Hybrid Encryption in a Multi-user Setting, Revisited

Federico Giacon[1(✉)], Eike Kiltz[1], and Bertram Poettering[2]

[1] Ruhr University Bochum, Bochum, Germany
{federico.giacon,eike.kiltz}@rub.de
[2] Royal Holloway, University of London, London, UK
bertram.poettering@rhul.ac.uk

Abstract. This paper contributes to understanding the interplay of security notions for PKE, KEMs, and DEMs, in settings with multiple users, challenges, and instances. We start analytically by first studying (a) the tightness aspects of the standard hybrid KEM+DEM encryption paradigm, (b) the inherent weak security properties of all deterministic DEMs due to generic key-collision attacks in the multi-instance setting, and (c) the negative effect of deterministic DEMs on the security of hybrid encryption.

We then switch to the constructive side by (d) introducing the concept of an *augmented data encapsulation mechanism* (ADEM) that promises robustness against multi-instance attacks, (e) proposing a variant of hybrid encryption that uses an ADEM instead of a DEM to alleviate the problems of the standard KEM+DEM composition, and (f) constructing practical ADEMs that are secure in the multi-instance setting.

Keywords: Hybrid encryption · Multi-user security · Tightness

1 Introduction

Hybrid encryption and its security. Public-key encryption (PKE) is typically implemented following a hybrid paradigm: To encrypt a message, first a randomized key encapsulation mechanism (KEM) is used to establish—independently of the message—a fresh session key that the receiver is able to recover using its secret key; then a deterministic data encapsulation mechanism (DEM) is used with the session key to encrypt the message. Both KEM and DEM output individual ciphertexts, and the overall PKE ciphertext is just their concatenation. Benefits obtained from deconstructing PKE into the two named components include easier implementation, deployment, and analysis. An independent reason that, in many cases, makes separating asymmetric from symmetric techniques actually necessary is that asymmetric cryptographic components

The full version can be found in the IACR eprint archive as article 2017/843 (https://eprint.iacr.org/2017/843).

M. Abdalla and R. Dahab (Eds.): PKC 2018, LNCS 10769, pp. 159–189, 2018.
https://doi.org/10.1007/978-3-319-76578-5_6

can typically deal only with messages of limited length (e.g., 2048 bit messages in RSA-based systems) or of specific structure (e.g., points on an elliptic curve). The paradigm of hybrid encryption, where the message-processing components are strictly separated from the asymmetric ones, side-steps these disadvantages.

Hybrid encryption was first studied on a formal basis in [11]. (Implicitly the concept emerged much earlier, for instance in PGP email encryption.) The central result on the security of this paradigm is that combining a secure KEM with a secure DEM yields a secure PKE scheme. Various configurations of sufficient definitions of 'secure' for the three components have been proposed [11,16,18], with the common property that the corresponding security reductions are tight.

Multi-user security of PKE and KEMs. Classic security definitions for PKE, like IND-CPA and IND-CCA, formalize notions of confidentiality of a single message encrypted to a single user. (For public-key primitives, we identify (receiving) users with public keys.) This does not well-reflect real-world requirements where, in principle, billions of senders might use the same encryption algorithm to send, concurrently and independently of each other, related or unrelated messages to billions of receivers. Correspondingly, for adequately capturing security aspects of PKE that is deployed at large scale, generalizations of IND-CPA/CCA have been proposed that formalize indistinguishability in the face of multiple users and multiple challenge queries [4] (the goal of the adversary is to break confidentiality of one message, not necessarily of all messages). On the one hand, fortunately, these generalized notions turn out to be equivalent to the single-user single-challenge case [4] (thus supporting the relevance of the latter). On the other hand, and unfortunately, all known proofs of this statement use reductions that are not tight, losing a factor of $n \cdot q_e$ where n is the number of users and q_e the allowed number of challenge queries per user. Of course this does not mean that PKE schemes with tightly equivalent single- and multi-user security cannot exist, and indeed [1,4,12,15,17,19,20] expose examples of schemes with tight reductions between the two worlds.

The situation for KEMs is the same as for PKE: While the standard security definitions [11,16] consider exclusively the single-user single-challenge case, natural multi-user multi-challenge variants have been considered and can be proven—up to a security loss with factor $n \cdot q_e$—equivalent to the standard notions.

Multi-instance security of DEMs. Besides scaled versions of security notions for PKE and KEMs, we also consider similarly generalized variants of DEM security. More specifically, we formalize a new[1] security notion for DEMs that assumes multiple independently generated instances and allows for one challenge encapsulation per instance. (For secret key primitives, we identify instances with secret keys.) The single-challenge restriction is due to the fact that overall we are interested in KEM+DEM composition and, akin to the single-instance

[1] We are not aware of prior work that explicitly develops multi-instance security models for DEMs; however, [22] (and others) discuss the multi-instance security of symmetric encryption, and [7] considers the multi-instance security of (nonce-based) AE.

case [11], a one-time notion for the DEM is sufficient (and, as we show, necessary) for proving security of the hybrid. As for PKE and KEMs, the multi-instance security of a DEM is closely coupled to its single-instance security; however, generically, if N is the number of instances, the corresponding reduction loses a factor of N.

A couple of works [8,22] observe that DEMs that possess a specific technical property[2] indeed have a lower security in the multi-instance setting than in the single-instance case. This is shown via attacks that assume a number of instances that is so large that, with considerable probability, different instances use the same encapsulation key; such key collisions can be detected, and message contents can be recovered. Note that, strictly speaking, the mentioned type of attack does *not* imply that the reduction of multi-instance to single-instance security is necessarily untight, as the attacks crucially depend on the DEM key size which is a parameter that does not appear in above tightness bounds. We finally point out that the attacks described in [8,22] are not general but target only specific DEMs. In this paper we show that the security of *any* (deterministic) DEM degrades as the number of considered instances increases.

1.1 Our Contributions

This paper contributes to understanding the interplay of security notions for PKE, KEMs, and DEMs, in settings with multiple users, challenges, and instances. We start analytically by first studying (a) the tightness aspects of the standard hybrid KEM+DEM encryption paradigm, (b) the inherent weak security properties of deterministic DEMs in the multi-instance setting, and (c) the negative effect of deterministic DEMs on the security of hybrid encryption. We then switch to the constructive side by (d) introducing the concept of an *augmented data encapsulation mechanism* (ADEM) that promises robustness against multi-instance attacks, (e) proposing a variant of hybrid encryption that uses an ADEM instead of a DEM to alleviate the problems of the standard KEM+DEM composition, and (f) constructing secure practical ADEMs. We proceed with discussing some of these results in more detail, in the order in which they appear in the paper.

Standard KEM+DEM Hybrid Encryption. In Sect. 3 we define syntax and security properties of PKE, KEMs, and DEMs; we also recall hybrid encryption. Besides unifying the notation of algorithms and security definitions, the main contribution of this section is to provide a new multi-instance security notion for DEMs that matches the requirements of KEM+DEM hybrid encryption in the multi-user multi-challenge setting. That is, hybrid encryption is secure, tightly, if KEM and DEM are simultaneously secure (in our sense). We further show that

[2] The cited work is not too clear about this property; loosely speaking the condition seems to be that colliding ciphertexts of the same message under random keys can be used as evidence that also the keys are colliding. One example for a DEM with this property is CBC encryption.

any attack on the multi-instance security of the DEM tightly implies an attack on the multi-user multi-challenge security of the hybrid scheme. This implication is particularly relevant in the light of the results of Sect. 4, discussed next.

Generic Key-Collision Attacks on Deterministic DEMs. In Sect. 4 we study two attacks that target arbitrary (deterministic) DEMs, leveraging on the multi-instance setting and exploiting the tightness gap between single-instance and multi-instance security. Concretely, inspired by the key-collision attacks (also known as birthday-bound attacks) from [7,8,22], in Sects. 4.1 and 4.2 we describe two attacks against arbitrary DEMs that break indistinguishability or even recover encryption keys with success probability $N^2/|\mathcal{K}|$, where N is the number of instances and $\mathcal{K}$ is the DEM's key space. (The reason for specifying two attacks instead of just one is that deciding which one is preferable may depend on the particular DEM.) As mentioned above, in hybrid encryption these attacks carry over to the overall PKE.

What are the options to thwart the described attacks on DEMs? One way to avoid key-collision attacks in practice is of course to increase the key length of the DEM. This requires the extra burden of also changing the KEM (it has to output longer keys) and hence might not be a viable option. (Observe that leaving the KEM as-is but expanding its key to, say, double length using a PRG is *not* going to work as our generic DEM attacks would immediately kick in against that construction as well.) Another way to go would be to randomize the DEM. Drawbacks of this approach are that randomness might be a scarce resource (in particular on embedded systems, but also on desktop computers there is a price to pay for requesting randomness[3]), and that randomized schemes necessarily have longer ciphertexts than deterministic ones. In Sects. 5 to 7 we explore an alternative technique to overcome key-collision attacks in hybrid encryption without requiring further randomness and without requiring changing the KEM. We describe our approach in the following.

KEM+ADEM Hybrid Encryption. In Sect. 5 we introduce the concept of an augmented data encapsulation mechanism (ADEM). It is a variant of a DEM that takes an additional input: the tag. The intuition is that ADEMs are safer to use for hybrid encryption than regular DEMs, in particular in the presence of session-key collisions: Even if two keys collide, security is preserved if the corresponding tags are different. Importantly, the two generic DEM attacks from Sect. 4 do not apply to ADEMs. In Sect. 5 we further consider *augmented hybrid encryption*, which constructs PKE from a KEM and an ADEM by using the KEM ciphertext as ADEM tag. The corresponding security reduction is tight.

[3] Obtaining entropy from a modern operating system kernel involves either file access or system calls; both options are considerably more costly than, say, doing an AES computation. While some modern CPUs have built-in randomness generators, the quality of the latter is difficult to assess and relying exclusively on them thus discouraged (see https://plus.google.com/+TheodoreTso/posts/SDcoemc9V3J).

Practical ADEM Constructions. Sections 6 and 7 are dedicated to the construction of practical ADEMs. The two constructions in Sect. 6 are based on the well-known counter mode encryption, instantiated with an ideal random function and using the tag as initial counter value. We prove tight, beyond-birthday security bounds of the form $N/|\mathcal{K}|$ for the multi-instance security of our ADEMs. That is, our constructions provably do not fall prey to key collision attacks, in particular not the ones from [8,22] and Sect. 4. Unfortunately, as they are based on counter mode, the two schemes *per se* are not secure against active adversaries. This is remedied in Sect. 7 where we show that an *augmented message authentication code*[4] (AMAC) can be used to generically strengthen a passively-secure ADEM to become secure against active adversaries. (We define AMACs and give a tightly secure construction in the same section.)

2 Notation

If S is a finite set, $s \xleftarrow{\$} S$ denotes the operation of picking an element of S uniformly at random and assigning the result to variable s. For a randomized algorithm A we write $y \xleftarrow{\$} \mathsf{A}(x_1, x_2, \ldots)$ to denote the operation of running A with inputs $x_1, x_2, \ldots$ and assigning the output to variable y. Further, we write $[\mathsf{A}(x_1, x_2, \ldots)]$ for the set of values that A outputs with positive probability. We denote the concatenation of strings with $\|$ and the XOR of same-length strings with $\oplus$. If $a \le b$ are natural numbers, we write $[a\,..\,b]$ for the range $\{a, \ldots, b\}$.

We say a sequence $v_1, \ldots, v_n$ has a (two-)collision if there are indices $1 \le i < j \le n$ such that $v_i = v_j$. More generally, the sequence has a k-collision if there exist $1 \le i_1 < \ldots < i_k \le n$ such that $v_{i_1} = \ldots = v_{i_k}$. We use predicate $\mathbf{Coll}_k[\,]$ to indicate k-collisions. For instance, $\mathbf{Coll}_2[1, 2, 3, 2]$ evaluates to *true* and $\mathbf{Coll}_3[1, 2, 3, 2]$ evaluates to *false*.

Let $\mathcal{L}$ be a finite set of cardinality $L = |\mathcal{L}|$. Sometimes we want to refer to the elements of $\mathcal{L}$ in an arbitrary but circular way, i.e., such that indices x and $x + L$ resolve to the same element. We do this by fixing an arbitrary bijection $[\![\cdot]\!]_L \colon \mathbb{Z}/L\mathbb{Z} \to \mathcal{L}$ and extending the domain of $[\![\cdot]\!]_L$ to the set $\mathbb{Z}$ in the natural way. This makes expressions like $[\![a + b]\!]_L$, for $a, b \in \mathbb{N}$, well-defined. We use the shortcut notation $[\![a \twoheadrightarrow l]\!]_L$ to refer to the span $\{[\![a + 1]\!]_L, \ldots, [\![a + l]\!]_L\}$ of length l. In particular we have $[\![a \twoheadrightarrow 1]\!]_L = \{[\![a + 1]\!]_L\}$.

Our security definitions are based on games played between a challenger and an adversary. These games are expressed using program code and terminate when the main code block executes 'return'; the argument of the latter is the output of the game. We write $\Pr[G \Rightarrow 1]$ or $\Pr[G \Rightarrow \mathit{true}]$ or just $\Pr[G]$ for the probability that game G terminates by executing a 'return' instruction with a value interpreted as true. Further, if E is some game-internal event, we write $\Pr[E]$ for the probability this event occurs. (Note the game is implicit in this notation.)

[4] The notion of an augmented MAC appeared recently in an unrelated context: An AMAC according to [3] is effectively keyed Merkle–Damgård hashing with an unkeyed output transform applied at the end. Importantly, while the notion of [3] follows the classic MAC syntax, ours does not (for having a separate tag input).

3 Traditional KEM/DEM Composition and Its Weakness

We define PKE, KEMs, and DEMs, and give security definitions that consider multi-user, multi-challenge, and multi-instance attacks. Using the techniques from [4] we show that the multi notions are equivalent to their single counterparts, up to a huge tightness loss. We show that hybrid encryption enjoys tight security also in the multi settings. We finally show how (multi-instance) attacks on the DEM can be leveraged to attacks on the PKE.

3.1 Syntax and Security of PKE, KEMs, and DEMs

Public-key encryption. A public-key encryption scheme $\mathsf{PKE} = (\mathsf{P.gen}, \mathsf{P.enc}, \mathsf{P.dec})$ is a triple of algorithms together with a message space $\mathcal{M}$ and a ciphertext space $\mathcal{C}$. The randomized key-generation algorithm $\mathsf{P.gen}$ returns a pair (pk, sk) consisting of a public key and a secret key. The randomized encryption algorithm $\mathsf{P.enc}$ takes a public key pk and a message $m \in \mathcal{M}$ to produce a ciphertext $c \in \mathcal{C}$. Finally, the deterministic decryption algorithm $\mathsf{P.dec}$ takes a secret key sk and a ciphertext $c \in \mathcal{C}$, and outputs either a message $m \in \mathcal{M}$ or the special symbol $\perp \notin \mathcal{M}$ to indicate rejection. The correctness requirement is that for all $(pk, sk) \in [\mathsf{P.gen}]$, $m \in \mathcal{M}$, and $c \in [\mathsf{P.enc}(pk, m)]$, we have $\mathsf{P.dec}(sk, c) = m$.

We adapt results from [4] to our notation, giving a game-based security definition for public-key encryption that formalizes multi-user multi-challenge indistinguishability: For a scheme PKE, to any adversary A and any number of users n we associate the distinguishing advantage $\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{PKE},\mathsf{A},n} := |\Pr[\mathsf{MUC\text{-}IND}^0_{\mathsf{A},n}] - \Pr[\mathsf{MUC\text{-}IND}^1_{\mathsf{A},n}]|$, where the two games are specified in Fig. 1. Note that if q_e resp. q_d specify upper bounds on the number of Oenc and Odec queries per user, then the single-user configurations $(n, q_e, q_d) = (1, 1, 0)$ and $(n, q_e, q_d) = (1, 1, \infty)$ correspond to standard definitions of IND-CPA and IND-CCA security for PKE.

Game $\mathsf{MUC\text{-}IND}^b_{\mathsf{A},n}$	**Oracle** $\mathsf{Oenc}(j, m_0, m_1)$	**Oracle** $\mathsf{Odec}(j, c)$
00 for all $j \in [1\,..\,n]$:	05 $c \xleftarrow{\$} \mathsf{P.enc}(pk_j, m_b)$	08 if $c \in C_j$: return $\perp$
01 $\quad(pk_j, sk_j) \xleftarrow{\$} \mathsf{P.gen}$	06 $C_j \leftarrow C_j \cup \{c\}$	09 $m \leftarrow \mathsf{P.dec}(sk_j, c)$
02 $\quad C_j \leftarrow \emptyset$	07 return c	10 return m
03 $b' \xleftarrow{\$} \mathsf{A}(pk_1, \ldots, pk_n)$		
04 return b'		

Fig. 1. PKE security games $\mathsf{MUC\text{-}IND}^b_{\mathsf{A},n}$, $b \in \{0, 1\}$, modeling multi-user multi-challenge indistinguishability for n users.

The following states that the multi-user multi-challenge notion is equivalent to the traditional single-user single-challenge case—up to a tightness loss linear in both the number of users and the number of challenges. The proof is in [4].

Lemma 1 [4]. *For any public-key encryption scheme* PKE*, any number of users* n*, and any adversary* A *that poses at most* q_e*-many* Oenc *and* q_d*-many*

Odec *queries per user, there exists an adversary* B *such that* $\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{PKE},\mathsf{A},n} \leq n \cdot q_e \cdot \mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{PKE},\mathsf{B},1}$*, where* B *poses at most one* Oenc *and* q_d*-many* Odec *queries. Further, the running time of* B *is at most that of* A *plus the time needed to perform* nq_e*-many* P.enc *operations and* nq_d*-many* P.dec *operations.*

Key encapsulation. A key-encapsulation mechanism $\mathsf{KEM} = (\mathsf{K.gen}, \mathsf{K.enc}, \mathsf{K.dec})$ for a finite session-key space $\mathcal{K}$ is a triple of algorithms together with a ciphertext space $\mathcal{C}$. The randomized key-generation algorithm K.gen returns a pair (pk, sk) consisting of a public key and a secret key. The randomized encapsulation algorithm K.enc takes a public key pk to produce a session key $K \in \mathcal{K}$ and a ciphertext $c \in \mathcal{C}$. Finally, the deterministic decapsulation algorithm K.dec takes a secret key sk and a ciphertext $c \in \mathcal{C}$, and outputs either a session key $K \in \mathcal{K}$ or the special symbol $\perp \notin \mathcal{K}$ to indicate rejection. The correctness requirement is that for all $(pk, sk) \in [\mathsf{K.gen}]$ and $(K, c) \in [\mathsf{K.enc}(pk)]$ we have $\mathsf{K.dec}(sk, c) = K$.

Like for PKE schemes we give a security definition for KEMs that formalizes multi-user multi-challenge indistinguishability: For a scheme KEM, to any adversary A and any number of users n we associate the distinguishing advantage $\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{KEM},\mathsf{A},n} := |\Pr[\mathsf{MUC\text{-}IND}^0_{\mathsf{A},n}] - \Pr[\mathsf{MUC\text{-}IND}^1_{\mathsf{A},n}]|$, where the two games are specified in Fig. 2. Note that if q_e resp. q_d specify upper bounds on the number of Oenc and Odec queries per user, then the single-user configurations $(n, q_e, q_d) = (1, 1, 0)$ and $(n, q_e, q_d) = (1, 1, \infty)$ correspond precisely to standard definitions of IND-CPA and IND-CCA security for KEMs.

Game $\mathsf{MUC\text{-}IND}^b_{\mathsf{A},n}$	**Oracle** $\mathsf{Oenc}(j)$	**Oracle** $\mathsf{Odec}(j, c)$
00 for all $j \in [1 \,..\, n]$:	05 $(K^0, c) \xleftarrow{\$} \mathsf{K.enc}(pk_j)$	09 if $c \in C_j$: return $\perp$
01 $\quad (pk_j, sk_j) \xleftarrow{\$} \mathsf{K.gen}$	06 $K^1 \xleftarrow{\$} \mathcal{K}$	10 $K \leftarrow \mathsf{K.dec}(sk_j, c)$
02 $\quad C_j \leftarrow \emptyset$	07 $C_j \leftarrow C_j \cup \{c\}$	11 return K
03 $b' \xleftarrow{\$} \mathsf{A}(pk_1, \ldots, pk_n)$	08 return (K^b, c)	
04 return b'		

Fig. 2. KEM security games $\mathsf{MUC\text{-}IND}^b_{\mathsf{A},n}$, $b \in \{0, 1\}$, modeling multi-user multi-challenge indistinguishability for n users.

Akin to the PKE case, our KEM multi-user multi-challenge notion is equivalent to its single-user single-challenge relative—again up to a tightness loss linear in the number of users and challenges. The proof can be found in the full version [14].

Lemma 2. *For any key-encapsulation mechanism* KEM*, any number of users* n*, and any adversary* A *that poses at most* q_e*-many* Oenc *and* q_d*-many* Odec *queries per user, there exists an adversary* B *such that* $\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{KEM},\mathsf{A},n} \leq n \cdot q_e \cdot \mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{KEM},\mathsf{B},1}$*, where* B *poses at most one* Oenc *and* q_d*-many* Odec *queries. Further, the running time of* B *is at most that of* A *plus the time needed to perform* nq_e*-many* K.enc *operations and* nq_d*-many* K.dec *operations.*

Data encapsulation. A data-encapsulation mechanism $\mathsf{DEM} = (\mathsf{D.enc}, \mathsf{D.dec})$ for a message space $\mathcal{M}$ is a pair of deterministic algorithms associated with a finite key space $\mathcal{K}$ and a ciphertext space $\mathcal{C}$. The encapsulation algorithm $\mathsf{D.enc}$ takes a key $K \in \mathcal{K}$ and a message $m \in \mathcal{M}$, and outputs a ciphertext $c \in \mathcal{C}$. The decapsulation algorithm $\mathsf{D.dec}$ takes a key $K \in \mathcal{K}$ and a ciphertext $c \in \mathcal{C}$, and outputs either a message $m \in \mathcal{M}$ or the special symbol $\perp \notin \mathcal{M}$ to indicate rejection. The correctness requirement is that for all $K \in \mathcal{K}$ and $m \in \mathcal{M}$ we have $\mathsf{D.dec}(K, \mathsf{D.enc}(K, m)) = m$.

As a security requirement for DEMs we formalize a multi-instance variant of the standard one-time indistinguishability notion: In our model the adversary can request one challenge encapsulation for each of a total of N independent keys; decapsulation queries are not restricted and can be asked multiple times for the same key. The corresponding games are in Fig. 3. Note that lines 05 and 09 ensure that the adversary cannot ask for decapsulations with respect to a key before having a challenge message encapsulated with it. (This matches the typical situation as it emerges in a KEM/DEM hybrid.) For a scheme DEM, to any adversary A and any number of instances N we associate the distinguishing advantage $\mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{A},N} := |\Pr[\mathsf{MIOT\text{-}IND}^0_{\mathsf{A},N}] - \Pr[\mathsf{MIOT\text{-}IND}^1_{\mathsf{A},N}]|$. Note that if Q_d specifies a global upper bound on the number of Odec queries, then the single-instance configurations $(N, Q_d) = (1, 0)$ and $(N, Q_d) = (1, \infty)$ correspond to standard definitions of OT-IND-CPA and OT-IND-CCA security for DEMs.

Game $\mathsf{MIOT\text{-}IND}^b_{\mathsf{A},N}$	**Oracle** $\mathsf{Oenc}(j, m_0, m_1)$	**Oracle** $\mathsf{Odec}(j, c)$
00 for all $j \in [1 \,..\, N]$:	05 if $C_j \neq \emptyset$: return $\perp$	09 if $C_j = \emptyset$: return $\perp$
01 $\quad K_j \xleftarrow{\$} \mathcal{K}$	06 $c \leftarrow \mathsf{D.enc}(K_j, m_b)$	10 if $c \in C_j$: return $\perp$
02 $\quad C_j \leftarrow \emptyset$	07 $C_j \leftarrow C_j \cup \{c\}$	11 $m \leftarrow \mathsf{D.dec}(K_j, c)$
03 $b' \xleftarrow{\$} \mathsf{A}$	08 return c	12 return m
04 return b'		

Fig. 3. DEM security games $\mathsf{MIOT\text{-}IND}^b_{\mathsf{A},N}$, $b \in \{0, 1\}$, modeling multi-instance one-time indistinguishability for N instances.

Similarly to the cases of PKE and KEMs, our multi-instance notion for DEMs is equivalent to its single-instance counterpart, with a tightness loss of N. The proof can be found in the full version [14].

Lemma 3. *For any data-encapsulation mechanism* DEM*, any number of instances* N*, and any adversary* A *that poses at most* Q_d*-many* Odec *queries in total, there exists an adversary* B *such that* $\mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{A},N} \leq N \cdot \mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{B},1}$*, where* B *poses at most one* Oenc *and* Q_d*-many* Odec *queries. Further, the running time of* B *is at most that of* A *plus the time needed to perform* N*-many* $\mathsf{D.enc}$ *operations and* Q_d*-many* $\mathsf{D.dec}$ *operations.*

3.2 Hybrid Encryption

The main application of KEMs and DEMs is the construction of public key encryption: To obtain a (hybrid) PKE scheme, a KEM is used to establish a

session key and a DEM is used with this key to protect the confidentiality of the message [11]. The details of this construction are in Fig. 4. It requires that the session key space of the KEM and the key space of the DEM coincide.

Proc P.gen	**Proc** P.enc(pk, m)	**Proc** P.dec$(sk, \langle c_1, c_2 \rangle)$
00 $(pk, sk) \xleftarrow{\$} \mathsf{K.gen}$	02 $(K, c_1) \xleftarrow{\$} \mathsf{K.enc}(pk)$	05 $K \leftarrow \mathsf{K.dec}(sk, c_1)$
01 return (pk, sk)	03 $c_2 \leftarrow \mathsf{D.enc}(K, m)$	06 if $K = \perp$: return $\perp$
	04 return $\langle c_1, c_2 \rangle$	07 $m \leftarrow \mathsf{D.dec}(K, c_2)$
		08 return m

Fig. 4. Hybrid construction of scheme PKE from schemes KEM and DEM. We write $\langle c_1, c_2 \rangle$ for the encoding of two ciphertext components into one.

The central composability result for hybrid encryption [11] says that if the KEM and DEM components are strong enough then also their combination is secure, with tight reduction. In Theorem 1 we give a generalized version of this claim: it considers multiple users and challenges, and implies the result from [11] as a corollary. Note that also our generalization allows for a tight reduction. The proof can be found in the full version [14].

Theorem 1. *Let* PKE *be the hybrid public-key encryption scheme constructed from a key-encapsulation mechanism* KEM *and a data-encapsulation mechanism* DEM *as in Fig. 4. Then for any number of users n and any PKE adversary* A *that poses at most q_e-many* Oenc *and q_d-many* Odec *queries per user, there exist a KEM adversary* B *and a DEM adversary* C *such that*

$$\mathbf{Adv}^{\text{muc-ind}}_{\mathsf{PKE},\mathsf{A},n} \leq 2\mathbf{Adv}^{\text{muc-ind}}_{\mathsf{KEM},\mathsf{B},n} + \mathbf{Adv}^{\text{miot-ind}}_{\mathsf{DEM},\mathsf{C},nq_e}.$$

The running time of B *is at most that of* A *plus the time required to run nq_e DEM encapsulations and nq_e DEM decapsulations. The running time of* C *is similar to the running time of* A *plus the time required to run nq_e KEM encapsulations, nq_e KEM decapsulations, and nq_e DEM decapsulations.* B *poses at most q_e-many* Oenc *and q_d-many* Odec *queries per user, and* C *poses at most nq_e-many* Oenc *and nq_d-many* Odec *queries in total.*

Theorem 1 bounds the distinguishing advantage of adversaries against hybrid PKE conditioned on its KEM and DEM components being secure. Note that from this result it cannot be deduced that deploying an insecure DEM (potentially in combination with a secure KEM) necessarily leads to insecure PKE. We show in Theorem 2 that also the latter implication holds. To ease the analysis, instead of requiring MUC-IND-like properties of the KEM, we rather assume that it has uniformly distributed session keys. Formally this means that for all public keys pk the distribution of $[(K, c) \xleftarrow{\$} \mathsf{K.enc}(pk)$; output $K]$ is identical with the uniform distribution on key space $\mathcal{K}$. The proof can be found in the full version [14].

Theorem 2. *For a key-encapsulation mechanism* KEM *and a data-encapsulation mechanism* DEM *let* PKE *be the corresponding hybrid encryption scheme. If* KEM *has uniform keys in* $\mathcal{K}$*, any attack on* DEM *can be converted to an attack on* PKE*. More precisely, for any* n, q_e *and any DEM adversary* A *that poses in total at most* nq_e*-many* Odec *queries, there exists an adversary* B *such that*

$$\mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{A},nq_e} \leq \mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{PKE},\mathsf{B},n} + \frac{{nq_e}^2}{2\,|\mathcal{K}|}.$$

The running time of B *is about that of* A*, and* B *poses at most* q_e*-many* Oenc *queries per user and* Q_d*-many* Odec *queries in total.*

4 Deterministic DEMs and Their Multi-instance Security

We give two generic key-collision attacks on the multi-instance security of (deterministic) DEMs. They have different attack goals (indistinguishability vs. key recovery) and succeed with slightly different probabilities. More precisely, in both cases the leading term of the success probability comes from the birthday bound and evaluates to roughly $N^2/|\mathcal{K}|$, and is thus much larger than the $N/|\mathcal{K}|$ that intuition might expect. By Theorem 2 the attacks can directly be lifted to ones targeting the multi-user multi-challenge security of a corresponding hybrid encryption scheme, achieving the same advantage.

4.1 A Passive Multi-instance Distinguishing Attack on DEMs

We describe an attack against multi-instance indistinguishability that applies generically to all DEMs. Notably, the attack is fully passive, i.e., the adversary does not pose any query to its Odec oracle. As technical requirements we assume a finite message space and a number of instances such that the inequalities $N^2 \leq 2\,|\mathcal{K}|$ and $|\mathcal{M}| \geq 3\,|\mathcal{K}| + N - 1$ are fulfilled. We consider these conditions extremely mild, since in practice $\mathcal{M}$ is very large and the value N can be chosen arbitrarily low by simply discarding some inputs.

For any value $N \in \mathbb{N}$ the details of our adversary $\mathsf{A} = \mathsf{A}_N$ are in Fig. 5a. It works as follows: It starts by picking uniformly at random messages $m_0, m_1^1, \ldots, m_1^N \in \mathcal{M}$ such that $m_1^1, \ldots, m_1^N$ are pairwise distinct. (Note the corresponding requirement $N \leq |\mathcal{M}|$ follows from above condition.) The adversary then asks for encapsulations of these messages in a way such that it obtains either N encapsulations of m_0 (if executed in game $\mathsf{MIOT\text{-}IND}^0$), or one encapsulation of each message m_1^j (if executed in game $\mathsf{MIOT\text{-}IND}^1$). If any two of the received ciphertexts collide, the adversary outputs 1; otherwise it outputs 0. The following theorem makes statements about advantage and running time of this adversary.

Theorem 3. *For a finite message space $\mathcal{M}$, let* DEM *be a DEM with key space $\mathcal{K}$. Suppose that $N^2 \leq 2|\mathcal{K}|$ and $|\mathcal{M}| \geq 3|\mathcal{K}| + N - 1$. Then adversary* A *from Fig. 5a breaks the N-instance indistinguishability of* DEM*, achieving the advantage*

$$\mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{A},N} \geq \frac{N(N-1)}{12\,|\mathcal{K}|}.$$

Its running time is $\mathcal{O}(N \log N)$, and it poses N-many Oenc *and no* Odec *queries.*

We remark that, more generally, the bound on $|\mathcal{M}|$ can be relaxed to $|\mathcal{M}| \geq 2\,|\mathcal{K}|\,(1+\delta) + N - 1$ for some $\delta \geq 0$ to obtain $\mathbf{Adv}^{\mathsf{miot\text{-}ind}}_{\mathsf{DEM},\mathsf{A},N} \geq \frac{\delta}{\delta+1} \cdot \frac{N(N-1)}{4|\mathcal{K}|}$.

Adversary A_N
00 $m_0 \xleftarrow{\$} \mathcal{M}$
01 for all $j \in [1\,..\,N]$:
02 $\quad m_1^j \xleftarrow{\$} \mathcal{M} \setminus \{m_1^1, \ldots, m_1^{j-1}\}$
03 $\quad c^j \leftarrow \mathsf{Oenc}(j, m_0, m_1^j)$
04 return 1 iff $\mathbf{Coll}_2[c^1, \ldots, c^N]$

(a) MIOT-IND adversary, Theorem 3.

Adversary A_N
00 for all $i \in [1\,..\,N]$:
01 $\quad K_i \xleftarrow{\$} \mathcal{K} \setminus \{K_1, \ldots, K_{i-1}\}$
02 $\quad c_i \leftarrow \mathsf{D.enc}(K_i, m_0)$
03 for all $j \in [1\,..\,N]$:
04 $\quad c'_j \leftarrow \mathsf{Oenc}(j, m_0)$
05 if $\exists (i,j) \in [1\,..\,N]^2$ s.t. $c_i = c'_j$:
06 $\quad$ return (K_i, j)
07 return $\perp$

(b) MIOT-KR adversary, Theorem 4.

Fig. 5. Adversaries against: (a) multi-instance indistinguishability and (b) multi-instance key recovery. Both ask for N encapsulations (resp. lines 03 and line 04) but do not use their decapsulation oracle.

Proof. The task of collecting N ciphertexts and checking for the occurrence of a collision can be completed in $\mathcal{O}(N \log N)$ operations. In the following we first assess the performance of the adversary when executed in games $\mathsf{MIOT\text{-}IND}^0$ and $\mathsf{MIOT\text{-}IND}^1$; then we combine the results.

Case $\mathsf{MIOT\text{-}IND}^0$. Adversary A receives N encapsulations of the same message m_0, created with N independent keys $K_1, \ldots, K_N$. If two of these keys collide then the corresponding (deterministic) encapsulations collide as well and A returns 1. Since $N(N-1) < N^2 \leq 2|\mathcal{K}|$ by the birthday bound we obtain

$$\Pr[\mathsf{MIOT\text{-}IND}^0_{\mathsf{A},N}] \geq \frac{N(N-1)}{4\,|\mathcal{K}|}.$$

Case $\mathsf{MIOT\text{-}IND}^1$. Adversary A receives encapsulations $c^1, \ldots, c^N$ of uniformly distributed (but distinct) messages $m_1^1, \ldots, m_1^N$. Denote with K_j the key used to compute c^j, let $\mathcal{M}_j \coloneqq \mathcal{M} \setminus \{m_1^1, \ldots, m_1^{j-1}\}$, and let further $\mathcal{C}_j \coloneqq \mathsf{D.enc}(K_j, \mathcal{M}_j)$ denote the image of $\mathcal{M}_j$ under (injective) function $\mathsf{D.enc}(K_j, \cdot)$. Observe this setup implies $|\mathcal{C}_j| = |\mathcal{M}_j|$ and $|\mathcal{C}_1| > \ldots > |\mathcal{C}_N|$. If further follows that each ciphertext c^j is uniformly distributed in set $\mathcal{C}_j$.

We aim at establishing an upper-bound on the collision probability of ciphertexts $c^1, \ldots, c^N$. The maximum collision probability is attained in the worst-case $\mathcal{C}_1 \supset \ldots \supset \mathcal{C}_N$, in which it is bounded by the collision probability of choosing N values uniformly from a set of cardinality $|\mathcal{C}_N| = |\mathcal{M}| - N + 1$. Using again the birthday bound and $|\mathcal{M}| \geq 3\,|\mathcal{K}| + N - 1$ we obtain

$$\Pr[\mathsf{MIOT\text{-}IND}^1_{\mathsf{A},N}] \leq \frac{1}{2} \cdot \frac{N(N-1)}{|\mathcal{M}| - N + 1} \leq \frac{N(N-1)}{6\,|\mathcal{K}|}.$$

Combining the two bounds yields the equation in our statement.

4.2 A Passive Multi-instance Key-Recovery Attack on DEMs

We give a generic attack on DEMs that aims at recovering keys rather than distinguishing encapsulations. Like in Sect. 4.1 the attack is passive. It is inspired by work of Zaverucha [22] and Chatterjee et al. [8]. However, our results are more general than theirs for not restricted to one specific DEM.

To formalize the notion of resilience against key recovery we correspondingly adapt the MIOT-IND game from Fig. 3 and obtain the MIOT-KR game specified in Fig. 6. The N-instance advantage of an adversary A is then defined as $\mathbf{Adv}^{\mathsf{miot\text{-}kr}}_{\mathsf{DEM},\mathsf{A},N} := \Pr[\mathsf{MIOT\text{-}KR}_{\mathsf{A},N}]$. The following theorem shows that for virtually all practical DEMs (including those based on CBC mode, CTR mode, OCB, etc., and even one-time pad encryption) there exist adversaries achieving a considerable key recovery advantage, conditioned on the DEM key space being small enough. Concretely, the adversaries we propose encapsulate $2N$ times the same message (N times with random but known keys, and N times with random but unknown keys) and detect collisions of ciphertexts.[5] As any ciphertext collision stems (in practice) from a collision of keys, this method allows for key recovery.[6]

Theorem 4. *Fix a DEM and denote its key space with $\mathcal{K}$ and its message space with $\mathcal{M}$. Let $m_0 \in \mathcal{M}$ be any fixed message. Fixing $N \in \mathbb{N}$ as a parameter, consider the adversary $\mathsf{A} = \mathsf{A}_N$ specified in Fig. 5b. We then have*

$$\mathbf{Adv}^{\mathsf{miot\text{-}kr}}_{\mathsf{DEM},\mathsf{A},N} \geq p(m_0) \cdot \min\left\{\frac{1}{2}, \frac{N^2}{2\,|\mathcal{K}|}\right\},$$

where $p(m_0)$ denotes the collision probability

$$p(m_0) := \Pr_{K_1, K_2 \xleftarrow{\$} \mathcal{K}}[K_1 = K_2 \mid \mathsf{D.enc}(K_1, m_0) = \mathsf{D.enc}(K_2, m_0)].$$

[5] While our setup is formally meaningful, in practice it would correspond to N parties, for a huge number N, encapsulating the same message m_0. This might feel rather unrealistic. However, we argue that a close variant of the attack might very well have the potential for practicality: All widely deployed DEMs are *online*, i.e., compute ciphertexts 'left-to-right'. For such DEMs, for our attack to be successful, it suffices that the N parties encapsulate (different) messages that have a common prefix, for instance a standard protocol header.

[6] The efficiency of this attack can likely be improved, on a heuristic basis, by deploying dedicated data structures like rainbow tables.

Game MIOT-KR$_{\mathsf{A},N}$	**Oracle** Oenc(j, m)	**Oracle** Odec(j, c)
00 for all $j \in [1 \,..\, N]$:	05 if $C_j \neq \emptyset$: return $\perp$	09 if $C_j = \emptyset$: return $\perp$
01 $\quad K_j \xleftarrow{\$} \mathcal{K}$	06 $c \leftarrow \mathsf{D.enc}(K_j, m)$	10 if $c \in C_j$: return $\perp$
02 $\quad C_j \leftarrow \emptyset$	07 $C_j \leftarrow C_j \cup \{c\}$	11 $m \leftarrow \mathsf{D.dec}(K_j, c)$
03 $(K, i) \xleftarrow{\$} \mathsf{A}$	08 return c	12 return m
04 return 1 iff $K = K_i$		

Fig. 6. DEM security game MIOT-KR$_{\mathsf{A},N}$ modeling resilience against key recovery, for N instances.

Its running time is $\mathcal{O}(N \log N)$, and it poses N-many Oenc *and no* Odec *queries.*

We further prove that in the case of DEMs based on one-time pad encryption we have $p(m_0) = 1$ for any m_0. Further, in the case of CBC-based encapsulation there exists a message m_0 such that $p(m_0) = |\mathcal{B}| / (|\mathcal{B}| + |\mathcal{K}| - 1)$, where $\mathcal{B}$ is the block space of the blockcipher and the latter is modeled as an ideal cipher.

Note that the performance of our attack crucially depends on the choice of message m_0, and that there does not seem to be a general technique for identifying good candidates. In particular, (artificial) DEMs can be constructed where $p(m_0)$ is small for some m_0 but large for others, or where $p(m_0)$ is small even for very long messages m_0. However, in many practical schemes the choice of m_0 is not determinant. After the proof we consider two concrete examples.

Proof. The running time of A is upper bounded by the search for collisions in line 05, since all other operations require at most linear time in N. We estimate the time bound: The list $c_1, \ldots, c_N$ is sorted, requiring time $\mathcal{O}(N \log N)$. Searching an element in the ordered list requires $\mathcal{O}(\log N)$ time. Repeating for all N searches requires $\mathcal{O}(N \log N)$. Combining these observations yields our statement.

We claim that the probability that the adversary does not output $\perp$ (in symbols, $\mathsf{A}_N \not\Rightarrow \perp$) is lower bounded by:

$$\Pr[\mathsf{A}_N \not\Rightarrow \perp] \geq 1 - \left(1 - \frac{N}{|\mathcal{K}|}\right)^N . \tag{1}$$

Since the DEM is deterministic, the probability to find any collision in line 05 is larger than the probability that any of the distinct N keys generated in lines 00–02 collides with one of the N keys $\tilde{K}_1, \ldots, \tilde{K}_N$ used by the MIOT-KR game to encapsulate. We compute the latter probability. Let $K \in \{\tilde{K}_1, \ldots, \tilde{K}_N\}$. We know that K is uniform in $\mathcal{K}$. Since $K_1, \ldots, K_N$ are distinct and independently chosen we can write: $\Pr[K \in \{K_1, \ldots, K_N\}] = N/|\mathcal{K}|$. Moreover, since the keys $\tilde{K}_1, \ldots, \tilde{K}_N$ are generated independently of each other, Eq. (1) follows.

Let now (i, j) be the indices for which the condition in line 05 is triggered, i.e., $c_i = c'_j$ and A_N outputs K_i. We can write:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{miot\text{-}kr}}_{\mathsf{DEM},\mathsf{A},N} &= \Pr[\mathsf{A}_N \not\Rightarrow \perp] \cdot \Pr[K_i = \tilde{K}_j \mid \mathsf{A}_N \not\Rightarrow \perp] \\ &\geq \left(1 - \left(1 - \frac{N}{|\mathcal{K}|}\right)^N\right) \cdot p(m_0).\end{aligned}$$

Applying known inequalities to the previous formula we obtain:

$$\mathbf{Adv}^{\mathsf{miot\text{-}kr}}_{\mathsf{DEM},\mathsf{A},N} \geq p(m_0) \cdot \left(1 - \left(1 - \frac{N}{|\mathcal{K}|}\right)^N\right) \geq p(m_0) \cdot \min\left\{\frac{1}{2}, \frac{N^2}{2\,|\mathcal{K}|}\right\}.$$

We compute $p(m_0)$ for two specific DEMs (one-time pad and CBC mode) and choices of m_0. We formalize the argument for CBC by considering single-block messages. We note that one can apply the same argument to other modes of operation, e.g., CTR. For notational simplicity we omit the description of the probability space, that is, uniform choice of $K_1, K_2 \in \mathcal{K}$.

One-time pad. The one-time pad DEM encapsulation is given by combining a key $K \in \mathcal{K} = \{0,1\}^k$ with a message $m \in \mathcal{M} = \{0,1\}^k$ using the XOR operation. In this case, if two ciphertexts for the same message collide, the same key must have been used to encapsulate. Thus $p(m_0) = 1$ for all m_0.

CBC with an ideal cipher. CBC-based DEM encapsulation consists of encrypting the message using a blockcipher in CBC mode with the zero initialization vector (IV). In the following analysis we assume an idealized blockcipher (ideal cipher model) represented by E. Note that since the IV is zero, encapsulating a single-block message m_0 under the key K is equivalent to enciphering m_0 with E_K. Let $\mathcal{B}$ be the block space. First we observe that for any single-block message m_0 we have

$$\begin{aligned}&\Pr[\mathsf{E}_{K_1}(m_0) = \mathsf{E}_{K_2}(m_0)] \\ &= \Pr[K_1 = K_2] + \Pr[K_1 \neq K_2] \Pr[\mathsf{E}_{K_1}(m_0) = \mathsf{E}_{K_2}(m_0) \mid K_1 \neq K_2] \\ &= |\mathcal{K}|^{-1} + (1 - |\mathcal{K}|^{-1})\,|\mathcal{B}|^{-1}.\end{aligned}$$

We then use the previous equality to compute $p(m_0)$ from its definition:

$$\begin{aligned}p(m_0) &= \frac{\Pr[K_1 = K_2]}{\Pr[\mathsf{E}_{K_1}(m_0) = \mathsf{E}_{K_2}(m_0)]} \\ &= \frac{|\mathcal{K}|^{-1}}{|\mathcal{K}|^{-1} + (1 - |\mathcal{K}|^{-1})\,|\mathcal{B}|^{-1}} = \frac{|\mathcal{B}|}{|\mathcal{B}| + |\mathcal{K}| - 1}.\end{aligned}$$

As an example, if $|\mathcal{B}| \geq |\mathcal{K}|$ then $p(m_0) > 1/2$ for any single-block message m_0.

5 Augmented Data Encapsulation

In the previous sections we showed that all deterministic DEMs, including those that are widely used in practice, might be less secure than expected in the face of

multi-instance attacks. We further showed that, in the setting of hybrid encryption, attacks on DEMs can be leveraged to attacks on the overall PKE. Given that the KEM+DEM paradigm is so important in practice, we next address the question of how this situation can be remedied. One option would of course be to increase the DEM key size (recall that good success probabilities in Theorems 3 and 4 are achieved only for not too large key spaces); however, increasing key sizes might not be a viable option in practical systems. (Potential reasons for this include that blockciphers like AES are slower with long keys than with short keys, and that ciphers like 3DES do not support key lengths that have a comfortable 'multi-instance security margin' in the first place.) A second option would be to augment the input given to the DEM encapsulation routine by an additional value. This idea was already considered in [22, p. 16] where, with the intuition of increasing the 'entropy' available to the DEM, it was proposed to use a KEM ciphertext as an initialization vector (IV) of a symmetric encryption mode. However, [22] does not contain any formalization or security analysis of this idea, and so it cannot be taken as granted that this strategy actually works. (And indeed, we show in Sect. 6.3 that deriving the starting value of blockcipher-based counter mode encryption from a KEM ciphertext is not ameliorating the situation for attacks based on indistinguishability.)

We formally explore the additional-input proposal for the DEM in this section. More precisely, we study two approaches of defining an *augmented data encapsulation mechanism* (ADEM), where we call the additional input the *tag*. The syntax is the same in both cases, but the security properties differ: either (a) the DEM encapsulator receives as the tag an auxiliary random (but public) string, or (b) the encapsulator receives as additional input a nonce (a 'number used once'). In both cases the decapsulation oracle operates with respect to the tag also used for encapsulation. After formalizing this we prove the following results: First, if the tag space is large enough, ADEMs that expect a nonce can safely replace ADEMs that expect a uniform tag. Second, ADEMs that expect a uniform tag can be constructed from ADEMs that expect a nonce by applying a random oracle to the latter. Our third result is that the augmented variant of hybrid encryption remains (tightly) secure.

Augmented data encapsulation. An augmented data encapsulation mechanism $\mathsf{ADEM} = (\mathsf{A.enc}, \mathsf{A.dec})$ for a message space $\mathcal{M}$ is a pair of deterministic algorithms associated with a finite key space $\mathcal{K}$, a tag space $\mathcal{T}$, and a ciphertext space $\mathcal{C}$. The encapsulation algorithm $\mathsf{A.enc}$ takes a key $K \in \mathcal{K}$, a tag $t \in \mathcal{T}$, and a message $m \in \mathcal{M}$, and outputs a ciphertext $c \in \mathcal{C}$. The decapsulation algorithm $\mathsf{A.dec}$ takes a key $K \in \mathcal{K}$, a tag $t \in \mathcal{T}$, and a ciphertext $c \in \mathcal{C}$, and outputs either a message $m \in \mathcal{M}$ or the special symbol $\perp \notin \mathcal{M}$ to indicate rejection. The correctness requirement is that for all $K \in \mathcal{K}$ and $t \in \mathcal{T}$ and $m \in \mathcal{M}$ we have $\mathsf{A.dec}(K, t, \mathsf{A.enc}(K, t, m)) = m$.

Augmented data encapsulation with uniform tags. The first security notion we formalize assumes that each encapsulation operation uses a fresh and uniformly picked tag (note this imposes the technical requirement that the tag space be finite). More precisely, while the tag may become public

after the encapsulation operation has completed, it may not be disclosed to the adversary before fixing the message to be encapsulated. We formalize this notion of uniform-tag multi-instance one-time indistinguishability for ADEMs via the games specified in Fig. 7. For a scheme ADEM, to any adversary A and any number of instances N we associate the distinguishing advantage $\mathbf{Adv}^{\text{u-miot-ind}}_{\mathsf{ADEM},\mathsf{A},N} := |\Pr[\text{U-MIOT-IND}^0_{\mathsf{A},N}] - \Pr[\text{U-MIOT-IND}^1_{\mathsf{A},N}]|$.

Game $\text{U-MIOT-IND}^b_{\mathsf{A},N}$	**Oracle** $\mathsf{Oenc}(j, m_0, m_1)$	**Oracle** $\mathsf{Odec}(j, c)$
00 for all $j \in [1 \,..\, N]$:	05 if $C_j \neq \emptyset$: return $\perp$	09 if $C_j = \emptyset$: return $\perp$
01 $\quad (K_j, t_j) \xleftarrow{\$} \mathcal{K} \times \mathcal{T}$	06 $c \leftarrow \mathsf{A.enc}(K_j, t_j, m_b)$	10 if $c \in C_j$: return $\perp$
02 $\quad C_j \leftarrow \emptyset$	07 $C_j \leftarrow C_j \cup \{c\}$	11 $m \leftarrow \mathsf{A.dec}(K_j, t_j, c)$
03 $b' \xleftarrow{\$} \mathsf{A}$	08 return (t_j, c)	12 return m
04 return b'		

Fig. 7. ADEM security games $\text{U-MIOT-IND}^b_{\mathsf{A},N}$, $b \in \{0, 1\}$, for N instances. The tags in line 11 are the same as the ones in line 06.

Augmented data encapsulation with nonces. Our second security notion for ADEMs requires the tag provided to each encapsulation operation to be unique (across all instances). The tag can be generated using any possible method (e.g., using some global type of counter). We formalize the corresponding security notion of nonce-based multi-instance one-time indistinguishability for ADEMs via the games specified in Fig. 8. For a scheme ADEM, to any adversary A and any number of instances N we associate the distinguishing advantage $\mathbf{Adv}^{\text{n-miot-ind}}_{\mathsf{ADEM},\mathsf{A},N} := |\Pr[\text{N-MIOT-IND}^0_{\mathsf{A},N}] - \Pr[\text{N-MIOT-IND}^1_{\mathsf{A},N}]|$.

Game $\text{N-MIOT-IND}^b_{\mathsf{A},N}$	**Oracle** $\mathsf{Oenc}(j, t, m_0, m_1)$	**Oracle** $\mathsf{Odec}(j, c)$
00 $T \leftarrow \emptyset$	06 if $C_j \neq \emptyset$: return $\perp$	12 if $C_j = \emptyset$: return $\perp$
01 for all $j \in [1 \,..\, N]$:	07 if $t \in T$: return $\perp$	13 if $c \in C_j$: return $\perp$
02 $\quad K_j \xleftarrow{\$} \mathcal{K}$	08 $T \leftarrow T \cup \{t\}$; $t_j \leftarrow t$	14 $m \leftarrow \mathsf{A.dec}(K_j, t_j, c)$
03 $\quad C_j \leftarrow \emptyset$	09 $c \leftarrow \mathsf{A.enc}(K_j, t_j, m_b)$	15 return m
04 $b' \xleftarrow{\$} \mathsf{A}$	10 $C_j \leftarrow C_j \cup \{c\}$	
05 return b'	11 return c	

Fig. 8. ADEM security games $\text{N-MIOT-IND}^b_{\mathsf{A},N}$, $b \in \{0, 1\}$, for N instances. The tags in line 14 are the same as the ones in line 09.

5.1 Relations Between ADEMs with Uniform and Nonce Tags

The two types of ADEMs we consider here can be constructed from each other. More concretely, the following lemma shows that if the tag space is large enough, ADEMs that expect a nonce can safely replace ADEMs that expect a uniform tag. The proof can be found in the full version [14].

Lemma 4. *Let* ADEM *be an augmented data encapsulation mechanism. If the cardinality of its tag space* $\mathcal{T}$ *is large enough and* ADEM *is secure with non-repeating tags, then it is also secure with random tags. More precisely, for any number of instances* N *and any adversary* A *there exist an adversary* B *that makes the same amount of queries such that* $\mathbf{Adv}^{\text{u-miot-ind}}_{\text{ADEM,A},N} \leq \mathbf{Adv}^{\text{n-miot-ind}}_{\text{ADEM,B},N} + N^2/(2\,|\mathcal{T}|)$*. The running time of the two adversaries is similar.*

The following simple lemma shows that ADEMs that expect a nonce can be constructed from ADEMs that expect a uniform tag by using each nonce to obtain a uniform, independent value from a random oracle. The proof is immediate since all queries to the random oracle have different input, thus the corresponding output is uniformly random and independently generated.

Lemma 5. *Let* $\mathsf{ADEM} = (\mathsf{A.enc}, \mathsf{A.dec})$ *be an augmented data encapsulation mechanism with tag space* $\mathcal{T}$*. Let* $H\colon \mathcal{T}' \to \mathcal{T}$ *denote a hash function, where* $\mathcal{T}'$ *is another tag space. Define* $\mathsf{ADEM}' = (\mathsf{A.enc}', \mathsf{A.dec}')$ *such that* $\mathsf{A.enc}'(K,t,m) := \mathsf{A.enc}(K, H(t), m)$ *and* $\mathsf{A.dec}'(K,t,c) := \mathsf{A.dec}(K, H(t), c)$*. Then if* H *is modeled as a random oracle and if* ADEM *is secure with random tags in* $\mathcal{T}$*, then* ADEM′ *is secure with non-repeating tags in* $\mathcal{T}'$*. Formally, for any number of instances* N *and any adversary* A *there exists an adversary* B *with* $\mathbf{Adv}^{\text{u-miot-ind}}_{\text{ADEM,A},N} = \mathbf{Adv}^{\text{n-miot-ind}}_{\text{ADEM}',\text{B},N}$*.*

5.2 Augmented Hybrid Encryption

A KEM and an ADEM can be combined to obtain a PKE scheme: the KEM establishes a session key and a first ciphertext component, and the ADEM is used on input the session key and the first ciphertext component (as tag) to protect the confidentiality of the message, creating a second ciphertext component. Figure 9 details this *augmented hybrid encryption*. It requires that the session key space of the KEM and the key space of the ADEM coincide. Further, the ciphertext space of the KEM needs to be a subset of the tag space of the ADEM.

Proc P.gen	**Proc** P.enc(pk, m)	**Proc** P.dec$(sk, \langle c_1, c_2\rangle)$
00 $(pk, sk) \xleftarrow{\$} \mathsf{K.gen}$	02 $(K, c_1) \xleftarrow{\$} \mathsf{K.enc}(pk)$	05 $K \leftarrow \mathsf{K.dec}(sk, c_1)$
01 return (pk, sk)	03 $c_2 \leftarrow \mathsf{A.enc}(K, c_1, m)$	06 if $K = \bot$: return $\bot$
	04 return $\langle c_1, c_2\rangle$	07 $m \leftarrow \mathsf{A.dec}(K, c_1, c_2)$
		08 return m

Fig. 9. Augmented hybrid construction of scheme PKE from schemes KEM and ADEM. We write $\langle c_1, c_2\rangle$ for the encoding of two ciphertext components into one.

The claim is that augmented hybrid encryption is more robust against attacks involving multiple users and challenges than standard hybrid encryption (see Fig. 4). The security condition posed on the ADEM requires that it be secure when operated with nonces, and the security property posed on the KEM

requires that it be both indistinguishable and have non-repeating ciphertexts (i.e., invoking the encapsulation twice on any public keys does virtually never result in colliding ciphertexts). Technically, the latter property is implied by indistinguishability. However, to obtain better bounds, we formalize it as a statistical condition: To any scheme KEM we assign the maximum ciphertext-collision probability

$$p := \max_{pk_1, pk_2} \Pr[(K_1, c_1) \xleftarrow{\$} \mathsf{K.enc}(pk_1); (K_2, c_2) \xleftarrow{\$} \mathsf{K.enc}(pk_2) : c_1 = c_2],$$

where the maximum is over all pairs pk_1, pk_2 of (potentially coinciding) public keys. Note that practical KEMs (ElGamal, RSA-based, Cramer–Shoup, ...) have much larger ciphertexts than session keys[7], so that the ciphertext-collision probability will always be negligible in practice. We proceed with a security claim for augmented hybrid encryption. The proof can be found in the full version [14].

Lemma 6. *Let* PKE *be the hybrid public-key encryption scheme constructed from a key-encapsulation mechanism* KEM *and an augmented data-encapsulation mechanism* ADEM *as in Fig. 9. Let* p *be the maximum ciphertext-collision probability of* KEM *over all possible public keys. Then for any* n *and any PKE adversary* A *that poses at most* q_e*-many* Oenc *and* q_d*-many* Odec *queries per user, there exist a KEM adversary* B *and an ADEM adversary* C *such that*

$$\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{PKE},\mathsf{A},n} \leq 2\mathbf{Adv}^{\mathsf{muc\text{-}ind}}_{\mathsf{KEM},\mathsf{B},n} + \mathbf{Adv}^{\mathsf{n\text{-}miot\text{-}ind}}_{\mathsf{ADEM},\mathsf{C},N} + 2\binom{N}{2}p,$$

where $N = nq_e$*. The running time of* B *is at most that of* A *plus the time required to run* nq_e *ADEM encapsulations and* nq_e *ADEM decapsulations. The running time of* C *is similar to that of* A *plus the time required to run* nq_e *KEM encapsulations,* nq_e *KEM decapsulations, and* nq_e *ADEM decapsulations.* B *poses at most* q_e*-many* Oenc *and* q_d*-many* Odec *queries per user, and* C *poses at most* nq_e*-many* Oenc *and* nq_d*-many* Odec *queries in total.*

6 Constructions of Augmented Data Encapsulation

We construct two augmented data-encapsulation mechanisms and analyze their security. The schemes are based on operating a function in counter mode. If the function is instantiated with an ideal random function then the ADEMs are secure beyond the birthday bound. (We also show that if the function is instead instantiated with an idealized blockcipher, i.e., a random permutation, the schemes' security may degrade.) Practical candidates for instantiating the ideal random function are for instance the compression functions of standardized Merkle–Damgård hash functions, e.g., of SHA2.[8,9] Another possibility is deriving the random function from an ideal cipher as in [21].

[7] This is no coincidence but caused by generic attacks against cyclic groups, RSA, etc.

[8] These compression functions are regularly modeled as having random behavior [2,13].

[9] The idea to construct a DEM from a hash function's compression function already appeared in the OMD schemes from [9].

6.1 Counter-Mode Encryption

Many practical DEMs are based on operating a blockcipher E in counter mode (CTR). Here, in brief, the encapsulation key is used as the blockcipher key, a sequence of message-independent input blocks is enciphered under that key, and the output blocks are XOR-ed into the message. More concretely, if under some key K a message m shall be encapsulated that, without requiring padding, evenly splits into blocks $v_1 \| \ldots \| v_l$, then the DEM ciphertext is the concatenation $w_1 \| \ldots \| w_l$ where $w_i = v_i \oplus E_K(i)$.

In the context of this paper, three properties of this construction are worth pointing out: (a) the 'counting' component of CTR mode serves a single purpose: preventing that two inputs to the blockcipher coincide; (b) any 'starting value' for the counter can be used; (c) security analyses of CTR mode typically model E as a pseudorandom function (as opposed to a pseudorandom permutation)[10].

In Fig. 10 we detail three ways of turning the principles of CTR mode into a DEM encapsulation routine. In all cases the underlying primitive is, syntactically, a function $F\colon \mathcal{K} \times \mathcal{B} \to \mathcal{D}$ that takes a key $K \in \mathcal{K}$ and maps some finite input space $\mathcal{B}$ into some finite group $(\mathcal{D}, \oplus)$. (Intuitively, $\mathcal{B}$ serves as a space of input blocks derived from a counter, and $\mathcal{D}$ as a space of pads that can be XORed into message blocks; note that if F is instantiated with a blockcipher we have $\mathcal{B} = \mathcal{D}$, but we explicitly allow other instantiations.) The most basic encapsulation routine based on CTR mode that we consider, and the one closest to our sketch above, is CTR0enc. Note that this DEM further assumes a bijection $[\![\cdot]\!]_L\colon \mathbb{Z}/L\mathbb{Z} \to \mathcal{L}$ with $\mathcal{L} = \mathcal{B}$. (Intuitively, this bijection turns a counter that is cyclic with period length L into input blocks for F; see Sect. 2 for notation.) We finally point out that all three variants of CTR mode that we formalize exclusively work with fixed-length multi-block messages (i.e., $\mathcal{M} = \mathcal{D}^l$). This choice, that we made for simplicity of exposition, is not really a restriction as 'any-length' CTR mode encryption can be simulated from 'block-wise' CTR mode encryption.

Proc CTR0enc(K, m)	**Proc** CTR+enc(K, t, m)	**Proc** CTR‖enc(K, t, m)
00 $(v_1, \ldots, v_l) \leftarrow m$	05 $(v_1, \ldots, v_l) \leftarrow m$	10 $(v_1, \ldots, v_l) \leftarrow m$
01 for all $i \in [1\,..\,l]$:	06 for all $i \in [1\,..\,l]$:	11 for all $i \in [1\,..\,l]$:
02 $\quad w_i \leftarrow v_i \oplus F(K, [\![i]\!]_L)$	07 $\quad w_i \leftarrow v_i \oplus F(K, [\![t+i]\!]_L)$	12 $\quad w_i \leftarrow v_i \oplus F(K, t \| [\![i]\!]_L)$
03 $c \leftarrow (w_1, \ldots, w_l)$	08 $c \leftarrow (w_1, \ldots, w_l)$	13 $c \leftarrow (w_1, \ldots, w_l)$
04 return c	09 return c	14 return c

Fig. 10. Encapsulation algorithms of the CTR0 DEM, the CTR+ ADEM, and the CTR‖ ADEM, for multi-block messages. In CTR0enc and CTR+enc we assume $[\![\cdot]\!]_L\colon \mathbb{Z}/L\mathbb{Z} \to \mathcal{L}$ with $\mathcal{L} = \mathcal{B}$, and in CTR‖enc we assume $[\![\cdot]\!]_L\colon \mathbb{Z}/L\mathbb{Z} \to \mathcal{L}$ and $\mathcal{T}$ such that $\mathcal{B} = \mathcal{T} \times \mathcal{L}$. The corresponding decapsulation routines is immediate.

[10] Technically, the PRP/PRF switching lemma [5] measures the price one has to pay for pursuing this modeling approach.

The two remaining procedures in Fig. 10 are ADEM encapsulation routines. The first one, CTR+enc, is the natural variant of CTR0enc where the tag space is $\mathcal{T} = [1 \,..\, L]$ and the tag specifies the starting value of the counter. The second, CTR‖enc, concatenates tag and counter. Here, the tag space $\mathcal{T}$ and parameter space $\mathcal{L}$ have to be arranged such that $\mathcal{B} = \mathcal{T} \times \mathcal{L}$.

We analyze the security of CTR+ and CTR‖ in the upcoming sections. Scheme CTR0 is not an ADEM and falls prey to our earlier attacks.

6.2 Security of Function-Based Counter Mode

We establish upper bounds on the advantage of U-MIOT-IND adversaries against the CTR+ and CTR‖ ADEMs.

Counter Mode with Tag-Controlled Starting Value. We limit the maximum amount of blocks in an encapsulation query to a fixed value ℓ. Prerequisites to our statement on CTR+ are two conditions on the number of instances relative to $\mathcal{K}$ and $\mathcal{T} = [1 \,..\, L]$. The bound is namely $N \leq \min\left\{ |\mathcal{K}|^{1/2}, (|\mathcal{T}|/(2\ell))^{1/(1+\delta)} \right\}$, for some arbitrary constant δ such that $1/N \leq \delta \leq 1$. Despite this restriction we consider our statement to be reflecting real-world applications: As an extreme example we see that the values $|\mathcal{K}| = |\mathcal{T}| = 2^{128}$, $N = 2^{56}$, $\ell = 2^{56}$, $q = 2^{64}$ and $\delta = 2/7$ fit above condition, yielding a maximum advantage of around 2^{-61}.

Theorem 5. *Suppose* $N \leq \min\left\{ |\mathcal{K}|^{1/2}, (|\mathcal{T}|/(2\ell))^{1/(1+\delta)} \right\}$, *for some* $1/N \leq \delta \leq 1$, *and suppose that* F *is modeled as a random oracle (using oracle* F*). Then for any adversary* A *against* N*-instance uniform-tag indistinguishability of* CTR+ *that poses at most* q *queries to* F*, no decapsulation queries, and encapsulates messages of length at most* ℓ *blocks we have:*

$$\mathbf{Adv}^{\text{u-miot-ind}}_{\text{CTR+},\mathsf{A},N} \leq \frac{1}{3}\frac{N}{|\mathcal{K}|} + \frac{4\ell - 2}{|\mathcal{T}|} + \frac{2q}{|\mathcal{K}|}\left(1 + \frac{1}{\delta}\right).$$

The core of the proof exploits that the outputs of (random oracle) F that are used to encapsulate are uniformly distributed in $\mathcal{D}$ and independent of each other. This requires forcing the inputs to be distinct in $\mathcal{L}$. We give further insight on some non-standard techniques the we use in the analysis in the proof.

Proof (of Theorem 5). The definition of the games $\mathsf{G}^{0,b}_{\mathsf{A},N}$, $\mathsf{G}^{1,b}_{\mathsf{A},N}$, $\mathsf{G}^{2,b}_{\mathsf{A},N}$ and $\mathsf{G}^{3,b}_{\mathsf{A},N}$ are found in Fig. 11. Except for some bookkeeping, game $\mathsf{G}^{0,b}_{\mathsf{A},N}$ is equivalent to game $\mathsf{U\text{-}MIOT\text{-}IND}^{b}_{\mathsf{A},N}$, where $b \in \{0, 1\}$. For $j \in [1 \,..\, N]$ we define $T_j = [\![t_j \twoheadrightarrow \ell]\!]_L$.

Game G^1**.** In game $\mathsf{G}^{1,b}_{\mathsf{A},N}$ we implicitly generate pairs of colliding keys. We loop over all pairs (j_1, j_2) such that $1 \leq j_1 < j_2 \leq N$. If both indices were not previously paired ($\mathsf{matched}[j_1] = \mathsf{matched}[j_2] = \mathit{false}$) and the corresponding keys collide ($K_{j_1} = K_{j_2}$) then the two indices are marked as paired. Moreover,

if the corresponding tag ranges collide ($T_{j_1} \cap T_{j_2} \neq \emptyset$) the flag bad_1 in line 10 is raised and the game aborts. We claim that

$$|\Pr[\mathsf{G}^{0,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{1,b}_{\mathsf{A},N}]| \leq \Pr[\mathrm{bad}_1] \leq \frac{2\ell - 1}{|\mathcal{T}|}. \tag{2}$$

To prove (2), we want to compute the probability $\Pr[\mathrm{bad}_1]$. Let m_{pairs} be the number of colliding key pairs in game $\mathsf{G}^{1,b}_{\mathsf{A},N}$, i.e., $2m_{\mathrm{pairs}}$ entries of flag matched are set to 1 at the end of the game. Then, for every $0 \leq i \leq \lfloor N/2 \rfloor$, $\Pr[\mathrm{bad}_1 \mid m_{\mathrm{pairs}} = i] \leq (2\ell - 1)i/|\mathcal{T}|$. This follows from the independent choices of the values K_j, t_j for each instance $j \in [1\,..\,N]$, and because for each pair of indices $j_1, j_2 \in [1\,..\,N], j_1 \neq j_2$, and for any choice of t_{j_1} there are exactly $2\ell - 1$ possible values of t_{j_2} such that $T_{j_1} \cap T_{j_2} \neq \emptyset$. The sets $\{m_{\mathrm{pairs}} = i\}$, $i \in 0, \ldots, \lfloor N/2 \rfloor$, partition the probability space, thus:

$$\begin{aligned}\Pr[\mathrm{bad}_1] &= \sum_{i=0}^{\lfloor N/2 \rfloor} \Pr[\mathrm{bad}_1 \mid m_{\mathrm{pairs}} = i] \Pr[m_{\mathrm{pairs}} = i] \\ &\leq \frac{2\ell - 1}{|\mathcal{T}|} \sum_{i=0}^{\lfloor N/2 \rfloor} i \Pr[m_{\mathrm{pairs}} = i] = \frac{2\ell - 1}{|\mathcal{T}|} \sum_{i=1}^{\lfloor N/2 \rfloor} \Pr[m_{\mathrm{pairs}} \geq i]. \end{aligned} \tag{3}$$

The last equality follows since the expected value of any random variable m with values in $\mathbb{N}$ can be written as $\sum_{i=0}^{\infty} i \Pr[m = i] = \sum_{i=1}^{\infty} \Pr[m \geq i]$. We show by induction that the terms of the sum are:

$$p_i := \Pr[m_{\mathrm{pairs}} \geq i] \leq \left(\frac{N^2}{2\,|\mathcal{K}|}\right)^i. \tag{4}$$

To prove (4), we consider a slightly different event. We say that *key K_i is bad* if $K_j = K_i$ for some $1 \leq i < j$. Let m_{badkeys} be the random variable counting the number of bad keys. Since every colliding key pair implies at least one bad key, then it can be shown that $\Pr[m_{\mathrm{pairs}} \geq i] \leq \Pr[m_{\mathrm{badkeys}} \geq i] \leq (N^2/2\,|\mathcal{K}|)^i$. For more details we refer to the full version [14].
Finally we prove (2) by combining (3) and (4), and by observing that from our hypothesis $N^2/|\mathcal{K}| \leq 1$:

$$\Pr[\mathrm{bad}_1] \leq \frac{2\ell - 1}{|\mathcal{T}|} \sum_{i=1}^{\lfloor N/2 \rfloor} \left(\frac{N^2}{2\,|\mathcal{K}|}\right)^i \leq \frac{2\ell - 1}{|\mathcal{T}|} \sum_{i=1}^{\infty} \frac{1}{2^i} = \frac{2\ell - 1}{|\mathcal{T}|}. \tag{5}$$

Game G^2. Game $\mathsf{G}^{2,b}_{\mathsf{A},N}$ is equivalent to $\mathsf{G}^{1,b}_{\mathsf{A},N}$, with the exception that it raises flag bad_2 in line 12 and aborts if any three keys collide. By the generalized birthday bound, and since $N^2/|\mathcal{K}| \leq 1$, we obtain

$$|\Pr[\mathsf{G}^{1,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{2,b}_{\mathsf{A},N}]| \leq \Pr[\mathrm{bad}_2] \leq \frac{1}{6}\frac{N^3}{|\mathcal{K}|^2} \leq \frac{1}{6}\frac{N}{|\mathcal{K}|}. \tag{6}$$

Game G^3. Game $\mathsf{G}^{3,b}_{\mathsf{A},N}$ is equivalent to $\mathsf{G}^{2,b}_{\mathsf{A},N}$, with the exception that the game raises flag bad_3 in line 23 and aborts if A makes a query (K, v) to F for which there exists an index $j \in [1\,..\,N]$ such that $K = K_j$ and $v \in T_j$. In the following we fix m_{inters} to be the random variable that counts the maximum number of sets $T_1, \ldots, T_N$ whose intersection is non-empty.
Fix a query (K, v) to F. For each $i \in [1\,..\,N]$ we have $\Pr[\exists j \in [1\,..\,N] : v \in T_j \wedge K = K_j \mid m_{\mathrm{inters}} = i] \leq i/|\mathcal{K}|$, because in the worst case v belongs to exactly m_{inters} of the sets $T_1, \ldots, T_N$. This bound yields

$$\begin{aligned}
&\Pr[\exists j \in [1\,..\,N] : v \in T_j \wedge K = K_j] \\
&= \sum_{i=1}^{N} \Pr[\exists j \in [1\,..\,N] : v \in T_j \wedge K = K_j \mid m_{\mathrm{inters}} = i] \cdot \Pr[m_{\mathrm{inters}} = i] \\
&\leq \sum_{i=1}^{N} \frac{i}{|\mathcal{K}|} \cdot \Pr[m_{\mathrm{inters}} = i] = \frac{1}{|\mathcal{K}|} \cdot \sum_{i=1}^{N} \Pr[m_{\mathrm{inters}} \geq i]. \qquad (7)
\end{aligned}$$

Some probabilistic considerations allow us to write $\Pr[m_{\mathrm{inters}} \geq i+1] \leq N^{i+1}\ell^i/|\mathcal{T}|^i$ (details in the full version [14]). For all $i \geq 1/\delta$ we can write $\frac{N^{i+1}\ell^i}{|\mathcal{T}|^i} \leq \left(\frac{N^{1+\delta}\ell}{|\mathcal{T}|}\right)^i \leq \frac{1}{2^i}$. Thus we can split the sum (7) into

$$\begin{aligned}
\frac{1}{|\mathcal{K}|} \cdot \sum_{i=1}^{N} \Pr[m_{\mathrm{inters}} \geq i] &\leq \frac{1}{|\mathcal{K}|}\left(\sum_{i=1}^{\lfloor 1/\delta \rfloor} \Pr[m_{\mathrm{inters}} \geq i] + \sum_{i=\lfloor 1/\delta \rfloor + 1}^{\infty} \frac{1}{2^{i-1}}\right) \\
&\leq \frac{1}{|\mathcal{K}|}\left(\frac{1}{\delta} + 1\right).
\end{aligned}$$

Since m_{inters} is constant for all q queries to F, a union bound gives us

$$|\Pr[\mathsf{G}^{2,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{3,b}_{\mathsf{A},N}]| \leq \Pr[\mathrm{bad}_3] \leq \frac{q}{|\mathcal{K}|}\left(\frac{1}{\delta} + 1\right). \qquad (8)$$

The theorem follows by combining the bounds in (2), (6), (8) for both $b = 0$ and $b = 1$ and the fact that game $\mathsf{G}^{3,b}_{\mathsf{A},N}$ is independent of the bit b.

Counter Mode with Tag Prefix. We have the following security statement on $\mathsf{CTR}\|$. Note it is slightly better than the one for $\mathsf{CTR}+$.

Theorem 6. *Suppose* $N \leq \min\left\{ |\mathcal{K}|^{1/2}, (|\mathcal{T}|/2)^{1/(1+\delta)} \right\}$, *for some* $1/N \leq \delta \leq 1$, *and suppose that* F *is modeled as a random oracle (using oracle* F*). Then for any adversary* A *against* N*-instance uniform-tag indistinguishability of* $\mathsf{CTR}\|$ *that poses at most* q *queries to* F *and no decapsulation queries we have:*

$$\mathbf{Adv}^{\text{u-miot-ind}}_{\mathsf{CTR}\|,\mathsf{A},N} \leq \frac{1}{3}\frac{N}{|\mathcal{K}|} + \frac{1}{|\mathcal{T}|} + \frac{2q}{|\mathcal{K}|}\left(1 + \frac{1}{\delta}\right).$$

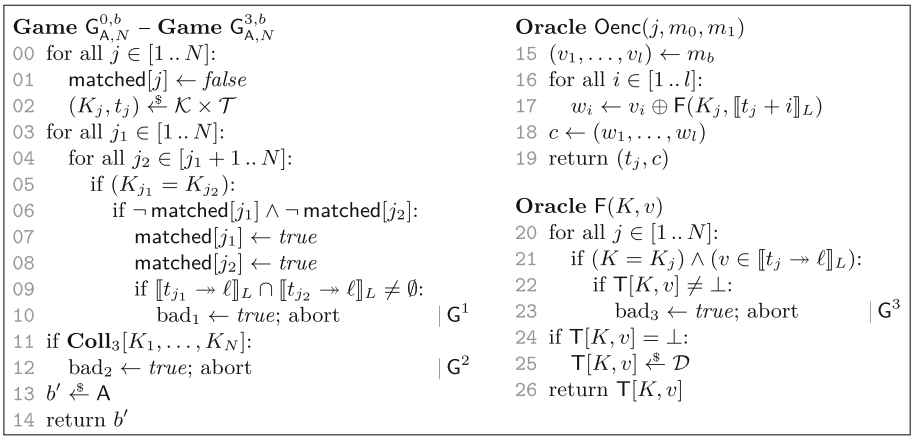

Fig. 11. The security game $\mathsf{G}^{0,b}_{\mathsf{A},N}$ for CTR+ in the random oracle model, and games $\mathsf{G}^{1,b}_{\mathsf{A},N}$, $\mathsf{G}^{2,b}_{\mathsf{A},N}$, and $\mathsf{G}^{3,b}_{\mathsf{A},N}$. Adversary A can query the oracle Oenc at most once for the same index j.

Proof. We refer to Fig. 12 for the definition of the games $\mathsf{G}^{0,b}_{\mathsf{A},N}$, $\mathsf{G}^{1,b}_{\mathsf{A},N}$, $\mathsf{G}^{2,b}_{\mathsf{A},N}$ and $\mathsf{G}^{3,b}_{\mathsf{A},N}$. Except for some bookkeeping, game $\mathsf{G}^{0,b}_{\mathsf{A},N}$ is equivalent to the security game $\mathsf{U\text{-}MIOT\text{-}IND}^b_{\mathsf{A},N}$, with $b \in \{0, 1\}$.

Game G^1. Game $\mathsf{G}^{1,b}_{\mathsf{A},N}$ is equivalent to $\mathsf{G}^{0,b}_{\mathsf{A},N}$, except when any three keys collide. By the generalized birthday bound, and since $N^2/|\mathcal{K}| \leq 1$, we obtain

$$|\Pr[\mathsf{G}^{0,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{1,b}_{\mathsf{A},N}]| \leq \Pr[\text{bad}_1] \leq \frac{1}{6}\frac{N^3}{|\mathcal{K}|^2} \leq \frac{1}{6}\frac{N}{|\mathcal{K}|}. \tag{9}$$

Game G^2. In game $\mathsf{G}^{2,b}_{\mathsf{A},N}$ we abort when two events occur simultaneously: a key 2-collision and collision of the corresponding tags. The probability to abort is by the generalized birthday bound, the independence of the two events, and the condition $N^2/|\mathcal{K}| \leq 1$:

$$|\Pr[\mathsf{G}^{1,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{2,b}_{\mathsf{A},N}]| \leq \Pr[\text{bad}_2] \leq \frac{N^2}{2\,|\mathcal{K}|}\frac{1}{|\mathcal{T}|} \leq \frac{1}{2}\frac{1}{|\mathcal{T}|}. \tag{10}$$

Game G^3. Game $\mathsf{G}^{3,b}_{\mathsf{A},N}$ is equivalent to $\mathsf{G}^{2,b}_{\mathsf{A},N}$, with the exception that the game raises flag bad$_3$ in line 16 if some specific condition is met. To get an upper bound on the probability to distinguish $\mathsf{G}^{2,b}_{\mathsf{A},N}$ and $\mathsf{G}^{3,b}_{\mathsf{A},N}$ we compute the probability that the adversary explicitly queries F for an input $(K, v \| [\![i]\!]_L)$ such that for some $j \in [1\,..\,N]$, $K = K_j$ and $v = t_j$. This leads to the equation:

$$|\Pr[\mathsf{G}^{2,b}_{\mathsf{A},N}] - \Pr[\mathsf{G}^{3,b}_{\mathsf{A},N}]| \leq \Pr[\text{bad}_3] \leq \frac{q}{|\mathcal{K}|}\left(\frac{1}{\delta} + 1\right). \tag{11}$$

Fix a query $(K, v \| [\![i]\!]_L)$ to F. Since the adversary knows all possible values of v used by Oenc after each call, the adversary must only guess the key. Assume that there are at most m_{coll} keys that use the same tag value v. Then the probability that flag bad_3 is triggered during this query is in the worst case $m_{\text{coll}}/|\mathcal{T}|$. We compute the probability of this event as follows.

$$\begin{aligned}
&\Pr[\exists j \in [1\,..\,N] : v = t_j \wedge K = K_j] \\
&= \sum_{i=1}^{N} \Pr[\exists j \in [1\,..\,N] : v = t_j \wedge K = K_j \mid m_{\text{coll}} = i] \cdot \Pr[m_{\text{coll}} = i] \\
&\leq \sum_{i=1}^{N} \frac{i}{|\mathcal{K}|} \cdot \Pr[m_{\text{coll}} = i] = \frac{1}{|\mathcal{K}|} \cdot \sum_{i=1}^{N} \Pr[m_{\text{coll}} \geq i]. \qquad (12)
\end{aligned}$$

The last equality follows since the expected value of any random variable m with values in $\mathbb{N}$ can be written as $\sum_{i=0}^{\infty} i \Pr[m = i] = \sum_{i=1}^{\infty} \Pr[m \geq i]$. Now we estimate the probability $\Pr[m_{\text{coll}} \leq i]$. Assume that $i \geq 1/\delta$. Then from the generalized birthday bound and the condition $N \leq (|\mathcal{T}|/2)^{1/(1+\delta)}$ we can write:

$$\Pr[m_{\text{coll}} \geq i+1] \leq \frac{N^{i+1}}{(i+1)!\,|\mathcal{T}|^i} \leq \left(\frac{N^{1+\delta}}{|\mathcal{T}|}\right)^i \leq \frac{1}{2^i}\,.$$

Considering this observation we split the sum in Eq. (12) into

$$\begin{aligned}
\frac{1}{|\mathcal{K}|} \cdot \sum_{i=1}^{N} \Pr[m_{\text{coll}} \geq i] &\leq \frac{1}{|\mathcal{K}|} \left(\sum_{i=1}^{\lfloor 1/\delta \rfloor} \Pr[m_{\text{coll}} \geq i] + \sum_{i=\lfloor 1/\delta \rfloor + 1}^{\infty} \frac{1}{2^{i-1}} \right) \\
&\leq \frac{1}{|\mathcal{K}|} \left(\frac{1}{\delta} + 1 \right).
\end{aligned}$$

Since m_{coll} is constant for all queries to F, a union bound yields our claim:

$$\Pr[\text{bad}_3] \leq \frac{q}{|\mathcal{K}|} \left(\frac{1}{\delta} + 1 \right).$$

The theorem follows by combining the bounds in (9), (10), (11) for both $b = 0$ and $b = 1$ and the fact that game $\mathsf{G}^{3,b}_{\mathsf{A},N}$ is independent of b.

6.3 On the Security of Permutation-Based Counter Mode

In above Theorem 5 we assessed the security of the CTR+ ADEM, defined with respect to a function $F\colon \mathcal{K} \times \mathcal{B} \to \mathcal{D}$. The analysis modeled F as an ideal random function and showed that using sets $\mathcal{K}$ and $\mathcal{B}$ of moderate size (e.g., of cardinality 2^{128}) is sufficient to let CTR+ achieve security. We next show that if F is instead instantiated with a blockcipher and modeled as an ideal family of

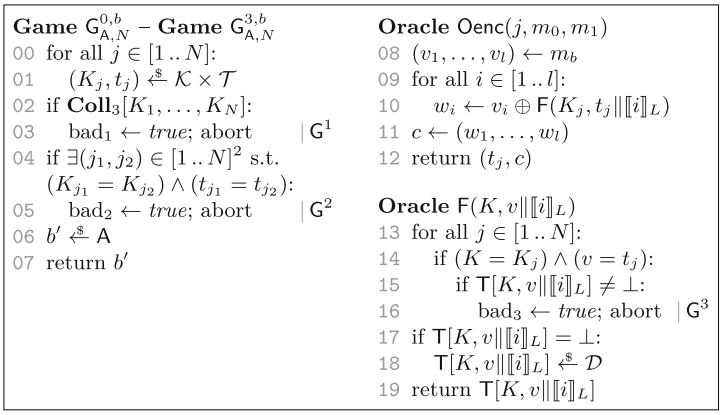

Fig. 12. The security game $\mathsf{G}^{0,b}_{\mathsf{A},N}$ for $\mathsf{CTR}\|$ in the random oracle model, and games $\mathsf{G}^{1,b}_{\mathsf{A},N}$, $\mathsf{G}^{2,b}_{\mathsf{A},N}$, and $\mathsf{G}^{3,b}_{\mathsf{A},N}$. Adversary A can query the oracle Oenc at most once for the same index j.

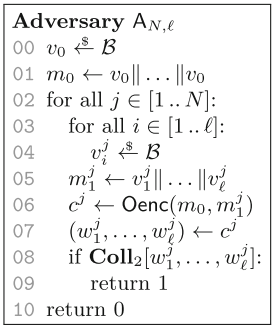

Fig. 13. Definition of adversary $\mathsf{A}_{N,\ell}$ against U-MIOT-IND security of CTR+ instantiated with a permutation $F(K, \cdot)$. In line 01 message m_0 is made of ℓ identical blocks.

permutations, then the minimum cardinality of $\mathcal{B} = \mathcal{D}$ for achieving security is considerably increased (e.g., to values around 2^{256}).

Our argument involves the analysis of a U-MIOT-IND adversary A that is specified in Fig. 13. Effectively, the idea of the attack is exploiting the tightness gap of the PRP/PRF switching lemma [5] via the multi-instance setting. More concretely, the adversary repeats the following multiple times (once for each instance): It asks either for the encapsulation of a message comprised of identical blocks, or for the encapsulation of a message consisting of uniformly-generated blocks. The adversary outputs 1 if any two blocks that form the ciphertext collide. If the ciphertext is the encapsulation of the identical-block message then the adversary does not find a collision, since $F(K, \cdot)$ is a permutation for each key $K \in \mathcal{K}$ and is evaluated on distinct input values. Otherwise the ciphertext blocks are random, and one can thus find a collision.

The theorem uses the technical condition that $N\ell(\ell-1)/|\mathcal{T}| \leq 4$, where ℓ is a parameter that determines the length of the encapsulated messages, measured in blocks. Note that adversaries that could process values N, ℓ that are too large to fulfill this bound will reach at least the same advantage as adversaries considered by the theorem, simply by refraining from posing queries. The stated lower-bound is roughly $N\ell^2/|\mathcal{T}|$ and effectively induced by N applications of the PRP/PRF switching lemma. Note that if the above condition is met with equality, the adversary's advantage is at least 1/2. Further, if $|\mathcal{T}| = |\mathcal{B}| = 2^{128}$, $\ell = 2^{40}$ (this corresponds to a message length of 16 terabytes) and we have $N = 2^{48}$ instances, the success probability of A is about 1/8, or larger.

Theorem 7. *Consider* CTR+ *instantiated with a family of permutations* $F(K,\cdot)$ *over* $\mathcal{B}$*, and let* $N \geq 2$*. Assume moreover that* $N\ell(\ell-1) \leq 4 \cdot |\mathcal{T}|$*. Then for the adversary* A *in Fig. 13 it holds:*

$$\mathbf{Adv}^{\text{u-miot-ind}}_{\text{CTR+},\mathsf{A},N} \geq \frac{N\ell(\ell-1)}{8 \cdot |\mathcal{T}|}.$$

The adversary has a running time of $\mathcal{O}(N\ell \log \ell)$*, makes* N *queries to* Oenc *for messages of length at most* ℓ *and makes no* Odec *queries.*

Proof. We start with the analysis of the running time of A: It is predominantly determined by the search for collisions among ℓ blocks for each of the N iterations of the main loop, hence the bound of $\mathcal{O}(N\ell \log \ell)$ on the time. We now compute the probability that the adversary outputs 1 depending on the game bit b.

Case U-MIOT-IND^0. For each instance $j \in [1\,..\,N]$ the adversary obtains an encapsulation of a sequence of identical blocks. All blocks composing c^j must be distinct, since for each key K, function $F(K,\cdot)$ is a permutation over $\mathcal{B}$. Therefore the output of this game is always 0 and we have $\Pr[\text{U-MIOT-IND}^0_{\mathsf{A},N}] = 0$.

Case U-MIOT-IND^1. Let p be the probability that there is a collision between ℓ random variables that are uniformly distributed in the set $\mathcal{B}$. We show that for each $j \in [1\,..\,N]$ the probability of A to output 1 when running the j-th iteration of the loop is p. From the definition of Oenc we can write $w_i^j = v_i^j \oplus F(K_j, [\![t_j + i]\!]_L)$ for each $i \in [1\,..\,\ell]$, where K_j and t_j are the key-tag pairs generated by the game $\text{U-MIOT-IND}^1_{\mathsf{A},N}$. The elements $v_1^j, \ldots, v_\ell^j$ are generated uniformly in $\mathcal{B}$ and independently of K_j, t_j, their index, and from each other. Hence the elements $w_1^j, \ldots, w_\ell^j$ are also uniformly distributed in $\mathcal{B}$ and mutually independent, even in the presence of colliding keys among $K_1, \ldots, K_N$. Since all blocks v_i^j with $i \in [1\,..\,\ell]$ and $j \in [1\,..\,N]$ are independently random, the probability that the adversary outputs 1 is:

$$\Pr[\text{U-MIOT-IND}^1_{\mathsf{A},N}] = 1 - (1-p)^N. \tag{13}$$

Since $\ell(\ell-1) \leq 2\,|\mathcal{B}| = 2\,|\mathcal{T}|$ by our hypotheses we can use the birthday bound to bound the probability p as $p \geq \ell(\ell-1)/(4 \cdot |\mathcal{B}|)$. With some simple algebra, and since $N\ell(\ell-1) \leq 4\,|\mathcal{T}| = 4\,|\mathcal{B}|$, we can bound Eq. 13 as:

$$\Pr[\mathsf{U\text{-}MIOT\text{-}IND}^1_{\mathsf{A},N}] \geq \min\left\{\frac{1}{2}, \frac{Np}{2}\right\} \geq \frac{N\ell(\ell-1)}{8\cdot|\mathcal{B}|} = \frac{N\ell(\ell-1)}{8\cdot|\mathcal{T}|}.$$

7 ADEMs Secure Against Active Adversaries

In the preceding section we proposed two ADEMs and proved them multi-instance secure against passive adversaries. However, the constructions are based on counter mode encryption and obviously vulnerable in settings with active adversaries that manipulate ciphertexts on the wire. In this section we alleviate the situation by constructing ADEMs that remain secure in the presence of active attacks. Concretely, in line with the encrypt-then-MAC approach [6], we show that an ADEM that is secure against active adversaries can be built from one that is secure against passive adversaries by tamper-protecting its ciphertexts using a message authentication code (MAC). More precisely, with the goal of *tightly* achieving multi-instance security, we use an *augmented message authentication code* (see footnote 4) (AMAC) where the generation and verification algorithms depend on an auxiliary input: the tag. In the combined construction, the same tag is used for both ADEM and AMAC. As before, using KEM ciphertexts as tags is a reasonable choice. We conclude the section by constructing a (tightly) secure AMAC based on a hash function.

7.1 Augmented Message Authentication

Augmented message authentication. An augmented message authentication code $\mathsf{AMAC} = (\mathsf{M.mac}, \mathsf{M.vrf})$ for a message space $\mathcal{M}$ is a pair of deterministic algorithms associated with a finite key space $\mathcal{K}$, a tag space $\mathcal{T}$, and a code space $\mathcal{C}$. The algorithm $\mathsf{M.mac}$ takes a key $K \in \mathcal{K}$, a tag $t \in \mathcal{T}$, and a message $m \in \mathcal{M}$, and outputs a code $c \in \mathcal{C}$. The verification algorithm $\mathsf{M.vrf}$ takes a key $K \in \mathcal{K}$, a tag $t \in \mathcal{T}$, a message $m \in \mathcal{M}$, and a code $c \in \mathcal{C}$, and outputs either *true* or *false*. The correctness requirement is that for all $K \in \mathcal{K}$, $t \in \mathcal{T}$, $m \in \mathcal{M}$ and $c \in [\mathsf{M.mac}(K, t, m)]$ we have $\mathsf{M.vrf}(K, t, m, c) = \mathit{true}$.

Augmented message authentication with nonces. We give a game-based authenticity model for AMACs.[11] In our model, for each of a total of N independent keys the adversary can request one MAC code computation but many verifications. The restriction is that for each key the MAC query has to precede all verification queries, and that always the same tag is used. Further, in

[11] In principle we could give two security definitions: one using uniform tags and one using nonce tags. In this paper we formalize only the latter, not the former, for mainly two reasons: (a) the nonce-based notion is not required for our results; (b) in the nonce setting it is not clear how to prove a result similar to the one of Theorem 8. The reason for (b) is that to simulate an encapsulation query for a U-MIOT-IND adversary using an AMAC oracle one must specify the tag that is also used to generate the DEM ciphertext, but this is only given as an output of the AMAC oracle.

line with the definition of nonce-based security for ADEMs, we require the tag provided in each MAC computation request to be unique (across all instances). We formalize the corresponding security notion of (strong) nonce-based multi-instance one-time unforgeability for AMACs via the game specified in Fig. 14. For a scheme AMAC, to any adversary A and any number of instances N we associate the advantage $\mathbf{Adv}^{\text{n-miot-uf}}_{\mathsf{AMAC},\mathsf{A},N} := \Pr[\mathsf{N\text{-}MIOT\text{-}UF}_{\mathsf{A},N}]$.

Game $\mathsf{N\text{-}MIOT\text{-}UF}_{\mathsf{A},N}$	**Oracle** $\mathsf{Omac}(j,t,m)$	**Oracle** $\mathsf{Ovrf}(j,m,c)$
00 $\mathsf{forged} \leftarrow 0$	07 if $C_j \neq \emptyset$: return $\perp$	13 if $C_j = \emptyset$: return $\perp$
01 $T \leftarrow \emptyset$	08 if $t \in T$: return $\perp$	14 if $(m,c) \in C_j$: return $\perp$
02 for all $j \in [1\,..\,N]$:	09 $T \leftarrow T \cup \{t\}$; $t_j \leftarrow t$	15 if $\mathsf{M.vrf}(K_j, t_j, m, c)$:
03 $\quad K_j \xleftarrow{\$} \mathcal{K}$	10 $c \leftarrow \mathsf{M.mac}(K_j, t_j, m)$	16 $\quad \mathsf{forged} \leftarrow 1$
04 $\quad C_j \leftarrow \emptyset$	11 $C_j \leftarrow C_j \cup \{(m,c)\}$	17 $\quad$ return *true*
05 run A	12 return c	18 return *false*
06 return forged		

Fig. 14. AMAC security game $\mathsf{N\text{-}MIOT\text{-}UF}_{\mathsf{A},N}$, modeling nonce-based multi-instance one-time unforgeability for N instances. The tags in line 15 are the same as the ones in line 10.

7.2 The ADEM-Then-AMAC Construction

Let ADEM and AMAC be an ADEM and an AMAC, respectively. Following the generic encrypt-then-MAC [6] composition technique, and assuming ADEM is secure against passive adversaries, we combine the two schemes to obtain the augmented data-encapsulation mechanism ADEM′, which we prove secure against active adversaries. More formally, if $\mathsf{ADEM} = (\mathsf{A.enc}, \mathsf{A.dec})$ and $\mathsf{AMAC} = (\mathsf{M.mac}, \mathsf{M.vrf})$ have key spaces $\mathcal{K}_{\mathsf{dem}}$ and $\mathcal{K}_{\mathsf{mac}}$, respectively, then the key space of ADEM′ is $\mathcal{K}_{\mathsf{dem}} \times \mathcal{K}_{\mathsf{mac}}$, and its algorithms are as in Fig. 15. Note that the tag space is the same for all three schemes (and that the message spaces have to be sufficiently compatible to each other).

Proc $\mathsf{A.enc}'(K,t,m)$	**Proc** $\mathsf{A.dec}'(K,t,c)$
00 $(K_{\mathsf{dem}}, K_{\mathsf{mac}}) \leftarrow K$	05 $(K_{\mathsf{dem}}, K_{\mathsf{mac}}) \leftarrow K$
01 $c_{\mathsf{dem}} \leftarrow \mathsf{A.enc}(K_{\mathsf{dem}}, t, m)$	06 $(c_{\mathsf{dem}}, c_{\mathsf{mac}}) \leftarrow c$
02 $c_{\mathsf{mac}} \leftarrow \mathsf{M.mac}(K_{\mathsf{mac}}, t, c_{\mathsf{dem}})$	07 if $\mathsf{M.vrf}(K_{\mathsf{mac}}, t, c_{\mathsf{dem}}, c_{\mathsf{mac}})$:
03 $c \leftarrow (c_{\mathsf{dem}}, c_{\mathsf{mac}})$	08 $\quad m \leftarrow \mathsf{A.dec}(K_{\mathsf{dem}}, t, c_{\mathsf{dem}})$
04 return c	09 $\quad$ return m
	10 return $\perp$

Fig. 15. Construction of ADEM′ from ADEM and AMAC.

The proof of the following theorem can be found in the full version [14].

Theorem 8. *Let* ADEM' *be constructed from* ADEM *and* AMAC *as described. Then for any number of instances* N *and any ADEM adversary* A *that poses at most* Q_d*-many* Odec *queries, there exist an AMAC adversary* B *and an ADEM adversary* C *such that*

$$\mathbf{Adv}^{\mathsf{n\text{-}miot\text{-}ind}}_{\mathsf{ADEM}',\mathsf{A},N} \leq 2\mathbf{Adv}^{\mathsf{n\text{-}miot\text{-}uf}}_{\mathsf{AMAC},\mathsf{B},N} + \mathbf{Adv}^{\mathsf{n\text{-}miot\text{-}ind}}_{\mathsf{ADEM},\mathsf{C},N}.$$

The running time of B *is at most that of* A *plus the time required to run* N*-many* ADEM *encapsulations and* Q_d*-many* ADEM *decapsulations. The running time of* C *is the same as the running time of* A*. Moreover,* B *poses at most* Q_d*-many* Ovrf *queries, and* C *poses no* Odec *query.*

7.3 A Multi-instance Secure AMAC

A random oracle directly implies a multi-instance secure AMAC, with a straightforward construction: the MAC code of a message is computed by concatenating key, tag, and message, and hashing the result. We formalize this as follows. Let $\mathcal{T}$ be a tag space and $\mathcal{M}$ a message space. Let $\mathcal{K}$ and $\mathcal{C}$ be arbitrary finite sets. Let $H\colon \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{C}$ be a hash function. Define function $\mathsf{M.mac}$ and a predicate $\mathsf{M.vrf}$ such that for all K, t, m, c we have $\mathsf{M.mac}(K, t, m) = H(K, t, m)$, and $\mathsf{M.vrf}(K, t, m, c) = \mathit{true}$ iff $H(K, t, m) = c$. Let finally $\mathsf{AMAC} = (\mathsf{M.mac}, \mathsf{M.vrf})$.

Note that hash functions based on the Merkle–Damgård design, like SHA256, do not serve directly as random oracles due to generic length-extension attacks [10], and indeed the ADEM' scheme from Fig. 15 is not secure if its AMAC is derived from such a function. Fortunately, Merkle–Damgård hashing can be modified to achieve indifferentiability from a random oracle [10]. Further, more recent hash functions like SHA3 are naturally resilient against length-extension attacks.

The proof of the following theorem can be found in the full version [14].

Theorem 9. *Let* $\mathcal{K}, \mathcal{T}, \mathcal{M}, \mathcal{C}$ *and* $\mathsf{AMAC} = (\mathsf{M.mac}, \mathsf{M.vrf})$ *be as above. If* H *behaves like a (non-programmable) random oracle, for any number of instances* N *and any adversary* A *we obtain*

$$\mathbf{Adv}^{\mathsf{n\text{-}miot\text{-}uf}}_{\mathsf{AMAC},\mathsf{A},N} \leq \frac{q}{|\mathcal{K}|} + \left(\frac{1}{|\mathcal{K}|} + \frac{1}{|\mathcal{C}|}\right) Q_v,$$

where q *is the number of direct calls to the random oracle by the adversary, and* Q_v *is the number of calls to the oracle* Ovrf*. Note that the bound does not depend on the number of* Omac *queries.*

Acknowledgments. We are grateful to Krzysztof Pietrzak and the anonymous reviewers for their valuable comments. The authors were partially supported by ERC Project ERCC (FP7/615074) and by DFG SPP 1736 Big Data.

References

1. Attrapadung, N., Hanaoka, G., Yamada, S.: A framework for identity-based encryption with almost tight security. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015 Part I. LNCS, vol. 9452, pp. 521–549. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_22
2. Bellare, M.: New proofs for NMAC and HMAC: security without collision resistance. J. Cryptol. **28**(4), 844–878 (2015)
3. Bellare, M., Bernstein, D.J., Tessaro, S.: Hash-function based PRFs: AMAC and its multi-user security. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016 Part I. LNCS, vol. 9665, pp. 566–595. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_22
4. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_18
5. Bellare, M., Kilian, J., Rogaway, P.: The security of cipher block chaining. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 341–358. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48658-5_32
6. Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44448-3_41
7. Bellare, M., Tackmann, B.: The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016 Part I. LNCS, vol. 9814, pp. 247–276. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_10
8. Chatterjee, S., Koblitz, N., Menezes, A., Sarkar, P.: Another look at tightness II: practical issues in cryptography. Cryptology ePrint Archive, Report 2016/360 (2016)
9. Cogliani, S., Maimuţ, D.Ş., Naccache, D., do Canto, R.P., Reyhanitabar, R., Vaudenay, S., Vizár, D.: OMD: a compression function mode of operation for authenticated encryption. In: Joux, A., Youssef, A. (eds.) SAC 2014. LNCS, vol. 8781, pp. 112–128. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13051-4_7
10. Coron, J.-S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: how to construct a hash function. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_26
11. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM J. Comput. **33**(1), 167–226 (2003)
12. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016 Part I. LNCS, vol. 9665, pp. 1–27. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_1
13. Gaži, P., Pietrzak, K., Tessaro, S.: Generic security of NMAC and HMAC with input whitening. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015 Part II. LNCS, vol. 9453, pp. 85–109. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48800-3_4
14. Giacon, F., Kiltz, E., Poettering, B.: Hybrid encryption in a multi-user setting, revisited. Cryptology ePrint Archive, Report 2017/843 (2017)

15. Gong, J., Chen, J., Dong, X., Cao, Z., Tang, S.: Extended nested dual system groups, revisited. In: Cheng, C.-M., Chung, K.-M., Persiano, G., Yang, B.-Y. (eds.) PKC 2016 Part I. LNCS, vol. 9614, pp. 133–163. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49384-7_6
16. Herranz, J., Hofheinz, D., Kiltz, E.: Some (in)sufficient conditions for secure hybrid encryption. Inf. Comput. **208**(11), 1243–1257 (2010)
17. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 590–607. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_35
18. Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_31
19. Libert, B., Joye, M., Yung, M., Peters, T.: Concise multi-challenge CCA-secure encryption and signatures with almost tight security. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014 Part II. LNCS, vol. 8874, pp. 1–21. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45608-8_1
20. Libert, B., Peters, T., Joye, M., Yung, M.: Compactly hiding linear spans. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015 Part I. LNCS, vol. 9452, pp. 681–707. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_28
21. Patarin, J.: Security in $O(2^n)$ for the xor of two random permutations–proof with the standard H technique. Cryptology ePrint Archive, Report 2013/368 (2013)
22. Zaverucha, G.: Hybrid encryption in the multi-user setting. Cryptology ePrint Archive, Report 2012/159 (2012)