, ciphertext-independent scheme, and one that has extremely high performance and deployability.

Note that although the value $x \oplus k$ fulfils a similar purpose as the ciphertext header in AE-hybrid, since this value is not needed in re-keying, it resides in the ciphertext body.

Theorem 12 (UP-IND-BI Security of XOR-KEM). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and let Π be the updatable AE scheme XOR-KEM built using π as above. Then for any UP-IND-BI adversary \mathcal{A} against Π , there exists an AE adversary \mathcal{B} against π such that:*

$$\text{Adv}_{\Pi, \kappa, t}^{\text{up-ind-bi}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\pi}^{\text{ae}}(\mathcal{B})$$

for all $\kappa \geq 0, t \geq 1$.

Proof. We consider a sequence of games G_0, \dots, G_{q+1} , where q is the number of queries made by \mathcal{A} to its LR oracle.

Let game G_0 correspond to the regular UP-IND-BI game. Let S_i correspond to the event that game G_i outputs **true**.

In game G_1 , a key $x^* \leftarrow_s \mathcal{K}$ is generated at the start of the game. This is used to encrypt the message in the first query made to the LR oracle. That is, on input (i, m_0, m_1) , the LR oracle computes:

$$x \leftarrow_s \mathcal{K}; \bar{C} \leftarrow_s (x \oplus k_i, \mathcal{E}(x^*, m_b)),$$

and returns \bar{C} .

All other queries are answered as in G_0 . The adversary has the same view in both games, unless the adversary recovers k_i (in which case the adversary can win with one extra query). Either way, we have $\Pr[S_0] = \Pr[S_1]$.

For $1 < \tau \leq q$, game G_τ is identical to $G_{\tau-1}$ except for the adversary's τ -th LR query. There, encryption is computed as:

$$x_\tau \leftarrow_s \mathcal{K}; \bar{C} \leftarrow_s (x_\tau \oplus k_i, \mathcal{E}(x^*, m_b))$$

where x^* is the same key generated at the start of G_1 . As before, these games are identical to the adversary, and therefore $\Pr[S_\tau] = \Pr[S_{\tau-1}]$.

In game G_{q+1} , we replace encryption by x^* with randomly sampled values. It is straightforward to construct an adversary \mathcal{B} in the AE-ROR game such that $|\Pr[S_{q+1}] - \Pr[S_q]| \leq \text{Adv}_\pi^{\text{ae}}(\mathcal{B})$.

Finally, in G_{q+1} the adversary can learn nothing about which m_b is encrypted, therefore $\Pr[S_{q+1}] = \frac{1}{2}$. Combining the above, we get the stated result. \square

XOR-KEM does not provide integrity guarantees in the face of compromised keys: an attacker who learns both $C = \text{Enc}(k_1, m)$ and $C' = \text{ReEnc}(\Delta_{1,t+1}, \bar{C}, C)$ can derive k_1 .

For ciphertext integrity of bidirectional schemes, we modify the `InvalidCTXT` conditions to check whether the ciphertext is a re-encryption using a bidirectional rekey token. We call this new predicate `InvalidCTXTBI` which returns **true** whenever `InvalidCTXT` returns **true**, and additionally returns **true** if:

- (\tilde{C}, \bar{C}) is the ciphertext output by running $\text{ReEnc}(\text{Invert}(\Delta_{j,i}), C')$ for $C' = (\tilde{C}', \bar{C}')$ where $\Delta_{j,i}$ was the result of a query $\text{ReKeyGen}(j, i)$ and $\text{InvalidCTXT}_{\text{BI}}(i, C') = \text{true}$.

We refer to the security game using `InvalidCTXTBI` as UP-INT-BI.

Unfortunately, proving that XOR-KEM even achieves UP-INT-BI security for no compromised keys ($\kappa = 0$) is not straightforward. It is clear that the adversary can produce trivial manipulations of the ciphertext header: $\tilde{C} \oplus z = (x \oplus k) \oplus z = (x \oplus z) \oplus k$. Thus, to prove UP-INT-BI security additionally requires us to assume that the AE scheme π used in the construction is secure against related-key attacks, in which the adversary can access the encryption and decryption functions of π under XOR-offsets of the unknown key.

Definition 9 (Related-Key Secure AE). Let $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an authenticated encryption scheme with keyspace $\mathcal{S}_{\mathcal{K}}$ and let Φ be a set of functions $\Phi = \{\phi : \mathcal{S}_{\mathcal{K}} \rightarrow \mathcal{S}_{\mathcal{K}}\}$. Let the Φ -restricted RKA security game be the game in which the adversary has access to a pair of oracles Enc, Dec which, on input (ϕ, x) , $\phi \in \Phi$, return for $b = 0$:

$$\text{Enc}(\phi, x) = \mathcal{E}(\phi(k), x), \text{Dec}(\phi, x) = \mathcal{D}(\phi(k), x),$$

and for $b = 1$:

$$\text{Enc}(\phi, x) = \$(\cdot), \text{Dec}(\phi, x) = \perp$$

where $k \leftarrow_s \mathcal{K}$ and $b \leftarrow_s \{0, 1\}$ are sampled at the start of the game.

The $\text{RKA}_{\pi, \Phi}^{\text{A}}$ game for encryption scheme π , family of functions Φ , and adversary \mathcal{A} outputs **true** if the adversary outputs the correct bit b at the end of the game.

We define the advantage of an adversary \mathcal{A} by:

$$\text{Adv}_{\pi, \Phi}^{\text{rka-ae}}(\mathcal{A}) = 2 \cdot \Pr[\text{RKA-AE}_{\pi, \Phi}^{\text{A}} \Rightarrow \text{true}] - 1.$$

Theorem 13 (UP-INT-BI Security of XOR-KEM). *Let $\pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme with keyspace $\{0, 1\}^n$ and let Π be the updatable AE scheme XOR-KEM built using π as defined above. Furthermore, let Φ be the set of permutations $\Phi := \{XOR_\Delta : \Delta \in \{0, 1\}^n\}$ where $XOR_\Delta : \{0, 1\}^n \rightarrow \{0, 1\}^n$ denotes the function $XOR_\Delta(k) = k \oplus \Delta$.*

Then for any adversary \mathcal{A} for the game UP-INT, there exists an adversary \mathcal{B} for the Φ -restricted RKA security game where:

$$\text{Adv}_{\Pi, 0, t}^{\text{up-int-bi}}(\mathcal{A}) \leq \text{Adv}_{\pi, \Phi}^{\text{rka-ae}}(\mathcal{B})$$

for all $t \geq 1$.

Proof. Let game G_0 be the original UP-INT-BI game.

Let game G_1 be identical to G_0 , except we replace all encryption queries with random strings, and condition on the event that every query (i, C) the adversary makes to the Try oracle does not result in setting the win flag to true.

We show that $|\Pr[S_0] - \Pr[S_1]| \leq \text{Adv}_{\pi, \Phi}^{\text{rka-ae}}(\mathcal{B})$. To see this, we construct an adversary \mathcal{B} which simulates the UP-INT-BI game.

For every encryption query to the underlying encryption scheme \mathcal{E} , the adversary instead makes a call to the RKA oracle $\text{Enc}(XOR_x, m)$, where x is the usual randomly sampled DEM key.

In effect, the key used for encrypting the data is $K + x$ where K is the key randomly generated in the RKA game.

When $b = 0$ in the RKA game, encryption queries are returned as $\mathcal{E}(K + x, m)$. This perfectly simulates G_0 .

When $b = 1$, decryption queries are all replaced by \perp . Therefore, every (i, C) submitted to the Try oracle was either seen before, and hence $\text{InvalidCTXT}_{\text{BI}}(i, C)$ returns true, or results in a call to $\text{Dec}(XOR_x(i), C)$ which returns \perp .

Therefore this perfectly simulates game G_1 .

Finally, $\Pr[S_1] = 0$, and thus we get the stated result. \square

We conjecture that an authenticated encryption scheme achieving this new notion could be obtained by using a RKA-secure PRF to build a CTR mode cipher, with the PRF also serving as the MAC in encrypt-then-mac mode.