

# Authenticated Encryption with Key Identification

Julia Len<sup>1</sup>, Paul Grubbs<sup>2</sup>, and Thomas Ristenpart<sup>1</sup>

<sup>1</sup> Cornell Tech

<sup>2</sup> University of Michigan

**Abstract.** Authenticated encryption with associated data (AEAD) forms the core of much of symmetric cryptography, yet the standard techniques for modeling AEAD assume recipients have no ambiguity about what secret key to use for decryption. This is divorced from what occurs in practice, such as in key management services, where a message recipient can store numerous keys and must identify the correct key before decrypting. To date there has been no formal investigation of their security properties or efficacy, and the ad hoc solutions for identifying the intended key deployed in practice can be inefficient and, in some cases, vulnerable to practical attacks.

We provide the first formalization of nonce-based AEAD that supports key identification (AEAD-KI). Decryption now takes in a vector of secret keys and a ciphertext and must both identify the correct secret key and decrypt the ciphertext. We provide new formal security definitions, including new key robustness definitions and indistinguishability security notions. Finally, we show several different approaches for AEAD-KI and prove their security.

**Keywords:** Key identification, authenticated encryption, key commitment, key robustness

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Preliminaries</b>	<b>5</b>
<b>3</b>	<b>Defining AEAD with Key Identification</b>	<b>7</b>
3.1	Key robustness .....	8
3.2	All-in-one confidentiality and integrity .....	9
<b>4</b>	<b>Key Labels</b>	<b>11</b>
<b>5</b>	<b>Static Key Identifiers</b>	<b>14</b>
5.1	Static Key Hint .....	15
5.2	Static Key Commitment .....	16
<b>6</b>	<b>Dynamic Key Identifiers</b>	<b>17</b>
6.1	Dynamic Key Hint .....	18
6.2	Dynamic Key Commitment .....	19
<b>A</b>	<b>Extended Full Robustness</b>	<b>22</b>
<b>B</b>	<b>KI-nAE implies INT-CTXT</b>	<b>23</b>
<b>C</b>	<b>Proof of Theorem 2</b>	<b>24</b>
<b>D</b>	<b>Proof of Theorem 3</b>	<b>27</b>
<b>E</b>	<b>Proofs of Theorem 5 and Theorem 7</b>	<b>30</b>
E.1	Proof of Theorem 5 .....	31
E.2	Proof of Theorem 7 .....	33
<b>F</b>	<b>Proofs of Theorem 9 and Theorem 11</b>	<b>35</b>
F.1	Proof of Theorem 9 .....	35
F.2	Proof of Theorem 11 .....	38

## 1 Introduction

Authenticated encryption with associated data (AEAD) is ubiquitously used in practice. Standard formalizations of AEAD schemes model a “single-key” setting where a single sender sends an encrypted message to a single receiver. Though simple to analyze, this model is increasingly divorced from practice in a number of important aspects.

A setting which has received little or no attention in the cryptographic literature is one where a message recipient will store numerous keys and must identify the correct key to use before decryption can proceed. This practice can be seen in cryptographic libraries such as Google’s Tink API [32], key management services (KMS) such as for Amazon Web Services (AWS) [5], and multi-user Shadowsocks [30], among others. Notably, AEAD schemes and their security models do not formally address the issue of key identification, producing a gap when translating from cryptographic theory to practice.

Two approaches for key identification are often used in practice. The first is trial decryption, where one attempts to decrypt the ciphertext under each of the keys held by the recipient. However, this is slow, even for a small number of keys. As such, a second approach is often used: the sender attaches a previously agreed-upon key identifier to each message it sends. For instance, this approach is used by Tink, which derives 5-byte strings as the identifier. The recipient can then efficiently use the identifier to look up the corresponding key. However, this approach does not work in settings where keys must remain anonymous, such as in anonymous messaging protocols. It is also unclear what security properties (for both approaches) are being achieved in the case where there is potential for adversarial modification of key identifiers and/or adversarial choice of some of the recipient keys. Adversarially chosen keys can arise in settings where the sender chooses a secret key to share with the recipient, which have resulted in several recent attacks [2, 15, 19, 23].

One example is multi-user Shadowsocks [30]. Shadowsocks is an anonymity proxy that works by having a client encrypt their traffic under a shared password with a server, which decrypts using AEAD. The multi-user mode allows multiple passwords to be specified for a single server, which means that incoming packets must be trial decrypted under every possible user password. Len et al. [23] describe an attack on this scheme where an attacker can insert a malicious password into this set of passwords, then mount a partitioning oracle attack that enables the attacker to learn some target user’s password. Fundamentally, the vulnerability is that Shadowsocks’s AEAD has no efficient and secure way to identify the appropriate key.

**Our contributions.** We initiate the formal study of AEAD that supports key identification. The starting point is nonce-based AEAD [28], which we extend to include in the formal syntax and semantics of encryption schemes the key identification task: decryption takes in a vector of secret keys as well as a nonce, associated data, and a ciphertext, and must both identify the correct secret key and decrypt the ciphertext. This change, while conceptually simple, immediately introduces a number of complexities. It forces scheme designers to specify how the right key is identified, requires changes to correctness notions, suggests that we must give new security definitions that speak to issues like adversaries forcing the wrong key to be identified, and more.

We formalize a new cryptographic primitive called AEAD-KI, or AEAD with key identification. Like AEAD, the primitive is composed of a triple of algorithms for key generation, encryption, and decryption. Key generation takes in what we call a key label so that AEAD-KI keys are composed of the traditional secret key as well as the key label. This label acts as optional public metadata for the key and models techniques in practice, e.g., URLs that suggest where to locate the key or other kinds of static identifiers. Encryption takes in a key, nonce, associated data, and message. Ciphertexts can opt to include a special component, called a key tag. Decryption for AEAD-KI, in turn, accepts a vector of keys, instead of one key as for AEAD. We use a vector instead of a set to preserve information about the order of keys, which could affect the decryption outcome. Given the ciphertext (including the key tag), decryption returns both the key that correctly decrypts the ciphertext as well as the resulting message. If decryption determines that no key correctly decrypts the ciphertext, it simply returns an error symbol  $\perp$ .

We next consider which security definitions best capture the AEAD-KI setting. Our first goal is to extend the standard AEAD security notions of confidentiality and ciphertext integrity to AEAD-KI. A good starting point is transforming the traditional all-in-one real-versus-random indistinguishability security notion for AEAD due to Rogaway and Shrimpton [29] to the AEAD-KI setting. Specifically, we can allow the attacker to interact with multiple encryption key instances, reminiscent of the multi-user setting for encryption [6]. However, this definition only allows for honest keys, which unfortunately does not capture attacks in which

Approach	Description	AEAD FROB?	Key anon.?	Section
Key labels	Key gen. labels each key, sent as part of ciphertext; brute-force decrypt with all keys matching key label in ciphertext	Yes	No	§4
Trial decryption	Special case of key labels where all labels are empty ( $\epsilon$ )	Yes	Yes	§4
Static key hint	Ciphertext includes deterministic non-CR hash of key	Yes	No	§5
Static key commitment	Ciphertext includes deterministic CR hash of key	No	No	§5
Dynamic key hint	Ciphertext includes PRF of key & nonce	Yes	Yes	§6
Dynamic key commitment	Ciphertext includes CR-PRF of key & nonce	No	Yes	§6

Fig. 1: Summary of various approaches to AEAD-KI that we consider. For each approach, we also list whether it must use a FROB AEAD scheme and whether it provides key anonymity.

a malicious key is somehow inserted into a recipient’s key vector. Indeed, attacks have already been shown in practice where malicious keys are given to a recipient to prevent decryption under honest keys [2, 15].

We therefore opt for a notion of security for which adversaries can insert malicious keys into key vectors used during decryption queries. This renders more complex how the security game should handle decryption oracles in order to distinguish between honest and malicious keys. To handle these subtleties, we introduce KI-nAE, a new security notion that uses a simulation-based approach for the all-in-one definition. This definition captures a wide class of interference attacks in which a malicious user somehow inserts a malicious key into a recipient’s key vector in an attempt to interfere with honest keys.

A security property intrinsic to the key identification setting is key robustness, a security goal first investigated in the context of public-key encryption [1] and later investigated for authenticated encryption [17] (see also [2, 8, 15, 19, 23]). Interestingly, robustness here functions as a form of correctness, as ensuring the correct key decrypts a given ciphertext in the presence of many (potentially adversarial) keys can only be guaranteed by robustness. We thus extend the AEAD robustness notion called full robustness (FROB) [17] to the AEAD-KI setting, which is straightforward. Interestingly, this extension however proves insufficient to rule out some attacks. In particular, when decryption is given the correct key within a key set, it should not fail to decrypt; such failures could leak information about the (honest) keys composing a key vector. One way to handle this is with an extended key robustness notion, but an observation due to Mihir Bellare is that one can instead extend correctness to rule out such decryption failures. See the body and Appendix A for more details.

**Approaches to AEAD-KI.** We then turn to analyzing security of existing key identification schemes as well as suggesting new ones. A summary of our analyses appears in Figure 1. We divide key identification into several categories. The first approach utilizes the key label of the key as the key tag itself. Decryption can then find all keys whose label matches the key tag of the ciphertext and perform trial decryptions. (Note the second category, trial decryption, is a special case where all labels are the empty string.) This reflects how, in practice, labels are sometimes not unique and instead used to label a set of keys. Using key labels obviates achieving key anonymity, since a ciphertext produced by a certain key will always be flagged by the key’s label. However, brute-force trial decryption, a special case of the key label approach, can achieve anonymity, although with the trade-off of increased computational costs. Since key labels are not unique, they do not provide key commitment to AEAD schemes that are not key committing. Nevertheless, we see this insecure construction arises in practice, such as in the Tink library. Our analysis shows that if instead the key label approach is used with a key-committing AEAD scheme, then the composition is key committing.

Next we turn to what we call “static” approaches, those where key tags are deterministically computed from the secret key. These are similar to “key check values”, legacy schemes for ensuring integrity of the key [18, 27, 31]. Since the key tag for a key never changes, this approach also does not allow for key anonymity. We further divide static key identifiers into two classes: static key hints and static key commitments. Key hints are key tags computed from the key in a non-key-committing way, typically using a non-collision resistant hash of the key. This means that AEAD-KI schemes using key hints, like key labels, will need to trial decrypt on all keys matching the key hint and use a FROB AEAD scheme to achieve key robustness. The benefit of key hints is that they can often be short and efficiently computed. In contrast, static key

commitments, the second class of static key identifiers, do commit to the key. These are typically a collision-resistant hash of the key. While static key commitments might be less efficient to compute than key hints, they can be used to build secure AEAD-KI from non-FROB AEAD schemes.

Finally, both key hints and key identifiers can be made “dynamic” to provide key anonymous counterparts of static schemes. This approach uses a nonce when computing the key tag so that the key tag for a key is unique for each encryption call. Like static identifiers, dynamic identifiers can be both key hints and key commitments, with security achieved when combining with any AEAD or FROB AEAD, respectfully.

**Further related work.** Farshim, Orlandi, and Rosie [17] first proposed the set of key robustness notions for symmetric primitives. Their strongest definition, full robustness, represents the goal of key commitment for AEAD schemes. Grubbs et al. [19] suggest a notion of compactly committing AEAD, which is useful in abuse moderation settings. Dodis et al. [15] describe an attack against Facebook Messenger’s abuse moderation tooling that relies on AES-GCM, which is not key committing. In this attack, a malicious sender can force an honest recipient to use a malicious key to decrypt a message to abusive content that cannot be reported. Len et al. [23] and Albertini et al. [2] showed that other commonly used AEAD schemes, such as AES-GCM-SIV, ChaCha20-Poly1305, and OCB3, are not key committing and they describe other practical attack scenarios that exploit non-key committing AEAD schemes.

Albertini et al. also propose two approaches to adding key commitment to AEAD schemes. One of these approaches is to compute a collision-resistant PRF of the key, both deterministically and using a nonce for key anonymity. Our static and dynamic identifiers parallel these schemes, but for the AEAD-KI setting. Bellare and Hoang [9] propose a spectrum of new definitions for commitment that capture not just committing to the key but also committing to the nonce, associated data, and plaintext message. They also propose new key-committing AEAD schemes based on AES-GCM and AES-GCM-SIV as well as a construction that transforms a legacy AEAD scheme into one that is key-committing using what they call a *committing PRF*. This construction is similar to that proposed by Albertini et al. (and indeed they note that Albertini et al.’s construction can be viewed as a specific instantiation of their scheme). Their work considers only the single-key setting, but their schemes can be used as FROB AEAD schemes in our AEAD-KI constructions.

Degabriele et al. [14] propose nonce-set AEAD, which is similar to our AEAD-KI formalism but instead for nonce sets. Their formalism considers decryption accepting a set of nonces and then returning the correct nonce along with the plaintext message.

Chan and Rogaway [13] formalize anonymous authenticated encryption, which requires that ciphertexts maintain strong privacy when considering nonces and associated data. They mention the need for robustness, although they only consider robustness for honestly generated keys, which is implied by the typical AEAD security notion as shown in [17]. Finally, Jaeger and Tyagi [22] consider multi-user simulation-based security definitions for various standard symmetric definitions where keys can be adaptively compromised.

## 2 Preliminaries

We follow the notational conventions used in [19]. We fix some alphabet  $\Sigma$ , e.g.  $\Sigma = \{0, 1\}$ . For any  $x \in \Sigma^*$ , let  $|x|$  denote its length. We write  $x \leftarrow_s X$  to denote uniformly sampling from a finite set  $X$ . We write  $X||Y$  to denote concatenation of two strings. For a string  $X$  of  $n$  bits, we will write  $X[i, \dots, j]$  for  $1 \leq i < j \leq n$  to mean the substring of  $X$  beginning at index  $i$  and ending at index  $j$ . For notational simplicity, we assume that one can unambiguously parse  $Z = X||Y$  into its two parts, even for strings of varying length. For strings  $X, Y \in \{0, 1\}^*$  we write  $X \oplus Y$  to denote taking the XOR of  $X[1, \dots, \min(|X|, |Y|)] \oplus Y[1, \dots, \min(|X|, |Y|)]$ . For some table of values where  $y_i$  is stored at key  $x_i$ , denoted as  $\mathbb{T}[x_i] \leftarrow y_i$ , for a set of keys  $X = \{x_i : 1 \leq i \leq \kappa\}$  for some integer  $\kappa$ ,  $\mathbb{T}[X]$  denotes the set  $\{\mathbb{T}[x_i] : x_i \in X\}$ . We denote a vector of elements as  $[\cdot]$ . We denote the value stored at index  $i$  in vector  $\mathbb{K}$  as  $\mathbb{K}[i]$ . For vector  $\mathbb{K}$ , we denote  $x \in \mathbb{K}$  as  $x$  is an element of  $\mathbb{K}$  and  $|\mathbb{K}|$  as the number of elements in  $\mathbb{K}$ . We also denote  $\mathbb{K}.\text{add}(x)$  to mean adding  $x$  to the end of vector  $\mathbb{K}$ . We denote  $[n]_\ell$  as the  $\ell$ -bit representation of the integer  $n$ .

We use code-based games [11] to formalize security notions. Variables’ types should be clear from context and are modeled as random variables in the probability distribution defined by the random coins used in execution.  $\Pr[G \Rightarrow y]$  denotes (over the random coins of  $G$ ) that the game  $G$  outputs the value  $y$ . For a scheme  $S$ , we will sometimes use “a  $G_S$  adversary” to describe an adversary in the game  $G$  instantiated with the scheme  $S$ . For an adversary  $\mathcal{A}$ ,  $G_S^{\mathcal{A}}$  denotes the game  $G$  instantiated with the scheme  $S$  and specific

<p><u>MU-nAE1<sub>AEAD</sub><sup>A</sup>:</u>  <math>i \leftarrow 0</math>  <math>b' \leftarrow \mathcal{A}^{\text{GENKEY, ENC, DEC}}</math>  Return <math>b'</math></p> <p><u>GENKEY():</u>  <math>K \leftarrow \text{AEAD.Kg}()</math>  <math>T[i] \leftarrow K; i \leftarrow i + 1</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>C \leftarrow \text{AEAD.Enc}(T[\text{id}], N, AD, M)</math>  <math>c[N, AD, C] \leftarrow \text{id}</math>  Return <math>C</math></p> <p><u>DEC(id, N, AD, C):</u>  If <math>c[N, AD, C] = \text{id}</math> then return <math>\perp</math>  <math>M \leftarrow \text{AEAD.Dec}(T[\text{id}], N, AD, C)</math>  Return <math>M</math></p>	<p><u>MU-nAE0<sub>AEAD</sub><sup>A</sup>:</u>  <math>i \leftarrow 0</math>  <math>b' \leftarrow \mathcal{A}^{\text{GENKEY, ENC, DEC}}</math>  Return <math>b'</math></p> <p><u>GENKEY():</u>  <math>i \leftarrow i + 1</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>C \leftarrow \{0, 1\}^{\text{len}(M)}</math>  Return <math>C</math></p> <p><u>DEC(id, N, AD, C):</u>  Return <math>\perp</math></p>
---	---

Fig. 2: Games MU-nAE1 and MU-nAE0 are used for MU-nAE<sub>AEAD</sub>, or multi-user real-or-random security, for scheme AEAD = (Kg, Enc, Dec).

adversary  $\mathcal{A}$ .  $\Pr [G_S^{\mathcal{A}} \Rightarrow \text{out}]$  denotes the probability that game G instantiated with scheme  $S$  and adversary  $\mathcal{A}$  outputs out.

**Authenticated encryption.** An authenticated encryption with associated data (AEAD) scheme AEAD = (Kg, Enc, Dec) consists of a triple of algorithms. Associated to any scheme AEAD is a key space  $\mathcal{K} \subseteq \Sigma^*$ , nonce space  $\mathcal{N} \subseteq \Sigma^*$ , header space  $\mathcal{A} \subseteq \Sigma^*$ , message space  $\mathcal{M} \subseteq \Sigma^*$ , and ciphertext space  $\mathcal{C} \subseteq \Sigma^*$ . The randomized key generation algorithm Kg outputs a secret key  $K \in \mathcal{K}$ . Encryption Enc is deterministic and takes as input a 4-tuple  $(K, N, AD, M) \in (\Sigma^*)^4$  and outputs ciphertext  $C$  or a distinguished error symbol  $\perp$ . We require that  $\text{Enc}(K, N, AD, M) \neq \perp$  if  $(K, N, AD, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ . Decryption Dec is deterministic and takes as input a tuple  $(K, N, AD, C) \in (\Sigma^*)^4$  and outputs value  $M$  or  $\perp$ . An AEAD scheme is correct if for any  $(K, N, AD, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$  it holds that  $\text{Dec}(K, N, AD, \text{Enc}(K, N, AD, M)) = M$ .

The security notion we consider for AEAD schemes is nonce-based real-or-random security under chosen-ciphertext attack [29]. We generalize this to a multi-user setting [12] where an adversary can interact with multiple instances of the AEAD scheme, which we call MU-nAE<sub>AEAD</sub>. The game pseudocode is presented in Figure 2. We restrict our attention to nonce-respecting adversaries, meaning they never query the same nonce twice to ENC for the same key identifier id, and they only query a key identifier  $\text{id} < i$  to ENC and DEC. The MU-nAE<sub>AEAD</sub> advantage of an adversary  $\mathcal{A}$  is defined as

$$\mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{A}) = |\Pr [\text{MU-nAE1}_{\text{AEAD}}^{\mathcal{A}} \Rightarrow 1] - \Pr [\text{MU-nAE0}_{\text{AEAD}}^{\mathcal{A}} \Rightarrow 1]|.$$

**Full robustness.** We use the full robustness notion for AEAD schemes from Farshim et al. [17], but adapted to the nonce-based setting. Albertini et al. [2] also provide a FROB notion for nonce-based AEAD, with slightly different syntax — our formulation is equivalent to theirs. Roughly, FROB security tasks an adversary with providing two keys and a ciphertext such that both keys successfully decrypt the ciphertext. We define the game in Figure 3.

The FROB<sub>AEAD</sub> advantage of an adversary  $\mathcal{A}$  is defined as

$$\mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{A}) = \Pr [\text{FROB}_{\text{AEAD}}^{\mathcal{A}} \Rightarrow 1].$$

**Pseudo-random functions.** We use a multi-user variant of the traditional pseudo-random function (PRF) definition (q.v., [7]) for two functions, where the adversary is given access to oracles for both functions. We define the games in Figure 3. The game gives the adversary an additional GENKEY oracle that allows it to generate multiple keys. The MU-PRF<sub>F</sub> advantage of an adversary  $\mathcal{A}$  is defined as

$$\mathbf{Adv}_{\text{F}_0, \text{F}_1}^{\text{mu-prf}}(\mathcal{A}) = |\Pr [\text{REAL}_{\text{F}_0, \text{F}_1}^{\mathcal{A}} \Rightarrow 1] - \Pr [\text{IDEAL}^{\mathcal{A}} \Rightarrow 1]|.$$

<p><u>FROB<sub>AEAD</sub><sup>A</sup>:</u>  <math>K_0, K_1, N, AD, C \leftarrow \mathcal{A}</math>          If <math>K_0 = K_1</math>:            Return 0  <math>M_0 \leftarrow \text{Dec}(K_0, N, AD, C)</math>  <math>M_1 \leftarrow \text{Dec}(K_1, N, AD, C)</math>          Return <math>M_0 \neq \perp \wedge M_1 \neq \perp</math></p>	<p><u>REAL<sub>F<sub>0</sub>, F<sub>1</sub></sub><sup>A</sup>:</u>  <math>i \leftarrow 0</math>  <math>b \leftarrow \mathcal{A}^{\text{GENKEY}, \text{Func0}, \text{Func1}}</math>          Return <math>b'</math></p> <p><u>GENKEY():</u>  <math>K \leftarrow \mathcal{S} \{0, 1\}^\kappa</math>  <math>T[i] \leftarrow K; i \leftarrow i + 1</math></p> <p><u>Func0(id, x):</u>          Return <math>F_0(T[id], x)</math></p> <p><u>Func1(id, x):</u>          Return <math>F_1(T[id], x)</math></p>	<p><u>IDEAL<sup>A</sup>:</u>  <math>i \leftarrow 0; T_0, T_1 \leftarrow []</math>  <math>b' \leftarrow \mathcal{A}^{\text{GENKEY}, \text{Func0}, \text{Func1}}</math>          Return <math>b'</math></p> <p><u>GENKEY():</u>  <math>i \leftarrow i + 1</math></p> <p><u>Func0(id, x):</u>          If <math>T_0[id, x] = \perp</math>: <math>T_0[id, x] \leftarrow \mathcal{S} \{0, 1\}^n</math>          Return <math>T_0[id, x]</math></p> <p><u>Func1(id, x):</u>          If <math>T_1[id, x] = \perp</math>: <math>T_1[id, x] \leftarrow \mathcal{S} \{0, 1\}^n</math>          Return <math>T_1[id, x]</math></p>
--	---	--

Fig. 3: (Left) Game FROB<sub>AEAD</sub> is the full robustness security notion for scheme AEAD = (Kg, Enc, Dec). (Center/Right) Games REAL and IDEAL are used for MU-PRF<sub>F<sub>0</sub>, F<sub>1</sub></sub>, or multi-user PRF security for functions  $F_0 = \{0, 1\}^\kappa \times \{0, 1\}^{\ell_0} \rightarrow \{0, 1\}^{n_0}$  and  $F_1 = \{0, 1\}^\kappa \times \{0, 1\}^{\ell_0} \rightarrow \{0, 1\}^{n_0}$ .

**Collision resistance.** The collision resistance (CR) game for function  $F = \{0, 1\}^\kappa \times \{0, 1\}^\ell \rightarrow \{0, 1\}^n$  measures the ability of an adversary to find two key-value pairs such that the evaluation of  $F$  on these inputs evaluates to the same output. More formally, the  $\text{CR}_F$  advantage of an adversary  $\mathcal{A}$  is defined as

$$\text{Adv}_F^{\text{CR}}(\mathcal{A}) = \Pr [ K_0, x_0, K_1, x_1 \leftarrow \mathcal{A} : (K_0, x_0) \neq (K_1, x_1) \wedge F(K_0, x_0) = F(K_1, x_1) ].$$

**Pre-image resistance.** The pre-image resistance game for function  $F = \{0, 1\}^\kappa \times \{0, 1\}^\ell \rightarrow \{0, 1\}^n$  measures the ability of an adversary to find the pre-image of a random range point for  $F$ . More formally, the  $\text{PRE}_F$  advantage of an adversary  $\mathcal{A}$  is defined as

$$\text{Adv}_F^{\text{PRE}}(\mathcal{A}) = \Pr [ y \leftarrow \mathcal{S} \{0, 1\}^n; K, x \leftarrow \mathcal{A}(y) : F(K, x) = y ].$$

### 3 Defining AEAD with Key Identification

We start by formalizing the notion of AEAD with key identification (AEAD-KI). AEAD-KI extends AEAD schemes to the setting where a recipient stores multiple keys and must therefore choose which key to use for decryption. At a high level, our formalization extends prior ones on AEAD in the following ways:

- We add a notion of key labels, which are potentially public, application-defined strings associated to secret keys. For notational simplicity, we will redefine a key to be a label, secret key pair.
- Decryption takes as input a vector of keys, instead of a single key. Decryption must determine both which key to use, and the corresponding plaintext. We model the keys used by decryption as a vector, instead of a set, to preserve information about the order.
- Ciphertexts may include a key identification tag, to assist decryption. We will explore a variety of ways to construct key identification tags, each with different security and performance profiles.

These changes to our conceptualization of syntax and semantics of AEAD necessarily require revisiting security as well. Later in this section we propose a new simulator-based all-in-one security definition that captures confidentiality and integrity for AEAD-KI schemes. Furthermore, we will see that the shift to AEAD-KI introduces a number of subtleties related to ciphertexts potentially being decryptable under more than one key. We will therefore provide new key robustness (also called key commitment) notions for AEAD-KI.

**Syntax and semantics.** An AEAD-KI scheme is a triple of algorithms combined with a key space  $\mathcal{K} \subseteq \Sigma^*$ , key label space  $\mathcal{L} \subseteq \Sigma^*$ , nonce space  $\mathcal{N} \subseteq \Sigma^*$ , associated data space  $\mathcal{A} \subseteq \Sigma^*$ , message space  $\mathcal{M} \subseteq \Sigma^*$ , ciphertext space  $\mathcal{C} \subseteq \Sigma^*$ , and key tag space  $\mathcal{T} \subseteq \Sigma^*$ . We will often leave the spaces implicit and clear from context. We write that an AEAD-KI scheme  $\text{AEKI} = (\text{Kg}, \text{Enc}, \text{Dec})$  consists of the algorithms:

- $K \leftarrow_s \text{AEKI.Kg}(\text{kid})$   
The randomized key generation algorithm takes as input the key label  $\text{kid}$  to use for the generated secret key. The key label operates as metadata for the secret key. Key generation outputs a key  $K \in \mathcal{L} \times \mathcal{K}$ , which is a pair composed of the key label and the secret encryption/decryption key. While the encryption key must be kept secret, the key label can be public.
- $(T_k, C) \leftarrow \text{AEKI.Enc}(K, N, AD, M)$   
The nonce-based deterministic encryption algorithm takes as input tuple  $(K, N, AD, M) \in (\Sigma^*)^4$  and outputs pair  $(T_k, C) \in \mathcal{T} \times \mathcal{C}$  or a distinguished error symbol  $\perp$ . Notice that encryption returns, in addition to the encrypted plaintext, a bit string  $T_k$ , which can be empty. We will refer to this as the key tag, as we discuss more below. Both the key tag and the encrypted plaintext form the ciphertext. We require that if  $(K, N, AD, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$  then  $\text{Enc}(K, N, AD, M) \neq \perp$ .
- $(K, M) \leftarrow \text{AEKI.Dec}(\mathbb{K}, N, AD, T_k, C)$   
The decryption algorithm's input is  $(\mathbb{K}, N, AD, T_k, C) \in \mathcal{K}^* \times (\Sigma^*)^4$  and its output is  $(K, M) \in \mathcal{K} \times \mathcal{M}$  or  $\perp$ . Decryption is deterministic. Notice that instead of a single key, decryption takes as input a vector of keys  $\mathbb{K}$ , and we denote vectors of length one or greater by  $\mathcal{K}^*$ . Furthermore, in addition to the plaintext, decryption returns the corresponding key that produced the plaintext. If no key can decrypt, then the error symbol  $\perp$  is returned.

**Correctness.** When extending AEAD to allow for multiple keys, a meaningful definition of correctness becomes more complex. We expect that when encrypting with a key  $K$  to produce ciphertext  $C$ , and then decrypting  $C$  with any vector  $\mathbb{K}$  that includes  $K$ , the original plaintext should be recovered. However, this cannot for practical schemes be guaranteed absolutely, since there may exist another key that successfully decrypts the ciphertext. Indeed, the correct outcome that we expect of a scheme that allows decryption to accept multiple keys now becomes more like a security property of key robustness, where we have computational guarantees that decryption succeeds for a single key. We therefore provide here a simpler, absolute correctness definition and later focus on capturing the behavior we want from AEAD-KI through key robustness, which we cover below.

**Definition 1.** *An AEAD-KI scheme is correct if the following hold:*

- (1) *For any  $(K, N, AD, M)$  it holds that  $\Pr[(K', M') = (K, M)] = 1$  where*

$$(K', M') \leftarrow \text{Dec}([K], N, AD, \text{Enc}(K, N, AD, M))$$

*and the probability is over the coins used by encryption;*

- (2) *For any  $(\mathbb{K}, N, AD, T_k, C)$  and  $(K, M) \leftarrow \text{Dec}(\mathbb{K}, N, AD, T_k, C)$  it must be that either  $(K, M) = \perp$  or  $K \in \mathbb{K}$ ; and*
- (3) *For any  $\mathbb{K}, \mathbb{K}'$  and any  $(N, AD, T_k, C)$ , let*

$$(K, M) \leftarrow \text{Dec}(\mathbb{K}, N, AD, T_k, C), (K', M') \leftarrow \text{Dec}(\mathbb{K}', N, AD, T_k, C).$$

*If  $(K, M) \neq \perp$  and  $K \in \mathbb{K}'$ , then  $(K', M') \neq \perp$ .*

The first condition lifts traditional perfect correctness for AEAD to the syntax of AEAD-KI for decryption with a single key. The second condition additionally asks that decryption only ever output a key that was in the key vector. The third correctness condition roughly requires that if  $\text{Dec}$  outputs some key  $K$  for a key set  $\mathbb{K}$ , any other  $\mathbb{K}'$  containing  $K$  must decrypt to non- $\perp$ . (Note we do not require decryption with  $\mathbb{K}'$  outputs  $K$ ; this property is guaranteed by our key robustness security notion, which we discuss next.) For all schemes we consider, correctness is easily established via inspection of their decryption algorithm; we therefore will omit explicit analysis.

### 3.1 Key robustness

As mentioned, in the AEAD-KI setting, key robustness is partly about correctness: we expect that the key used to encrypt a plaintext should be the only one to correctly decrypt the resulting ciphertext. However, when decryption allows for multiple keys—some of which may be adversarially-chosen—this property cannot

<p style="margin: 0;">KI-FROB<sub>AEKI</sub><sup>A</sup>:</p> <p style="margin: 0;"><math>\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C \leftarrow_{\\$} \mathcal{A}</math></p> <p style="margin: 0;"><math>(K_0, M_0) \leftarrow \text{Dec}(\mathbb{K}_0, N, AD, T_k, C)</math></p> <p style="margin: 0;"><math>(K_1, M_1) \leftarrow \text{Dec}(\mathbb{K}_1, N, AD, T_k, C)</math></p> <p style="margin: 0;"><math>(\text{kid}_0, K_0^*) \leftarrow K_0; (\text{kid}_1, K_1^*) \leftarrow K_1</math></p> <p style="margin: 0;">If <math>(K_0, M_0) \neq \perp \wedge (K_1, M_1) \neq \perp \wedge K_0^* \neq K_1^*</math>:</p> <p style="margin: 0; padding-left: 20px;">Return 1</p> <p style="margin: 0;">Return 0</p>
---

Fig. 4: Game KI-FROB<sub>AEKI</sub> represents full robustness for AEAD-KI scheme AEKI = (Kg, Enc, Dec).

be satisfied without some form of key robustness. Briefly, key robustness guarantees that only a single key can be used to decrypt a given ciphertext.

Farshim et al. [16] first defined several key robustness notions for AEAD schemes. Their strongest notion is called full robustness (FROB), and requires an adversary to discover two keys that each successfully decrypt an adversarially chosen nonce, associated data, and ciphertext (see Section 2). Bellare and Hoang [9] recently introduced even stronger notions, such as their CMT3, which require commitment to not just the key but also nonces and associated data. For simplicity we stick with adapting the FROB notion. Analogous adaptations can be made to lift CMT3 to the key identification setting, but some schemes would require modification to meet them (e.g., including nonce and associated data in key check value computations).

We define KI-FROB security for an AEAD-KI scheme via the game shown in Figure 10. It requires an adversary to find two key vectors and a nonce, associated data, and ciphertext. Decryption is run with each key vector, and the adversary wins should both decryptions succeed and the returned secret key, message pairs are distinct. Note here that we are focused on the secret key, not key including key label, thereby explicitly excluding as a win having distinct key labels. This is to allow schemes that use multiple key labels for the same key. The KI-FROB<sub>AEKI</sub> advantage of an adversary  $\mathcal{A}$  is defined as

$$\text{Adv}_{\text{AEKI}}^{\text{ki-frob}}(\mathcal{A}) = \Pr [\text{KI-FROB}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1].$$

Furthermore, note an important property of our KI-FROB security notion: any scheme meeting it is in some sense agnostic to the *ordering* of keys in the key vector input to decryption. If two different orderings of the same key vector caused different keys to be output, these two key vectors would give a KI-FROB win. (Note that correctness condition (3) above implies that two different orderings of the same key vector must both either output  $\perp$  or both non- $\perp$ .) Looking ahead, it also means our KI-nAE definition need not account for distinguishing attacks caused by the order of keys in the key vectors; they are ruled out for any KI-FROB scheme.

Finally, we provide an alternative key robustness definition unique to this setting which we call *extended full robustness* (XFROB) in Appendix A. While we show that this notion is stronger than KI-FROB for AEAD-KI schemes that only meet correctness conditions (1) and (2), for any correct AEAD-KI schemes that specifically meet condition (3) the two definitions are equivalent.

### 3.2 All-in-one confidentiality and integrity

Just as for an AEAD scheme, we expect an AEAD-KI scheme to maintain confidentiality and integrity. Towards formalizing what this means, one starting point is existing indistinguishability style security definitions for AEAD (e.g., [29]). However, an important modeling question is how to handle key vectors during decryption. This suggests we should start instead with a multi-user style AEAD security notion [12] which allows the adversary to request generation of many keys and obtain encryption of plaintexts under keys of their choice.

To model key anonymity, we may additionally require that adversaries not be able to distinguish between encryptions under different honest keys. To model chosen-ciphertext attacks and, in particular, ciphertext integrity, we face additional choices about how much control to give adversaries over key vectors during decryption. One option would be to allow adversaries to choose the key vector but only allow honestly generated keys to be added to the vector. Unfortunately, this would not capture attacks in which a malicious key is somehow inserted into a recipient's key vector, a scenario that arises in practice. For instance, Albertini

<p><u>KI-nAE1<sub>AEKI</sub><sup>A</sup>:</u>  <math>j \leftarrow 0</math>  <math>b \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{GENHONESTKEY, ENC, DEC}}</math>  Return <math>b</math></p> <p><u>GENHONESTKEY(kid):</u>  <math>K \leftarrow_{\mathcal{S}} \text{AEKI.Kg}(\text{kid})</math>  <math>T[j] \leftarrow K ; j \leftarrow j + 1</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>(T_k, C) \leftarrow \text{AEKI.Enc}(T[\text{id}], N, AD, M)</math>  Return <math>(T_k, C)</math></p> <p><u>DEC(<math>\mathbb{K}</math>, N, AD, <math>T_k</math>, C):</u>  <math>\mathbb{K}^* \leftarrow []</math>  For <math>(\text{honest}, \text{data}) \in \mathbb{K}</math>:    If <math>\text{honest} = \text{true}</math>: <math>\mathbb{K}^*.\text{add}(T[\text{data}])</math>    Else: <math>\mathbb{K}^*.\text{add}(\text{data})</math>  <math>(K, M) \leftarrow \text{AEKI.Dec}(\mathbb{K}^*, N, AD, T_k, C)</math>  Return <math>M</math></p>	<p><u>KI-nAE0<sub>AEKI, S, L<sub>Enc</sub></sub><sup>A</sup>:</u>  <math>\sigma_s \leftarrow_{\mathcal{S}} \text{S.Init}()</math>  <math>b \leftarrow_{\mathcal{S}} \mathcal{A}^{\text{GENHONESTKEY, ENC, DEC}}</math>  Return <math>b</math></p> <p><u>GENHONESTKEY(kid):</u>  <math>\sigma_s \leftarrow_{\mathcal{S}} \text{S.Kg}(\text{kid}, \sigma_s)</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>\mathcal{L} \leftarrow_{\mathcal{L}_{\text{Enc}}}(\text{id}, M)</math>  <math>(T_k, C, \sigma_s) \leftarrow_{\mathcal{S}} \text{S.Enc}(N, AD, \mathcal{L}, \sigma_s)</math>  <math>\mathcal{C}[\text{id}, N, AD, T_k, C] \leftarrow M</math>  Return <math>(T_k, C)</math></p> <p><u>DEC(<math>\mathbb{K}</math>, N, AD, <math>T_k</math>, C):</u>  For <math>(\text{honest}, \text{data}) \in \mathbb{K}</math>:    If <math>\text{honest} = \text{true} \wedge \mathcal{C}[\text{data}, N, AD, T_k, C] \neq \perp</math>:      Return <math>\mathcal{C}[\text{data}, N, AD, T_k, C]</math>  If <math>\exists (\text{honest}, \text{data}) \in \mathbb{K}</math> s.t. <math>\text{honest} = \text{false}</math>:    <math>(M, \sigma_s) \leftarrow_{\mathcal{S}} \text{S.Dec}(\mathbb{K}, N, AD, T_k, C, \sigma_s)</math>    Return <math>M</math>  Return <math>\perp</math></p>
---	---

Fig. 5: Game KI-nAE<sub>AEKI</sub> is the all-in-one security notion for AEAD-KI schemes.

et al. [2] describe a vulnerability arising from such a scenario in the context of key rotation within key management services.

We therefore opt for a stronger notion of security for which adversaries can insert malicious keys into key vectors. This renders more complex how the security game should handle decryption oracles, because we need to somehow demarcate between decryption queries that correspond to ciphertext forgeries and ones that should not: an adversary can always generate ciphertexts that decrypt under an adversary-chosen key.

To handle these subtleties, we use a simulation-based approach and an all-in-one confidentiality and ciphertext integrity notion. The pseudocode games for the resulting KI-nAE<sub>AEKI</sub> security notion are shown in Figure 5. KI-nAE<sub>AEKI</sub> is parameterized by the simulator, a stateful tuple of algorithms  $\mathcal{S} = (\text{Init}, \text{Kg}, \text{Enc}, \text{Dec})$ . The adversary is given access to an honest key generation oracle GENHONESTKEY, an encryption oracle ENC, and a decryption oracle DEC. DEC accepts as input a key vector  $\mathbb{K}$  as well as ciphertext tuple  $(N, AD, T_k, C)$ . The key vector  $\mathbb{K}$  is composed of tuples  $(\text{honest}, \text{data})$ , where  $\text{honest} = \text{true}$  indicates that the key is honestly generated and  $\text{data}$  is the game-generated key identifier for the key. If  $\text{honest} = \text{false}$ , then the key is malicious and  $\text{data}$  is itself the key.

The game KI-nAE1 models interactions with the real scheme, and therefore calls the relevant AEKI algorithm to answer oracle queries. The ideal game KI-nAE0 instead uses simulator  $\mathcal{S}$  to generate the oracle outputs for the adversary. Encryption provides leakage to the simulator. The encryption leakage algorithm  $\mathcal{L}_{\text{Enc}}(\text{id}, M)$  takes as input the game-generated key identifier and the plaintext message and outputs the encryption leakage. We specify two concrete leakage functions,  $\mathcal{L}_{\text{Enc}}^{\text{id}}$  and  $\mathcal{L}_{\text{Enc}}^{\text{anon}}$ . The non-key anonymous algorithm  $\mathcal{L}_{\text{Enc}}^{\text{id}}$  returns as leakage both the game-generated key identifier and the size of the plaintext, while  $\mathcal{L}_{\text{Enc}}^{\text{anon}}$  only returns the size of the plaintext. In the latter case, this of course means that the simulator will have no knowledge of which honest key the adversary chose for encryption.

The decryption oracle in KI-nAE0 works in two parts. In the first part, the table of previous ENC outputs is scanned for each honest key. If the queried ciphertext is in the table, the message is returned. Otherwise, if there are malicious keys in  $\mathbb{K}$ , the simulator is given its state, the ciphertext tuple, and  $\mathbb{K}$ .

It may not be immediately obvious why this DEC is the “right” one. The main advantage in defining DEC as we have is that for honest keys this decryption oracle ensures our definition implies a variant of ciphertext integrity for AEAD-KI: no matter the simulator’s behavior, crafting a new valid ciphertext for an honest key automatically gives a distinguisher between real and ideal games. We discuss this in more detail and formally prove that this definition implies ciphertext integrity in Appendix B.

<u>KL.Kg(kid):</u> $K \leftarrow \mathcal{K}$ Return (kid, $K$ )	<u>KL.Enc(<math>K, N, AD, M</math>):</u> $(\text{kid}, K^*) \leftarrow K$ $C \leftarrow \text{AEAD.Enc}(K^*, N, AD \parallel \text{kid}, M)$ Return (kid, $C$ )	<u>KL.Dec(<math>\mathbb{K}, N, AD, T_k, C</math>):</u> For (kid, $K$ ) $\in \mathbb{K}$ : If kid = $T_k$ : $M \leftarrow \text{AEAD.Dec}(K, N, AD \parallel \text{kid}, C)$ If $M \neq \perp$ : Return ( $K, M$ ) Return $\perp$
--	--	---

Fig. 6: A typical key label scheme KL which is parameterized by an AEAD scheme.

We have two versions of KI-nAE, one key anonymous and one not. For a simulator  $\mathcal{S}$ , the more general (non-anonymous)  $\text{KI-nAE}_{\text{AEKI}}$  advantage of an adversary  $\mathcal{A}$  with respect to  $\mathcal{S}$  is defined as

$$\text{Adv}_{\text{AEKI}, \mathcal{S}}^{\text{ki-nae}}(\mathcal{A}) = \left| \Pr \left[ \text{KI-nAE1}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[ \text{KI-nAE0}_{\text{AEKI}, \mathcal{S}, L_{\text{Enc}}^{\text{id}}}^{\mathcal{A}} \Rightarrow 1 \right] \right|.$$

Meanwhile, the key anonymous  $\text{KI-nAE-KA}_{\text{AEKI}}$  advantage of an adversary  $\mathcal{A}$  with respect to  $\mathcal{S}$  is defined as

$$\text{Adv}_{\text{AEKI}, \mathcal{S}}^{\text{ki-nae-anon}}(\mathcal{A}) = \left| \Pr \left[ \text{KI-nAE1}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[ \text{KI-nAE0}_{\text{AEKI}, \mathcal{S}, L_{\text{Enc}}^{\text{anon}}}^{\mathcal{A}} \Rightarrow 1 \right] \right|.$$

We further note that while we chose to base our definitions on those for nonce-based AEAD from [12, 29], we can also adapt these definitions to other settings, such as those for randomized symmetric encryption and the nonce-hiding framework for the AE2 definitions proposed by Bellare et al. [10].

**Malicious keys and ciphertexts.** As discussed before, we opt to capture the adversary providing both malicious keys and malicious ciphertexts since this models real-world settings. However, we must explicitly give the simulator all malicious keys because the adversary knows these keys as well and could otherwise trivially distinguish. This makes it difficult for KI-nAE to capture key robustness notions like KI-FROB. We therefore opt for separate key robustness notions, mirroring the AEAD setting. We believe it is possible to give an all-in-one definition that also implies KI-FROB; we leave this difficult, but interesting, modelling question to future work.

## 4 Key Labels

One simple technique for key identification is assigning to each key a static label and then prepending the key's label to every ciphertext it produces. In practice, labels typically are randomly-generated strings, URLs indicating where to fetch the key, or even user identifiers. This approach is widely used by both key management services (KMS), such as the Amazon Web Services (AWS) KMS [5], Microsoft Azure Key Vault [25], and Oracle Key Vault [26]; as well as cryptography libraries, such as Google's Tink [32]. Relatedly, the popular cryptographic library Libsodium [24] also recommends using a key label as a way to add robustness to AEAD schemes.

While key labels appear efficient and straightforward, they have not been formally analyzed. For instance, key labels are often used with non-FROB AEAD schemes, such as AES-GCM and ChaCha20-Poly1305. We will see in our analysis that key labels do not automatically produce a KI-FROB AEAD-KI scheme.

We generalize the key labels construction as  $\text{KL}[\text{AEAD}]$  and provide the pseudocode in Figure 6. The scheme is parameterized by an AEAD scheme, used for encryption and decryption. To simplify notation, we also refer to the scheme as KL when the specific AEAD scheme used can be arbitrary or is obvious from context. Key generation allows the caller to specify the key label kid, which is then stored as part of the key. We model the label as input to generalize label creation out-of-band and to enable adversarial inputs. Meanwhile, encryption simply uses kid as the key tag. Decryption iterates through  $\mathbb{K}$  to find the first key with an identifier matching the key tag that successfully decrypts the ciphertext.

**Analyzing robustness.** Utilizing a key label at first seems like a trivial and practical way to add robustness to any AEAD-KI scheme. A ciphertext with some key identifier can only be decrypted by the corresponding key. However, this method fails when multiple keys can have the same label and a non-FROB AEAD scheme is chosen. Notice, for instance, that KL does not enforce uniqueness of labels. An adversary could choose

two AEAD keys so that they have the same label, compute an AEAD key multi-collision ciphertext [23] for these keys, and attach the label to the ciphertext as the key tag. Decryption will then proceed successfully for both keys because their labels match the key tag of the ciphertext.

In the following theorem, we show that an FROB AEAD is both sufficient and necessary for KL to be KI-FROB.

**Theorem 1.** *Let  $\mathcal{A}$  be a KI-FROB adversary for scheme  $\text{KL}[\text{AEAD}]$ . Then we give FROB adversary  $\mathcal{B}$  for AEAD such that*

$$\mathbf{Adv}_{\text{KL}[\text{AEAD}]}^{\text{ki-frob}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B}).$$

*Furthermore, let  $\mathcal{C}$  be an FROB adversary for AEAD. Then we give a KI-FROB adversary  $\mathcal{D}$  for  $\text{KL}[\text{AEAD}]$  such that*

$$\mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{C}) \leq \mathbf{Adv}_{\text{KL}[\text{AEAD}]}^{\text{ki-frob}}(\mathcal{D}).$$

$\mathcal{B}$  runs in time that of  $\mathcal{A}$  and  $\mathcal{D}$  runs in time that of  $\mathcal{C}$ .

*Proof.* We first construct FROB adversary  $\mathcal{B}$  against AEAD as follows.  $\mathcal{B}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ . Let  $\text{Dec}(\mathbb{K}_0, N, AD, T_k, C)$  and  $\text{Dec}(\mathbb{K}_1, N, AD, T_k, C)$  return  $(K_0, M_0)$  and  $(K_1, M_1)$ , respectively, where  $K_0 = (\text{kid}_0, K_0^*)$  and  $K_1 = (\text{kid}_1, K_1^*)$ . Note that  $\mathcal{A}$  can only win the FROB game if  $(K_0, M_0), (K_1, M_1) \neq \perp$  and  $K_0^* \neq K_1^*$ , so let this be the case.  $\mathcal{B}$  can return  $K_0^*, K_1^*, N, AD, C$ , where again both keys successfully decrypt  $N, AD, C$ . Thus,  $\mathcal{B}$  wins game FROB for AEAD when  $\mathcal{A}$  wins game KI-FROB for  $\text{KL}[\text{AEAD}]$ .

We now construct KI-FROB adversary  $\mathcal{D}$  against KL as follows.  $\mathcal{D}$  runs  $\mathcal{C}$ , which returns  $K_0, K_1, N, AD, C$ . Then  $\mathcal{D}$  creates KL key vectors  $\mathbb{K}_0 \leftarrow [(0^n, K_0)]$  and  $\mathbb{K}_1 \leftarrow [(0^n, K_1)]$  and finally returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, 0^n, C$ . Since both  $K_0, K_1$  have the same identifier, which matches the given key tag, the scheme KL will try both keys when decrypting and successfully decrypt. Thus, we have that  $\mathcal{D}$  wins game KI-FROB for  $\text{KL}[\text{AEAD}]$  when  $\mathcal{C}$  wins game FROB for AEAD.  $\square$

**Analyzing KI-nAE.** We now show that  $\text{KL}[\text{AEAD}]$  is KI-nAE secure for leakage algorithm  $\mathbf{L}_{\text{Enc}}^{\text{id}}$  when AEAD is both FROB and MU-nAE secure. Notably, encryption cannot be key anonymous: the label is static across calls to ENC and the simulator can only simulate this by knowing which key was queried for encryption. We provide the theorem statement and proof sketch below. The full proof is provided in Appendix C.

**Theorem 2.** *Let  $\mathcal{A}$  be a KI-nAE adversary making at most  $q$  queries to its oracles and querying at most  $m$  malicious keys. Then we give KI-nAE simulator  $\mathcal{S}$  and adversaries  $\mathcal{B}, \mathcal{C}$  such that*

$$\mathbf{Adv}_{\text{KL}[\text{AEAD}], \mathcal{S}}^{\text{ki-nae}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B}) + \mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{C}) + mq/|\mathcal{K}|.$$

*Adversaries  $\mathcal{B}, \mathcal{C}$  run in time that of  $\mathcal{A}$  with an  $\mathcal{O}(q)$  overhead and simulator  $\mathcal{S}$  runs in time  $\mathcal{O}(mq)$ .*

*Proof sketch:* The KI-nAE simulator simulates ENC queries by keeping track of the label for each key and returning it as the key tag with a random string of the correct length for the encrypted plaintext. For DEC queries, the simulator iterates through the list of key data given in the key vector  $\mathbb{K}$  and checks for malicious keys, for which it is directly given the secret key and can decrypt the ciphertext for any that has a key label matching the key tag. If there are no malicious keys or none that correctly decrypt, then the simulator returns  $\perp$ .

We bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which DEC keeps iterating through  $\mathbb{K}$  if  $\mathcal{A}$  provides a malicious key in  $\mathbb{K}$  that was honestly generated by a call to GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish between these games by  $mq/|\mathcal{K}|$ , since there are at most  $m$  malicious keys and at most  $q$  honest keys and  $\mathcal{A}$  can at best guess one of the honest keys. We next transition to a game in which DEC skips any malicious key that can decrypt a ciphertext output from a call to ENC. We bound the ability of  $\mathcal{A}$  to distinguish between these games by the FROB security of AEAD. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext through ENC. We bound the ability of the adversary to distinguish between these games by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

**Using unique random identifiers.** Thus far, we have relied on an FROB AEAD scheme, due to the fact that KL allows duplicate key identifiers, which follows practice (e.g., the Tink library). One might instead suggest somehow enforcing uniqueness of key labels at the “application layer” and using a non-FROB AEAD. We argue below that this approach either fails to meet natural security goals or greatly increases the complexity of the application layer; thus, FROB AEADs are superior as the “base” AEAD for KL-like constructions.

To study this question, we first need to express application-level enforcement of unique key labels in our AEKI formalism. A simple way to do this is to have Dec check the identifiers for all keys in  $\mathbb{K}$  and output  $\perp$  if two different keys have the same identifier. This approach does not meet condition (3) of our correctness notion above. Any key vector with repeated labels will cause decryption to fail, even if the key vector contains the correct key.

More subtly, this approach also fails to provide KI-FROB for the AEKI scheme. To see why, note that because decryption is stateless, uniqueness cannot be checked across different invocations. Thus, if the underlying AEAD is not FROB, an adversary can choose two keys, produce a key multi-collision ciphertext, assign both keys the same label that matches the ciphertext’s key tag, and then put the keys in separate key vectors. Decryption will succeed for both key vectors, even though their keys have non-unique labels.

Preventing this attack and providing KI-FROB for the AEKI scheme requires stateful decryption: namely, the application must track all key identifier-secret key pairs seen across *all* decryption operations. This seems difficult to implement correctly and efficiently, and is certain to increase the complexity of the application. Thus, we believe it is better to use FROB AEAD to cryptographically guarantee KI-FROB security of AEKI.

**Analyzing trial decryption.** A special case of the key labels scheme is brute-force trial decryption, which we refer to as TD[AEAD]. This scheme simply assigns the empty string  $\varepsilon$  as the key label for all keys, meaning there are no key identifiers, and decryption must always trial decrypt for all keys in the key vector. The Tink library, for instance, allows keys to also have “raw” labels, which indicate they have no identifier.

Notably, the multi-user Shadowsocks protocol we describe in the introduction falls into this category. The attack described by Len et al. [23] is made possible by the fact that the Shadowsocks protocol uses a non-FROB AEAD scheme. The benefit of our KI-FROB definition is that it demonstrates that such a scheme is insecure when not using a FROB AEAD scheme.

While trial decryption still requires the use of an FROB AEAD scheme, it is key anonymous. In the below theorem, we show that trial decryption meets our stronger key anonymous encryption leakage model. We provide the full proof in Appendix D.

**Theorem 3.** *Let  $\mathcal{A}$  be a KI-nAE-KA adversary making at most  $q$  queries to its oracles and querying at most  $m$  malicious keys. Then we give KI-nAE-KA simulator  $\mathcal{S}$  and adversaries  $\mathcal{B}, \mathcal{C}$  such that*

$$\mathbf{Adv}_{\text{TD}[\text{AEAD}], \mathcal{S}}^{\text{ki-nae-anon}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B}) + \mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{C}) + mq/|\mathcal{K}|.$$

*Adversaries  $\mathcal{B}, \mathcal{C}$  run in time that of  $\mathcal{A}$  with an  $\mathcal{O}(q)$  overhead and simulator  $\mathcal{S}$  runs in time  $\mathcal{O}(mq)$ .*

*Proof sketch:* The KI-nAE simulator simulates ENC queries by returning a random string of the correct length for the encrypted plaintext. For DEC queries, the simulator iterates through the list of key data given in the key vector  $\mathbb{K}$  and checks for malicious keys, for which it is directly given the secret key and can decrypt itself. If there are no malicious keys or none that correctly decrypt, then the simulator returns  $\perp$ .

We bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which DEC keeps iterating through  $\mathbb{K}$  if  $\mathcal{A}$  provides a malicious key in  $\mathbb{K}$  that was honestly generated by a call to GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish between these games by  $mq/|\mathcal{K}|$ , since there are at most  $m$  malicious keys and at most  $q$  honest keys and  $\mathcal{A}$  can at best guess one of the honest keys. We next transition to a game in which DEC skips any malicious key that can decrypt a ciphertext output from a call to ENC. We bound the ability of  $\mathcal{A}$  to distinguish between these games by the FROB security of AEAD. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext through ENC. We bound the ability of the adversary to distinguish between these games by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

<u>KCV.Kg(kid):</u> $K \leftarrow \mathcal{K}$ Return (kid, $K$ )	<u>KCV.Enc(<math>K, N, AD, M</math>):</u> $(\text{kid}, K^*) \leftarrow K$ $\text{kcv} \leftarrow F_{\text{kcv}}(K^*)$ ; $K_e \leftarrow \text{KDF}(K^*)$ $T_k \leftarrow \text{kid} \parallel \text{kcv}$ $C \leftarrow \text{AEAD.Enc}(K_e, N, AD \parallel T_k, M)$ Return ( $T_k, C$ )	<u>KCV.Dec(<math>\mathbb{K}, N, AD, T_k, C</math>):</u> For $K \in \mathbb{K}$ : $(\text{kid}^*, K^*) \leftarrow K$ $\text{kcv}^* \leftarrow F_{\text{kcv}}(K^*)$ ; $K_e \leftarrow \text{KDF}(K^*)$ If $\text{kid}^* \parallel \text{kcv}^* = T_k$ : $M \leftarrow \text{AEAD.Dec}(K_e, N, AD \parallel T_k, C)$ If $M \neq \perp$ : Return ( $K, M$ ) Return $\perp$
---	---	---

Fig. 7: A key check value scheme KCV parameterized by encryption scheme AEAD = (Kg, Enc, Dec) with associated key space  $\mathcal{K} = \{0, 1\}^k$ , key check value function  $F_{\text{kcv}} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ , and encryption key derivation function  $\text{KDF} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ .

While key labels are a simple and commonly used approach in practice for identifying keys, here we have shown that a formal analysis surfaces subtleties. In particular, key labels do not provide key anonymity, unless all keys have the same label as in the less efficient trial decryption-based scheme. Moreover, a key label approach must rely on the underlying AEAD scheme being FROB. In the next section, we explore a different tactic called static key identifiers, which computes the identifier from the key itself.

## 5 Static Key Identifiers

In this section we describe a class of AEAD-KI that uses what we call *static key identifiers*. This technique computes a static identifier from the key that along with the key label is used as the key tag. In practice, this static identifier is often referred to as a *key check value*, also known as a key checksum value. We formalize this approach as the AEAD-KI scheme  $\text{KCV}[\text{AEAD}, F_{\text{kcv}}, \text{KDF}]$ , shown in Figure 7. The scheme has key space  $\{0, 1\}^k$  and is parameterized by AEAD scheme  $\text{AEAD} = (\text{Kg}, \text{Enc}, \text{Dec})$  with associated key space  $\mathcal{K} = \{0, 1\}^k$ , key check value function  $F_{\text{kcv}} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ , and encryption key derivation function  $\text{KDF} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ . For simplicity, we will sometimes refer to the scheme as KCV when the parameters are obvious.

Key generation generates a secret key and attaches the input label  $\text{kid}$  as part of the AEAD-KI key. For static identifier schemes, the key tag is composed of both the key label and the key check value. This models what happens in practice, as schemes may use the key label as a way to locate a key or set of keys and then use the key check value as a commitment or integrity check of the key. For instance, AWS uses the Amazon Resource Name (ARN) as a URL for looking up keys, while an extra commitment string verifies this is the correct key [4].

Encryption derives the key check value using the function  $F_{\text{kcv}}$  and separately derives the AEAD secret key using the function  $\text{KDF}$ . Encryption adds the key tag to the authentication scope of AEAD by appending it to the authenticated data. Decryption iterates through the key vector to compute each key’s identifier using  $F_{\text{kcv}}$  and find the first key with one matching  $T_k$ . If AEAD decryption with this key succeeds, then the corresponding decrypted plaintext is returned.

Key check values have been widely used in practice [2, 4, 18, 27, 31] to derive a value from the key, typically using a hash function or block cipher, that can then be used to confirm the integrity of or identify the key during decryption. These static key identifiers can be used in two ways: as static key hints or as static key commitments. Static key hints use a non-CR function to derive an identifier from the key, meaning that key hints are not unique to a single key. Because they are not used to commit to a key, they can be short and efficiently computed, while still enabling AEAD-KI schemes to narrow down the scope of keys to check during decryption. For instance, one common technique is taking 24 bits from the AES evaluation of the key over some fixed string. However, key hints must be used with an FROB AEAD scheme in practice to guarantee robust key identification.

Conversely, static key commitments do commit to the key and therefore do not need to be used with an FROB AEAD scheme. In practice, this means key commitments must employ a CR key check value function, which can be more computationally intensive and require longer key tags. For instance, Albertini et al. [2] suggest a variant of this method as their “generic solution” for adding key commitment to AEAD schemes. Their Type I and II schemes in particular feature a static identifier by computing two SHA256

Application	Key Check Value	KH or KC?
AWS Encryption SDK [4]	SHA256( $K \parallel 0x436f6d6d69740102$ )	KC
GlobalPlatform [18]	$\text{msb}_{24}(\text{AES}_K((\mathbb{1}_8)^{16}))$	KH
Telegram [31]	$\text{lsb}_{64}(\text{SHA1}(K))$	KH
PKCS#11 [27]	$\text{msb}_{24}(\text{AES}_K(0^{128}))$	KH

Fig. 8: Key check value functions used in practice, labeled as either key hints (KH) or key commitments (KC).

hashes over the key and using these values for the key identifier and AEAD encryption key. This scheme has been adopted as the default method of key identification by AWS [4]. We summarize some sample schemes used for both key hints and key commitments in Figure 8. Later in this section we will show detailed security results for both types of static key identifiers.

**Using a key derivation function.** Whenever a fixed value is computed from the key as a key tag and then composed with an AEAD scheme that uses the same key, there is the potential that the AEAD scheme uses the same value for its internal computation. Since the key tag is sent in the clear, this could lead to confidentiality or integrity vulnerabilities. Indeed, Iwata and Wang [21] have shown that this does happen in practice. They describe forgery attacks for several variants of CBC-MAC proposed by ISO/IEC 9797-1:2011 [20] when used with the key check value suggested by ANSI X9.24-1:2009 [3].

To simplify the analysis of composing  $F_{\text{kcv}}$  with an AEAD scheme, we also use key derivation function KDF to derive an independent AEAD encryption key. In this section, we will show that KDF will need to be a CR PRF. For similar reasons, Albertini et al. also use a CR PRF to derive a separate AEAD key. This of course results in extra overhead and could be unnecessary if the key identifier is never used in the internal computation of AEAD. One could analyze specific key check value functions and AEAD schemes that can be used together without the need for a separate key derivation function; we leave this as an open problem.

## 5.1 Static Key Hint

Static key hints are a sub-class of static key identifiers which compute the key check value using a non-collision-resistant PRF. This means they can be short and efficiently computed, e.g. using a universal hash or by truncating the output of a hash function or block cipher. While key hints cannot be used to commit to a key, they can be used to narrow down the search space of the given key vector during decryption. However, in order to ensure that key robustness is achieved, these key hints must rely on using an FROB AEAD scheme, as we show below.

**Analyzing robustness.** Here we show that static key hints rely on using an FROB AEAD scheme as well as collision-resistant function KDF to achieve KI-FROB.

**Theorem 4.** *Let  $\mathcal{A}$  be an KI-FROB adversary for scheme  $\text{KCV}[\text{AEAD}, F_{\text{kcv}}, \text{KDF}]$ . Then we give adversaries  $\mathcal{B}, \mathcal{C}$  such that*

$$\text{Adv}_{\text{KCV}[\text{AEAD}, F_{\text{kcv}}, \text{KDF}]}^{\text{ki-frob}}(\mathcal{A}) \leq \text{Adv}_{\text{KDF}}^{\text{cr}}(\mathcal{B}) + \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{C}).$$

*Adversaries  $\mathcal{B}, \mathcal{C}$  run in time that of  $\mathcal{A}$ .*

*Proof.* We prove the theorem using a sequence of game hops. Let game  $G_0$  be the game FROB with the call to the decryption algorithm Dec replaced by the pseudocode for  $\text{KCV}[\text{AEAD}, F_{\text{kcv}}, \text{KDF}].\text{Dec}()$ . Next we transition to game  $G_1$ , which is identical to  $G_0$  except that when KDF is called to derive the encryption key from the keys in  $\mathbb{K}_1$ , it checks if there was some other different key  $K^*$  prior to this call that output to the same encryption key derived from the keys in  $\mathbb{K}_0$ . If this happens, then  $G_1$  will return 0.

We can upper bound the difference in advantage of  $\mathcal{A}$  in  $G_0$  and  $G_1$  by the probability that KDF finds a collision. We then provide the CR adversary  $\mathcal{B}$  such that its advantage in the CR game for KDF upper bounds this probability.  $\mathcal{B}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ .  $\mathcal{B}$  then checks if there is some key  $K_0 \in \mathbb{K}_0$  such that  $(\text{kid}_0, K_0^*) \leftarrow K_0$  and  $K_e \leftarrow \text{KDF}(K_0^*)$  and some key  $K_1 \in \mathbb{K}_1$  such that  $(\text{kid}_1, K_1^*) \leftarrow K_1$  and  $K_e \leftarrow \text{KDF}(K_1^*)$ , and returns  $K_0^*, K_1^*$ . Notice that whenever KDF finds the collision in  $G_1$ , then  $\mathcal{B}$  wins the CR game for KDF.

Finally, we upper bound the advantage of  $\mathcal{A}$  in game  $G_1$  by the advantage of the following FROB adversary  $\mathcal{C}$  against AEAD.  $\mathcal{C}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ . Let  $\text{Dec}(\mathbb{K}_0, N, AD, T_k, C)$  and  $\text{Dec}(\mathbb{K}_1, N, AD, T_k, C)$  return  $(K_0, M_0)$  and  $(K_1, M_1)$ , respectively. Also let  $(\text{kid}_0, K_0^*) \leftarrow K_0$  and  $(\text{kid}_1, K_1^*) \leftarrow K_1$ . Note that  $\mathcal{A}$  can only win game  $G_1$  if  $(K_0, M_0) \neq \perp$  and  $(K_1, M_1) \neq \perp$  and  $K_0^* \neq K_1^*$ , so let this be the case.  $\mathcal{C}$  can return  $K_0^*, K_1^*, N, AD, C$ , where both keys successfully decrypt  $N, AD, C$ . Thus,  $\mathcal{C}$  wins FROB for AEAD when  $\mathcal{A}$  wins  $G_1$  for KCV.  $\square$

**Analyzing KI-nAE.** We now show that  $\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$  is KI-nAE secure for leakage algorithm  $\text{L}_{\text{Enc}}^{\text{id}}$  when AEAD is both FROB and MU-nAE secure,  $\text{F}_{\text{kcv}}$  and KDF are multi-user PRFs, and KDF is pre-image resistant. Notably, this means that encryption is not key anonymous. We provide the theorem statement and proof sketch below. We show the full proof in Appendix E.

**Theorem 5.** *Let  $\mathcal{A}$  be a KI-nAE adversary making at most  $q$  queries to its oracles, of which  $q_k$  are to GENHONESTKEY, and querying at most  $m$  malicious keys. Then we give KI-nAE simulator  $\mathcal{S}$  and adversaries  $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$  such that*

$$\text{Adv}_{\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}], \mathcal{S}}^{\text{ki-nae}}(\mathcal{A}) \leq \text{Adv}_{\text{F}_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B}) + q_k \cdot \text{Adv}_{\text{KDF}}^{\text{pre}}(\mathcal{C}) + \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{D}) + \text{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{E}) + \frac{q_k^2}{2^{\kappa+1}}.$$

$\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$  run in time that of  $\mathcal{A}$  with  $\mathcal{O}(q)$  overhead and  $\mathcal{S}$  runs in time  $\mathcal{O}(mq)$ .

*Proof sketch:* The KI-nAE simulator simulates ENC queries by keeping track of the label for each key and generating a random  $n$ -bit string as the key check value  $\text{kcv}$ . It then returns the label appended with the key check value as the key tag along with a random string of the correct length for the encrypted plaintext. For DEC queries, the simulator iterates through the list of key data given in the key vector  $\mathbb{K}$  and checks for malicious keys, for which it is directly given the secret key and decrypts itself for any that have a matching key tag. If there are no malicious keys or none that correctly decrypt, then the simulator returns  $\perp$ .

We bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which calls to  $\text{F}_{\text{kcv}}$  and KDF for honest keys are replaced with calls to random functions. We bound the ability of  $\mathcal{A}$  to distinguish these games by the MU-PRF security of  $\text{F}_{\text{kcv}}$  and KDF. We next transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if for KDF they are the pre-image of some honestly generated AEAD encryption key  $K_e$  computed in GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish these games by the pre-image resistance security of KDF, multiplied by a factor of  $q_k$ . Then we transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if they can decrypt some honestly generated ciphertext output by GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish between these games by the FROB security of AEAD. We then transition to a game in which we eliminate collisions when key  $K$  is chosen at random from the key space  $\mathcal{K}$ , for which we use the birthday bound  $q_k^2/2^{\kappa+1}$  to bound. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext from a call to ENC. We bound the distinguishing advantage by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

## 5.2 Static Key Commitment

Static key commitments are the second subclass of static key identifiers. They compute the key check value using using a collision-resistant PRF. While this means they must be longer and less efficient than key hints, they can commit to the key, and thus can be used with non-FROB AEADs, as we now prove.

**Theorem 6.** *Let  $\mathcal{A}$  be a KI-FROB adversary for  $\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$ . Then we give CR adversary  $\mathcal{B}$ , running in time that of  $\mathcal{A}$ , for  $\text{F}_{\text{kcv}}$  such that*

$$\text{Adv}_{\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]}^{\text{ki-frob}}(\mathcal{A}) \leq \text{Adv}_{\text{F}_{\text{kcv}}}^{\text{cr}}(\mathcal{B}).$$

<p><b>nKCV.Kg(kid):</b></p> <p><math>K \leftarrow \mathcal{K}</math></p> <p>Return <math>(\varepsilon, K)</math></p>	<p><b>nKCV.Enc(<math>K, N, AD, M</math>):</b></p> <p><math>(\varepsilon, K^*) \leftarrow K; (N_0, N_1) \leftarrow N</math></p> <p><math>\text{kcv} \leftarrow \text{F}_{\text{kcv}}(K^*, N_0); K_e \leftarrow \text{KDF}(K^*)</math></p> <p><math>T_k \leftarrow \text{kcv}</math></p> <p><math>C \leftarrow \text{AEAD.Enc}(K_e, N_1, AD \  T_k, M)</math></p> <p>Return <math>(T_k, C)</math></p>	<p><b>nKCV.Dec(<math>\mathbb{K}, N, AD, T_k, C</math>):</b></p> <p><math>(N_0, N_1) \leftarrow N</math></p> <p>For <math>K \in \mathbb{K}</math>:</p> <p><math>(\varepsilon, K^*) \leftarrow K</math></p> <p><math>\text{kcv}^* \leftarrow \text{F}_{\text{kcv}}(K^*, N_0); K_e \leftarrow \text{KDF}(K^*)</math></p> <p>If <math>\text{kcv}^* = T_k</math>:</p> <p><math>M \leftarrow \text{AEAD.Dec}(K_e, N_1, AD \  T_k, C)</math></p> <p>If <math>M \neq \perp</math>: Return <math>(K, M)</math></p> <p>Return <math>\perp</math></p>
--	---	--

Fig. 9: A nonce-based key check value scheme nKCV parameterized by AEAD = (Kg, Enc, Dec) with key space  $\mathcal{K} = \{0, 1\}^k$ ; key check value function  $\text{F}_{\text{kcv}} : \{0, 1\}^k \times \{0, 1\}^r \rightarrow \{0, 1\}^n$ , and encryption key derivation function  $\text{KDF} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ .

*Proof.* We construct adversary  $\mathcal{B}$  as follows. It runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ . Let the values returned by decryption of the ciphertext for each key vector be  $(K_0, M_0), (K_1, M_1)$ . Also let  $(\text{kid}_0, K_0^*) \leftarrow K_0$  and  $(\text{kid}_1, K_1^*) \leftarrow K_1$ . We know that  $(K_0, M_0), (K_1, M_1) \neq \perp$  and  $K_0^* \neq K_1^*$  for  $\mathcal{A}$  to win. This also means that  $\text{F}_{\text{kcv}}(K_0^*) = \text{F}_{\text{kcv}}(K_1^*)$ .  $\mathcal{B}$  can then return  $K_0, K_1$  as a collision for  $\text{F}_{\text{kcv}}$ .  $\square$

**Analyzing KI-nAE.** While Albertini et al. do not explicitly prove security for this scheme, they claim that it meets their real-or-random AE security definition. However, any encryption scheme that attaches a fixed string to a ciphertext trivially cannot meet this definition. The benefit of our KI-nAE definition is that it captures security for non-key anonymous schemes. Indeed, our result shows that if  $\text{F}_{\text{kcv}}$  is a pre-image resistant multi-user PRF, KDF is a multi-user PRF, and AEAD is MU-nAE-secure, then KCV is KI-nAE-secure, for leakage  $\text{L}_{\text{Enc}}^{\text{id}}$ . We provide the theorem statement and proof sketch here; the full proof is in Appendix E.

**Theorem 7.** *Let  $\mathcal{A}$  be a KI-nAE adversary making at most  $q$  queries to its oracles, of which  $q_k$  are to GENHONESTKEY, and querying at most  $m$  malicious keys. Then we give adversaries  $\mathcal{B}, \mathcal{C}, \mathcal{D}$  such that*

$$\text{Adv}_{\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}], \mathcal{S}}^{\text{ki-nae}}(\mathcal{A}) \leq \text{Adv}_{\text{F}_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B}) + q_k \cdot \text{Adv}_{\text{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C}) + \text{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{D}) + \frac{q_k^2}{2^{\kappa+1}}.$$

$\mathcal{B}, \mathcal{C}, \mathcal{D}$  run in time that of  $\mathcal{A}$  with a  $\mathcal{O}(q)$  overhead and  $\mathcal{S}$  runs in time  $\mathcal{O}(mq)$ .

*Proof sketch:* The proof uses the same KI-nAE simulator as that for Theorem 5. We again bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which calls to  $\text{F}_{\text{kcv}}$  and KDF for honest keys are replaced with calls to random functions. We bound the ability of  $\mathcal{A}$  to distinguish these games by the MU-PRF security of  $\text{F}_{\text{kcv}}$  and KDF. We next transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if for  $\text{F}_{\text{kcv}}$  they are the pre-image of some honestly generated key check value  $\text{kcv}$  computed in GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish these games by the pre-image resistance security of  $\text{F}_{\text{kcv}}$ , multiplied by a factor of  $q_k$ . We then transition to a game in which we eliminate collisions when key  $K$  is chosen at random from the key space  $\mathcal{K}$ , for which we use the birthday bound  $q_k^2/2^{\kappa+1}$  to bound. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext from a call to ENC. We bound the distinguishing advantage by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

While static key identifiers are versatile in that they can be used either as key hints or key commitments, they unfortunately do not provide key anonymity. Next, we will see how dynamic identifiers enable anonymous key identification.

## 6 Dynamic Key Identifiers

In this section we describe the key anonymous counterpart to static key identifiers, which we call *dynamic key identifiers*. A dynamic identifier is computed from the secret key during encryption using part of the

input nonce. Unlike the static identifier approach, this scheme cannot use key labels as part of the key tag because key labels are fixed for a key and would therefore break key anonymity. We formalize dynamic key identifiers as an AEAD-KI scheme with  $\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$ , shown in Figure 9. The scheme has key space  $\{0, 1\}^\kappa$  and is parameterized by encryption scheme  $\text{AEAD} = (\text{Kg}, \text{Enc}, \text{Dec})$  with associated key space  $\mathcal{K} = \{0, 1\}^k$ , key check value function  $\text{F}_{\text{kcv}} : \{0, 1\}^\kappa \times \{0, 1\}^r \rightarrow \{0, 1\}^n$ , and encryption key derivation function  $\text{KDF} : \{0, 1\}^\kappa \rightarrow \{0, 1\}^k$ . For simplicity, we may refer to the scheme as  $\text{nKCV}$  when the parameters are obvious.

Encryption now takes in a nonce for which one part is used to derive the key check value and the other is used in the AEAD computation. These nonces do not need to be distinct. Encryption derives the key check value and AEAD key from the secret key using the functions  $\text{F}_{\text{kcv}}$  and  $\text{KDF}$ , respectively. The key check value  $\text{kcv}$  is computed on part of the nonce so that it changes for each encryption call. Meanwhile, the AEAD key is computed on just the secret key, meaning the AEAD key is fixed for each  $\text{nKCV}$  secret key. Unlike for  $\text{KCV}$ , here the key check value by itself forms the key tag. Encryption adds the key tag to the authentication scope of AEAD by appending it to the authenticated data. Decryption iterates through the key vector to compute each key's identifier using  $\text{F}_{\text{kcv}}$  and find the first key with one matching  $T_k$ .

## 6.1 Dynamic Key Hint

Dynamic key hints are a subclass of dynamic key identifiers which compute the key check value using using a non-collision-resistant PRF. Similar to static key hints, they can be short and more efficiently computed than key commitments. They are useful for narrowing down the search space of the given key vector during decryption. However, in order to ensure that key robustness is achieved, these key hints must rely on using an FROB AEAD scheme, as we show below.

**Analyzing robustness.** Here we show that dynamic key hints rely on using a CR function  $\text{KDF}$  and an FROB AEAD scheme to achieve KI-FROB.

**Theorem 8.** *Let  $\mathcal{A}$  be a KI-FROB adversary for scheme  $\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$ . Then we give adversaries  $\mathcal{B}, \mathcal{C}$  running in time that of  $\mathcal{A}$  such that*

$$\text{Adv}_{\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]}^{\text{ki-frob}}(\mathcal{A}) \leq \text{Adv}_{\text{KDF}}^{\text{cr}}(\mathcal{B}) + \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{C}).$$

*Proof.* We prove the theorem using a sequence of game hops. Let game  $G_0$  be the game KI-FROB with the call to the decryption algorithm  $\text{Dec}$  replaced by the pseudocode for  $\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}].\text{Dec}()$ . Next we transition to game  $G_1$ , which is identical to  $G_0$  except that when  $\text{KDF}$  is called to derive the encryption key from the keys in  $\mathbb{K}_1$ , it checks if there was some other different key  $K^*$  prior to this call that output to the same encryption key derived from the keys in  $\mathbb{K}_0$ . If this happens, then  $G_1$  will return 0.

We can upper bound the difference in advantage of  $\mathcal{A}$  in  $G_0$  and  $G_1$  by the probability that  $\text{KDF}$  finds a collision. We then provide the CR adversary  $\mathcal{B}$  such that its advantage in the CR game for  $\text{KDF}$  upper bounds this probability.  $\mathcal{B}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ .  $\mathcal{B}$  then checks if there is some key  $K_0 \in \mathbb{K}_0$  such that  $(\text{kid}_0, K_0^*) \leftarrow K_0$  and  $K_e \leftarrow \text{KDF}(K_0^*)$  and some key  $K_1 \in \mathbb{K}_1$  such that  $(\text{kid}_1, K_1^*) \leftarrow K_1$  and  $K_e \leftarrow \text{KDF}(K_1^*)$ , and returns  $K_0^*, K_1^*$ . Notice that whenever  $\text{KDF}$  finds the collision in  $G_1$ , then  $\mathcal{B}$  wins the CR game for  $\text{KDF}$ .

Finally, we upper bound the advantage of  $\mathcal{A}$  in game  $G_1$  by the advantage of the following FROB adversary  $\mathcal{C}$  against AEAD.  $\mathcal{C}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ . Let  $\text{Dec}(\mathbb{K}_0, N, AD, T_k, C)$  and  $\text{Dec}(\mathbb{K}_1, N, AD, T_k, C)$  return  $(K_0, M_0)$  and  $(K_1, M_1)$ , respectively. Also let  $(\varepsilon, K_0^*) \leftarrow K_0$ ,  $(\varepsilon, K_1^*) \leftarrow K_1$ , and  $N_0 \| N_1 \leftarrow N$ . Note that  $\mathcal{A}$  can only win game  $G_1$  if  $(K_0, M_0) \neq \perp$  and  $(K_1, M_1) \neq \perp$  and  $K_0^* \neq K_1^*$ , so let this be the case.  $\mathcal{C}$  can return  $K_0^*, K_1^*, N_1, AD, C$ , where again both keys decrypt  $N_1, AD, C$ . Thus,  $\mathcal{C}$  wins FROB for AEAD when  $\mathcal{A}$  wins  $G_1$  for  $\text{nKCV}$ .  $\square$

**Analyzing KI-nAE.** We now show that  $\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$  is KI-nAE-KA secure for leakage algorithm  $\text{L}_{\text{Enc}}^{\text{anon}}$  when AEAD is both FROB and MU-nAE secure,  $\text{KDF}$  is a pre-image resistant multi-user PRF, and  $\text{F}_{\text{kcv}}$  is a multi-user PRF. Notably, this means that encryption is key anonymous. We assume that the adversary  $\mathcal{A}$  is a nonce-respecting adversary that never queries the same  $N_0$  or  $N_1$  to  $\text{ENC}$ . We provide the theorem statement and proof sketch below. The proof is provided in Appendix F.

**Theorem 9.** *Let  $\mathcal{A}$  be a KI-nAE adversary making at most  $q$  queries to its oracles, of which  $q_k$  are to GENHONESTKEY, and querying at most  $m$  malicious keys. Then we give adversaries  $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$  such that*

$$\mathbf{Adv}_{\mathbf{nKCV}[\mathbf{AEAD}, \mathbf{F}_{\mathbf{kcv}}, \mathbf{KDF}], \mathbf{S}}^{\mathbf{ki-nae-anon}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathbf{F}_{\mathbf{kcv}}, \mathbf{KDF}}^{\mathbf{mu-prf}}(\mathcal{B}) + q_k \cdot \mathbf{Adv}_{\mathbf{KDF}}^{\mathbf{pre}}(\mathcal{C}) + \mathbf{Adv}_{\mathbf{AEAD}}^{\mathbf{frob}}(\mathcal{D}) + \mathbf{Adv}_{\mathbf{AEAD}}^{\mathbf{mu-nae}}(\mathcal{E}) + \frac{q_k^2}{2^{\kappa+1}}.$$

$\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$  run in time that of  $\mathcal{A}$  with a  $\mathcal{O}(q)$  overhead and  $\mathbf{S}$  runs in time  $\mathcal{O}(mq)$ .

*Proof sketch:* The KI-nAE simulator can simulate ENC queries by generating a random  $\mathbf{n}$ -bit string as the key check value  $\mathbf{kcv}$ . It can then return the key check value as the key tag along with a random string of the correct length for the encrypted plaintext. For DEC queries, the simulator can iterate through the list of key data given in the key vector  $\mathbb{K}$  and check for malicious keys, for which it is directly given the secret key and can decrypt itself for any that have a matching key tag. If there are no malicious keys or none that correctly decrypt, then the simulator returns  $\perp$ .

We bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which calls to  $\mathbf{F}_{\mathbf{kcv}}$  and  $\mathbf{KDF}$  for honest keys are replaced with calls to random functions. We bound the ability of  $\mathcal{A}$  to distinguish these games by the MU-PRF security of  $\mathbf{F}_{\mathbf{kcv}}$  and  $\mathbf{KDF}$ . We next transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if for  $\mathbf{KDF}$  they are the pre-image of some honestly generated AEAD encryption key  $K_e$  computed in GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish these games by the pre-image resistance security of  $\mathbf{KDF}$ , multiplied by a factor of  $q_k$ . Then we transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if they can decrypt some honestly generated ciphertext output by GENHONESTKEY. We bound the ability of  $\mathcal{A}$  to distinguish between these games by the FROB security of AEAD. We then transition to a game in which we eliminate collisions when key  $K$  is chosen at random from the key space  $\mathcal{K}$ , for which we use the birthday bound  $q_k^2/2^{\kappa+1}$  to bound. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext from a call to ENC. We bound the distinguishing advantage by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

## 6.2 Dynamic Key Commitment

Dynamic key commitments are the second subclass of dynamic key identifiers; they instead compute the key check value using using a collision-resistant PRF. While this means they must be longer and less efficient than their key hint counterpart, they can be used to commit to the key. This also means that they can be used with non-FROB AEAD schemes. Dynamic key commitments parallel the Type III generic construction from Albertini et al. [2]. They are also similar to the UtC transform proposed by Bellare and Hoang [9], although this scheme is only considered in the traditional single-key setting. The UtC transform uses a committing PRF that takes as input a key and nonce and outputs pair  $(P, L)$  such that  $P$  is a string that commits to the key.  $L$  is then used as the encryption key. We note that a committing PRF can be used in place of  $\mathbf{KDF}$  and  $\mathbf{F}_{\mathbf{kcv}}$ , although our formalism allows for analyzing the security requirements for deriving the key tag separately from deriving the key.

Furthermore, the NonceWrap scheme proposed by Chan and Rogaway [13] can be considered a type of dynamic key commitment scheme. Their scheme encrypts the ciphertext as

$$C = \text{AES}(K_1, N \| 0^3 2) \| \text{AES-GCM}(K_2, N, AD, M),$$

where the first string is a 128-bit “header”. During decryption, the correct key is found from a set of possible keys by re-computing the header and verifying the 32-bit all-zeros string remains intact. Interestingly, this scheme may be considered a key commitment scheme when the nonce must be specified along with the ciphertext, as in our formalization of KI-FROB. However, if, as the setting in this work intends, the nonce does not need to be specified, then this scheme does not meet KI-FROB and a key-committing AEAD should be used instead.

**Analyzing robustness.** Here we show that dynamic key commitments rely only on the collision-resistance of the function  $\mathbf{F}_{\mathbf{kcv}}$  to achieve KI-FROB. This means that AEAD does not in fact have to be FROB.

**Theorem 10.** *Let  $\mathcal{A}$  be a KI-FROB adversary for  $\text{KCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$ . Then we give CR adversary  $\mathcal{B}$ , running in time that of  $\mathcal{A}$ , for  $\text{F}_{\text{kcv}}$  such that*

$$\text{Adv}_{\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]}^{\text{ki-frob}}(\mathcal{A}) \leq \text{Adv}_{\text{F}_{\text{kcv}}}^{\text{cr}}(\mathcal{B}).$$

$\mathcal{B}$  runs in time that of  $\mathcal{A}$ .

*Proof.* We construct CR adversary  $\mathcal{B}$  against  $\text{F}_{\text{kcv}}$  as follows.  $\mathcal{B}$  runs  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ . Let  $\text{Dec}(\mathbb{K}_0, N, AD, T_k, C)$  and  $\text{Dec}(\mathbb{K}_1, N, AD, T_k, C)$  return  $(K_0, M_0)$  and  $(K_1, M_1)$ , respectively. Also let  $(\varepsilon, K_0^*) \leftarrow K_0$ ,  $(\varepsilon, K_1^*) \leftarrow K_1$ , and  $N_0 \| N_1 \leftarrow N$ . We know that  $(K_0, M_0) \neq \perp$  and  $(K_1, M_1) \neq \perp$  and  $K_0^* \neq K_1^*$  for  $\mathcal{A}$  to win.  $\mathcal{B}$  can then return  $(K_0^*, N_0), (K_1^*, N_0)$  as a collision for  $\text{F}_{\text{kcv}}$  since  $\text{F}_{\text{kcv}}(K_0^*, N_0) = \text{F}_{\text{kcv}}(K_1^*, N_0) = T_k$ . We therefore have that  $\mathcal{B}$  wins game CR for  $\text{F}_{\text{kcv}}$  when  $\mathcal{A}$  wins game KI-FROB for  $\text{nKCV}$ .  $\square$

**Analyzing KI-nAE.** We now show that  $\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}]$  is KI-nAE-KA secure for leakage algorithm  $\text{L}_{\text{Enc}}^{\text{anon}}$  when AEAD MU-nAE secure, KDF is a multi-user PRF, and  $\text{F}_{\text{kcv}}$  is a CR multi-user PRF. Again, this means that encryption is key anonymous. We assume that the adversary  $\mathcal{A}$  is a nonce-respecting adversary that never queries the same  $N_0$  or  $N_1$  across queries to ENC. We provide the theorem statement below; the full proof is provided in Appendix F.

**Theorem 11.** *Using  $\text{L}_{\text{Enc}}^{\text{anon}}$ , let  $\mathcal{A}$  be a KI-nAE adversary making at most  $q$  queries to its oracles, of which  $q_k$  are to GENHONESTKEY and  $q_e$  are to ENC, and querying at most  $m$  malicious keys. Then we give adversaries  $\mathcal{B}, \mathcal{C}, \mathcal{D}$  such that*

$$\text{Adv}_{\text{nKCV}[\text{AEAD}, \text{F}_{\text{kcv}}, \text{KDF}], \mathcal{S}}^{\text{ki-nae-anon}}(\mathcal{A}) \leq \text{Adv}_{\text{F}_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B}) + q_e \cdot \text{Adv}_{\text{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C}) + \text{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{D}) + \frac{q_k^2}{2^{\kappa+1}}.$$

$\mathcal{B}, \mathcal{C}, \mathcal{D}$  run in time that of  $\mathcal{A}$  with a  $\mathcal{O}(q)$  overhead and  $\mathcal{S}$  runs in time  $\mathcal{O}(mq)$ .

*Proof sketch:* The proof uses the same KI-nAE simulator as that for Theorem 9. We again bound the advantage of  $\mathcal{A}$  with a sequence of game hops. We first transition to a game in which calls to  $\text{F}_{\text{kcv}}$  and KDF for honest keys are replaced with calls to random functions. We bound the ability of  $\mathcal{A}$  to distinguish these games by the MU-PRF security of  $\text{F}_{\text{kcv}}$  and KDF. We next transition to a game in which malicious keys in  $\mathbb{K}$  queried to DEC are skipped if for  $\text{F}_{\text{kcv}}$  they are the pre-image of some honestly generated key check value  $\text{kcv}$  computed in ENC. We bound the ability of  $\mathcal{A}$  to distinguish these games by the pre-image resistance security of  $\text{F}_{\text{kcv}}$ , multiplied by a factor of  $q_e$ . We then transition to a game in which we eliminate collisions when key  $K$  is chosen at random from the key space  $\mathcal{K}$ , for which we use the birthday bound  $q_k^2/2^{\kappa+1}$  to bound. Finally, we transition to a game in which ENC generates a random string as the encrypted plaintext and DEC skips honest keys in the key vector if they were not used to produce the queried ciphertext from a call to ENC. We bound the distinguishing advantage by the MU-nAE security of AEAD. Since this last game guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext, iterating through  $\mathbb{K}$  in order in this game is identical to iterating through the honest keys first and then the malicious keys, proving our claim.  $\square$

## Acknowledgments

The authors thank Mihir Bellare for suggesting an improved correctness notion and various improvements in security definitions, as well as other helpful feedback on an early draft of the paper. The authors also thank Ian Miers and Nirvan Tyagi for their help in the early stages of this project. Finally, the authors are grateful to the anonymous reviewers of Asiacrypt 2022 for their feedback and suggestions. This work was supported in part by NSF grant CNS-2120651 and the NSF Graduate Research Fellowship under Grant No. DGE-2139899.

## References

1. Abdalla, M., Bellare, M., Neven, G.: Robust encryption. In: Theory of Cryptography Conference. pp. 480–497. Springer (2010)

2. Albertini, A., Duong, T., Gueron, S., Kölbl, S., Luykx, A., Schmiege, S.: How to abuse and fix authenticated encryption without key commitment. In: USENIX Security (2022)
3. ANSI: Retail financial services symmetric key management Part 1: Using symmetric techniques. Standard, ANSI X9.24-1:2009 (2009)
4. Improved client-side encryption: Explicit KeyIds and key commitment. <https://aws.amazon.com/blogs/security/improved-client-side-encryption-explicit-keyids-and-key-commitment/> (2020)
5. Amazon Web Services Key Management Service, <https://aws.amazon.com/kms/>
6. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: Security proofs and improvements. In: International Conference on the Theory and Applications of Cryptographic Techniques. pp. 259–274. Springer (2000)
7. Bellare, M., Canetti, R., Krawczyk, H.: Pseudorandom functions revisited: The cascade construction and its concrete security. In: Proceedings of 37th Conference on Foundations of Computer Science. pp. 514–523. IEEE (1996)
8. Bellare, M., Hoang, V.T.: Efficient schemes for committing authenticated encryption. In: Eurocrypt. pp. 845–875. Lecture Notes in Computer Science, Springer (2022)
9. Bellare, M., Hoang, V.T.: Efficient schemes for committing authenticated encryption. In: Advances in Cryptology - EUROCRYPT (2022)
10. Bellare, M., Ng, R., Tackmann, B.: Nonces are noticed: AEAD revisited. In: Advances in Cryptology - CRYPTO. Lecture Notes in Computer Science (2019)
11. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 409–426. Springer (2006)
12. Bellare, M., Tackmann, B.: The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In: Robshaw, M., Katz, J. (eds.) CRYPTO. Lecture Notes in Computer Science, vol. 9814, pp. 247–276. Springer (2016)
13. Chan, J., Rogaway, P.: Anonymous AE. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT. vol. 11922, pp. 183–208. Springer (2019)
14. Degabriele, J.P., Karadžić, V., Melloni, A., Münch, J.P., Stam, M.: Rugged pseudorandom permutations and their applications (2022), <https://rwc.iacr.org/2022/program.php>, Real World Crypto
15. Dodis, Y., Grubbs, P., Ristenpart, T., Woodage, J.: Fast message franking: From invisible salamanders to encryption. In: CRYPTO. pp. 155–186 (2018)
16. Farshim, P., Libert, B., Paterson, K.G., Quaglia, E.A.: Robust encryption, revisited. In: Public Key Cryptography. Lecture Notes in Computer Science, vol. 7778, pp. 352–368. Springer (2013)
17. Farshim, P., Orlandi, C., Rosie, R.: Security of symmetric primitives under incorrect usage of keys. IACR Transactions on Symmetric Cryptology (2017)
18. GlobalPlatform Technology Card Specification Version 2.3.1. Standard, GlobalPlatform (March 2018), [https://globalplatform.org/wp-content/uploads/2018/05/GPC\\_CardSpecification\\_v2.3.1\\_PublicRelease\\_CC.pdf](https://globalplatform.org/wp-content/uploads/2018/05/GPC_CardSpecification_v2.3.1_PublicRelease_CC.pdf)
19. Grubbs, P., Lu, J., Ristenpart, T.: Message franking via committing authenticated encryption. In: Katz, J., Shacham, H. (eds.) CRYPTO. Lecture Notes in Computer Science, vol. 10403, pp. 66–97. Springer (2017)
20. ISO/IEC: Information technology - security techniques - message authentication codes (MACs) - part 1: Mechanisms using a block cipher. Standard, ISO/IEC 9797-1:2011 (2011)
21. Iwata, T., Wang, L.: Impact of ANSI X9. 24-1: 2009 key check value on ISO/IEC 9797-1: 2011 MACs. In: International Workshop on Fast Software Encryption. pp. 303–322. Springer (2014)
22. Jaeger, J., Tyagi, N.: Handling adaptive compromise for practical encryption schemes. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO. Lecture Notes in Computer Science, vol. 12170, pp. 3–32. Springer (2020)
23. Len, J., Grubbs, P., Ristenpart, T.: Partitioning Oracle Attacks. In: USENIX Security (2021)
24. libsodium AEAD, [https://doc.libsodium.org/secret-key\\_cryptography/aead](https://doc.libsodium.org/secret-key_cryptography/aead)
25. Microsoft Key Vault, <https://azure.microsoft.com/en-us/services/key-vault/#product-overview>
26. Oracle Key Vault, <https://www.oracle.com/security/database-security/key-vault/>
27. PKCS #11 cryptographic token interface base specification version 2.40. Standard, OASIS (April 2015), <http://docs.oasis-open.org/pkcs11/pkcs11-base/v2.40/os/pkcs11-base-v2.40-os.pdf>
28. Rogaway, P.: Nonce-based symmetric encryption. In: Fast Software Encryption – FSE. pp. 348–358. Springer (2004)
29. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT. Lecture Notes in Computer Science, vol. 4004, pp. 373–390. Springer (2006)
30. Shadowsocks. <https://shadowsocks.org/en/index.html> (2020)
31. Telegram mobile protocol, <https://core.telegram.org/mtproto/description>
32. Google Tink library, <https://developers.google.com/tink>

<p style="margin: 0;"><b>XFROB<sub>AEKI</sub><sup>A</sup>:</b></p> <p style="margin: 0;"><math>\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C \leftarrow \mathcal{A}</math></p> <p style="margin: 0;"><math>(K_0, M_0) \leftarrow \text{Dec}(\mathbb{K}_0, N, AD, T_k, C)</math></p> <p style="margin: 0;"><math>(K_1, M_1) \leftarrow \text{Dec}(\mathbb{K}_1, N, AD, T_k, C)</math></p> <p style="margin: 0;">If <math>(K_0, M_0) = \perp</math>: Return 0</p> <p style="margin: 0;"><math>(\text{kid}_0, K_0^*) \leftarrow K_0</math>; <math>(\text{kid}_1, K_1^*) \leftarrow K_1</math></p> <p style="margin: 0;">If <math>K_0 \in \mathbb{K}_1 \wedge ((K_1, M_1) = \perp \vee K_0^* \neq K_1^*)</math>:</p> <p style="margin: 0; padding-left: 20px;">Return 1</p> <p style="margin: 0;">If <math>K_0 \notin \mathbb{K}_1 \wedge (K_1, M_1) \neq \perp</math>:</p> <p style="margin: 0; padding-left: 20px;">Return 1</p> <p style="margin: 0;">Return 0</p>
---

Fig. 10: Game  $\text{XFROB}_{\text{AEKI}}$  represents extended full robustness for AEAD-KI scheme  $\text{AEKI} = (\text{Kg}, \text{Enc}, \text{Dec})$ .

## A Extended Full Robustness

In an earlier draft of this work, we had provided a weaker correctness definition for AEAD-KI that did not include condition (3). This required the need for an alternative stronger key robustness definition than KI-FROB, which we call *extended full robustness* (XFROB). An observation due to Mihir Bellare through personal correspondence is that one can instead extend correctness to rule out the need for a stronger robustness notion. Thus, in the main body we include the stronger correctness definition and FROB as the main key robustness goal. For completeness, however, we describe here the stronger XFROB definition, for which one can weaken correctness by removing condition (3), and compare it with KI-FROB.

**XFROB.** We first observe that the key identification setting has subtleties that make KI-FROB not as strong of a robustness goal as we ultimately may want, given the weaker correctness definition. The winning condition for KI-FROB only captures when decryption succeeds for both key vectors. However, we have to consider the possibility that the particular keys in a key vector, or the order in which they appear, could have some adverse effect that causes spurious decryption failures. Intuitively, decryption should not behave inconsistently for related key vectors: imagine a scheme that decrypted properly a ciphertext using  $K \in \mathbb{K}$ , but that fails to decrypt entirely on  $\mathbb{K}'$  where  $K \in \mathbb{K}'$ . We expect that in most situations, this would be considered unexpected behavior, and adversaries should not be able to force it. Yet KI-FROB does *not* rule out such a scheme. We discuss this further below.

To capture ruling out such issues, we present a new key robustness definition for AEAD-KI, which we call extended full robustness (XFROB). We formalize it in Figure 10. The  $\text{XFROB}_{\text{AEKI}}$  advantage of an adversary  $\mathcal{A}$  is defined as

$$\text{Adv}_{\text{AEKI}}^{\text{xfrob}}(\mathcal{A}) = \Pr [\text{XFROB}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1].$$

In the following two theorems, we show that XFROB security is strictly stronger than KI-FROB security.

**Theorem 12.** *Let  $\mathcal{A}$  be a KI-FROB adversary. Then we give adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\text{AEKI}}^{\text{ki-frob}}(\mathcal{A}) \leq \text{Adv}_{\text{AEKI}}^{\text{xfrob}}(\mathcal{B})$$

where  $\mathcal{B}$  runs in time that of  $\mathcal{A}$ .

*Proof.* We construct XFROB adversary  $\mathcal{B}$  as follows. Adversary  $\mathcal{B}$  runs KI-FROB adversary  $\mathcal{A}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ , and simply returns what  $\mathcal{A}$  does. By inspecting games KI-FROB and XFROB, we see that any winning condition for KI-FROB is also a winning condition for XFROB, so  $\mathcal{B}$  wins game XFROB with at least the probability that  $\mathcal{A}$  wins game KI-FROB.  $\square$

For the following theorem, let  $\text{AEKI}$  be a KI-FROB AEAD-KI scheme. Now consider the AEAD-KI scheme  $\text{AEKI}^*$ , which is identical to  $\text{AEKI}$  except that it modifies decryption so that if some fixed key  $K^* \in \mathbb{K}$  is the first key in  $\mathbb{K}$ , where  $|\mathbb{K}| > 1$ , then decryption immediately returns  $\perp$ . Otherwise, decryption proceeds normally. Notice that  $\text{AEKI}^*$  meets the weaker correctness definition covering only conditions (1) and (2) if  $\text{AEKI}$  itself meets this weaker correctness definition, since  $\text{AEKI}^*$  only exhibits “incorrect” behavior when decryption is presented with multiple keys.

**Theorem 13.** *A concrete adversary  $\mathcal{A}$  exists such that  $\text{Adv}_{\text{AEKI}^*}^{\text{xfrob}}(\mathcal{A}) = 1$ . Furthermore, let  $\mathcal{A}_{\text{KI-FROB}}$  be a KI-FROB adversary for  $\text{AEKI}^*$ . Then we give KI-FROB adversary  $\mathcal{B}_{\text{KI-FROB}}$  against  $\text{AEKI}^*$  such that  $\text{Adv}_{\text{AEKI}^*}^{\text{ki-frob}}(\mathcal{A}_{\text{KI-FROB}}) \leq \text{Adv}_{\text{AEKI}^*}^{\text{ki-frob}}(\mathcal{B}_{\text{KI-FROB}})$ , where  $\mathcal{B}_{\text{KI-FROB}}$  runs in time that of  $\mathcal{A}_{\text{KI-FROB}}$ .*

*Proof.* We construct the following XFROB adversary  $\mathcal{A}$  against  $\text{AEKI}^*$ .  $\mathcal{A}$  chooses some key  $K_0 \neq K^*$ , some nonce  $N$ , associated data  $AD$ , and plaintext  $M$ , and then computes  $(T_k, C) \leftarrow \text{AEKI}^*. \text{Enc}(K_0, N, AD, M)$ . Then  $\mathcal{A}$  returns

$$[K_0], [K^*, K_0], N, AD, T_k, C.$$

The key vector  $[K_0]$  will decrypt the ciphertext correctly, but notice that the second key vector  $[K^*, K_0]$  will fail to decrypt the ciphertext since  $K^*$  is the first key. Thus, the XFROB game returns 1 with probability 1.

We construct  $\mathcal{B}_{\text{KI-FROB}}$  by running  $\mathcal{A}_{\text{KI-FROB}}$ , which returns  $\mathbb{K}_0, \mathbb{K}_1, N, AD, T_k, C$ , and then returning what  $\mathcal{A}_{\text{KI-FROB}}$  returns. Notice that neither of  $\mathbb{K}_0, \mathbb{K}_1$  can have  $K^*$  as the first key because otherwise decryption will fail. Therefore, decrypting using  $\mathbb{K}_0, \mathbb{K}_1$  results in the same outcome for  $\text{AEKI}$  as for  $\text{AEKI}^*$ . This means  $\mathcal{B}_{\text{KI-FROB}}$  wins game KI-FROB whenever  $\mathcal{A}_{\text{KI-FROB}}$  wins.  $\square$

**Discussion: Comparing XFROB and KI-FROB.** The benefit of XFROB over KI-FROB is that it guarantees that if the correct key is in the key vector, then decryption cannot fail. For instance, consider an AEAD-KI scheme where a particular key added to the key vector immediately triggers decryption failure. Or alternatively, consider a scheme where decryption might fail if keys are ordered in a certain way. Now imagine that an adversary can insert malicious keys into a key vector that holds honest keys the adversary does not know. If the adversary can submit a ciphertext that should decrypt under the malicious key yet decryption fails, then the adversary might learn something about other keys in the key vector. While we do not know of any such attacks in the wild, it is important to note that a non-XFROB scheme might inadvertently reveal information about the composition of the key vector. KI-FROB allows for these failures, but XFROB instead rules them out.

On the other hand, XFROB might not always be a desirable security target. In particular, imagine a scheme where decryption performs a type check on the key vector to ensure it conforms to a specific policy. As we will see later, one example of this is checking that all keys have unique key identifiers and forcing decryption to fail if this requirement is not met. This scheme trivially does not meet XFROB: we can construct an adversary that creates a ciphertext and two key vectors such that both hold the key used to create the ciphertext but one has keys with non-unique identifiers. XFROB expects that both key vectors would return the correct key, yet one key vector would result in decryption failure. In most settings, checks like this are harmless and could even prove beneficial. However, while type checks like this might seem like a natural implementation choice, it is interesting to note here they have subtle security implications. Indeed, it is possible that in certain settings as described above these checks could lead to some leakage of the composition of keys in the key vector.

We note, however, that correctness condition (3) eliminates the need for XFROB: any two key vectors that share keys between them cannot have only one of the key vectors successfully decrypt a ciphertext. Without this, our XFROB definition becomes equivalent to the KI-FROB definition.

## B KI-nAE implies INT-CTXT

Here we show that our KI-nAE definition implies authenticity and therefore serves as an all-in-one definition for the AEAD-KI setting. We begin by proposing a ciphertext integrity game extended to this setting and then prove that KI-nAE implies this definition.

**INT-CTXT for AEAD-KI.** Our ciphertext integrity game INT-CTXT extends the typical ciphertext integrity definition to the AEAD-KI setting. The pseudocode is shown in Figure 11. In particular, the adversary is given access to oracles GENKEY, ENC, and DEC. GENKEY enables the adversary to generate secret keys that are not revealed to the adversary. These keys are enumerated with a counter  $j$  that is incremented each time a key is generated and which serves as a key identifier within the game. ENC then expects an integer id as input which enables the adversary to refer to which key to use for encryption. The adversary gets back the key tag and ciphertext generated from encryption. The game also stores the tuple  $(\text{id}, N, AD, T_k, C)$  in set  $\mathcal{C}$  to keep track of which ciphertexts were generated by an honest key.

<p><b>INT-CTXT<sub>AEKI</sub><sup>A</sup>:</b>  <math>j \leftarrow 0 ; \mathbb{T} \leftarrow [ ] ; \mathbf{C} \leftarrow \emptyset</math>  <math>\text{win} \leftarrow 0</math>  <math>\mathcal{A}^{\text{GENKEY, ENC, DEC}}</math>  Return win</p> <p><b>GENKEY(kid):</b>  <math>K \leftarrow \text{AEKI.Kg}(\text{kid})</math>  <math>\mathbb{T}[j] \leftarrow K ; j \leftarrow j + 1</math></p>	<p><b>ENC(id, N, AD, M):</b>  <math>(T_k, C) \leftarrow \text{AEKI.Enc}(\mathbb{T}[\text{id}], N, AD, M)</math>  <math>\mathbf{C} \leftarrow \mathbf{C} \cup \{(id, N, AD, T_k, C)\}</math>  Return <math>(T_k, C)</math></p>	<p><b>DEC(<math>\mathbb{I}, N, AD, T_k, C</math>):</b>  If <math>\exists \text{id} \in \mathbb{I}</math> s.t. <math>(\text{id}, N, AD, T_k, C) \in \mathbf{C}</math>: Return <math>\mathbb{K} \leftarrow [ ]</math>  For <math>\text{id} \in \mathbb{I}</math>: <math>\mathbb{K}.\text{add}(\mathbb{T}[\text{id}])</math>  <math>(K, M) \leftarrow \text{AEKI.Dec}(\mathbb{K}, N, AD, T_k, C)</math>  If <math>(K, M) \neq \perp</math>:      <math>\text{win} \leftarrow 1</math>  Return <math>M</math></p>
--	---	---

Fig. 11: Game INT-CTXT<sub>AEKI</sub> captures the goal of ciphertext integrity for AEAD-KI scheme AEKI = (Kg, Enc, Dec).

DEC takes as input a vector  $\mathbb{I}$  of internal game key identifiers. If an ENC query with an identifier in  $\mathbb{I}$  was used to generate the queried ciphertext, then DEC immediately returns. Otherwise, it forms the key vector  $\mathbb{K}$  of the secret keys corresponding to the identifiers in  $\mathbb{I}$  and calls AEKI.Dec() to decrypt the specified ciphertext. If decryption succeeds, then the game sets the flag win to 1. The game returns the value of win after running the adversary. This means the adversary wins only when it can produce a tuple  $(N, AD, T_k, C)$  such that it can successfully be decrypted by a secret key that was not used to produce it by a call to ENC.

The INT-CTXT<sub>AEKI</sub> advantage of an adversary  $\mathcal{A}$  is defined as

$$\text{Adv}_{\text{AEKI}}^{\text{ki-frob}}(\mathcal{A}) = \Pr [ \text{INT-CTXT}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1 ] .$$

**Theorem 14.** *Let  $\mathcal{A}$  be a INT-CTXT adversary making at most  $q$  queries to its oracles. Then we give adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\text{AEKI}}^{\text{int-ctxt}}(\mathcal{A}) \leq \text{Adv}_{\text{AEKI}, \mathcal{S}}^{\text{ki-nae}}(\mathcal{B}).$$

Adversary  $\mathcal{B}$  runs in time that of  $\mathcal{A}$  with an  $\mathcal{O}(q)$  overhead.

*Proof.* We construct KI-nAE adversary  $\mathcal{B}$  against AEKI as follows.  $\mathcal{B}$  runs INT-CTXT adversary  $\mathcal{A}$ . When  $\mathcal{A}$  makes a GENKEY query,  $\mathcal{B}$  makes a GENHONESTKEY query. When  $\mathcal{A}$  makes an ENC query,  $\mathcal{B}$  calls its own ENC and returns to  $\mathcal{A}$  what its oracle returns.  $\mathcal{B}$  also keeps track of what tuples  $(\text{id}, N, AD, T_k, C)$  have been output in a table  $\mathbf{C}$ , like in game INT-CTXT. When  $\mathcal{A}$  makes a DEC query,  $\mathcal{B}$  checks whether  $\exists \text{id} \in \mathbb{I}$  such that  $(\text{id}, N, AD, T_k, C)$  are in  $\mathbf{C}$  and, if so, returns.

Otherwise,  $\mathcal{B}$  creates key vector  $\mathbb{K}$  so that its elements are  $(\text{true}, \text{id})$  for each  $\text{id} \in \mathbb{I}$ . Adversary  $\mathcal{B}$  then queries its own oracle DEC with  $\mathbb{K}$  and  $N, AD, T_k, C$ , returning to  $\mathcal{A}$  the result of this query. If any DEC query returns something other than  $\perp$ , then  $\mathcal{B}$  returns 1, else  $\mathcal{B}$  returns 0.

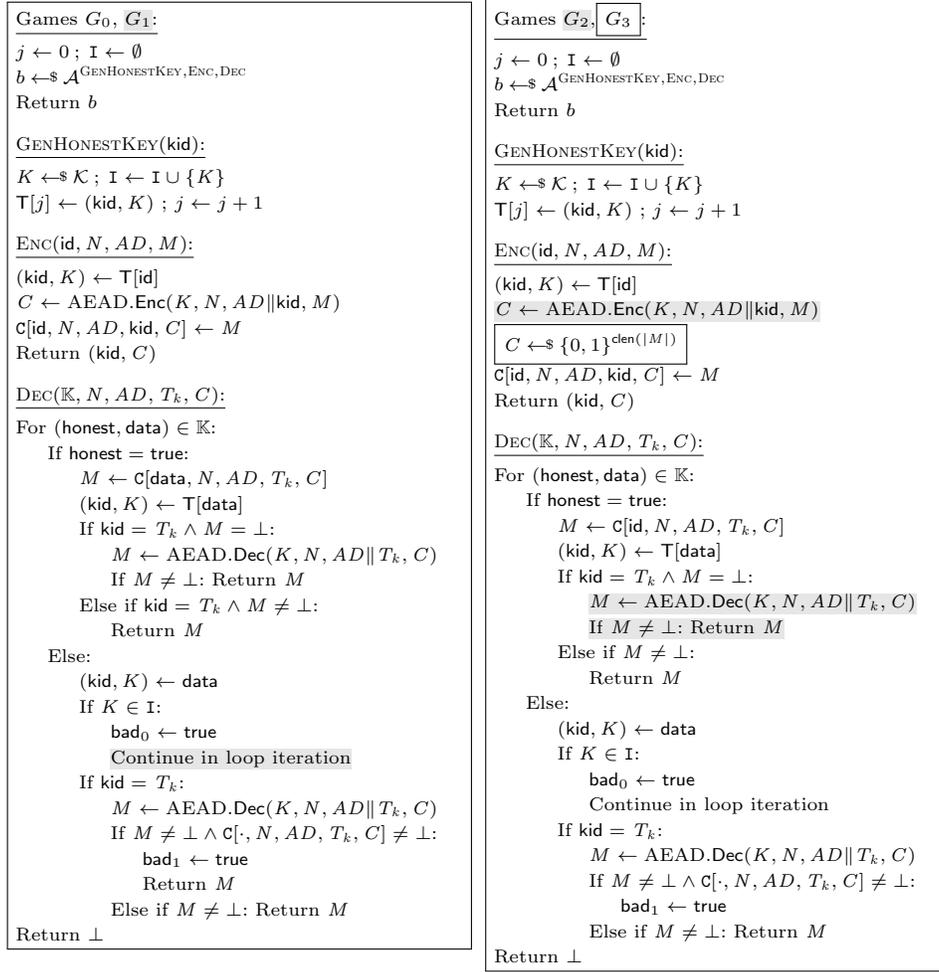
We then have that  $\Pr [ \text{KI-nAE1}_{\text{AEKI}}^{\mathcal{A}} \Rightarrow 1 ] = \Pr [ \text{INT-CTXT} \Rightarrow 1 ]$  because in the real world the probability that  $\mathcal{B}$  returns 1 is the probability that  $\mathcal{A}$  submits  $N, AD, T_k, C$  that can be successfully decrypted by a key that did not produce it. We also have that  $\Pr [ \text{KI-nAE0}_{\text{AEKI}, \mathcal{S}, \text{L}_{\text{Enc}}^{\text{id}}}^{\mathcal{A}} \Rightarrow 1 ] = 0$ . To understand why, first notice that  $\mathcal{B}$  only queries DEC with honest keys, so the simulator will never be called in DEC. Furthermore,  $\mathcal{B}$  only queries DEC on ciphertexts not generated by ENC by the keys specified in  $\mathbb{I}$ , which means DEC always returns  $\perp$ . This means  $\mathcal{B}$  always returns 0, completing the proof.  $\square$

## C Proof of Theorem 2

*Proof.* We construct a KI-nAE simulator as shown below. We will argue that such a simulator satisfies the KI-nAE definition using  $\text{L}_{\text{Enc}}^{\text{id}}$  for scheme KL.

Assume by contradiction that  $\mathcal{S}$  does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build FROB adversary  $\mathcal{B}$  against AEAD and MU-nAE adversary  $\mathcal{C}$  against AEAD. The proof follows from the following claims:

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr [ G_0 ] - \Pr [ G_1 ]| \leq mq/|\mathcal{K}|$
3.  $|\Pr [ G_1 ] - \Pr [ G_2 ]| \leq \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B})$

Fig. 12: Games  $G_0$  through  $G_3$  for the proof of Theorem 2.

4.  $|\Pr[G_2] - \Pr[G_3]| \leq \text{Adv}_{\text{AEAD}}^{\text{mu-nae}}(C)$
5. Game  $G_3$  is equivalent to KI-nAE0 with simulator  $S$

**Claim 1:**  $G_0$  is a syntactic rewriting of KI-nAE1 for scheme KL. The encryption oracle ENC additionally keeps track of all unique key-message pairs that output the ciphertext  $(N, AD, \text{kid}, C)$  for the key with internal game identifier id.

The decryption oracle DEC iterates through the key vector to find the first key, if any, that produces a matching key tag. If the key is honest but there is no entry for the ciphertext and key tag in the table, then the game uses AEAD to decrypt. If there is an entry, then the game returns the plaintext stored that has a key matching the key used to decrypt, which by the correctness of AEAD should match the decryption output. Finally, if the key is malicious, then the game again uses AEAD to decrypt.

**Claim 2:** Game  $G_1$  is identical to  $G_0$  until the flag  $\text{bad}_0$  is set. In  $G_0$  the flag  $\text{bad}_0$  is set when during a call to DEC the malicious AEAD encryption key is actually an honest key previously generated during a call to GENHONESTKEY. If  $\text{bad}_0$  is set in  $G_1$ , DEC ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_0 \Rightarrow 1] - \Pr[G_1 \Rightarrow 1]| \leq \Pr[G_1 \text{ sets } \text{bad}_0].$$

At best  $\mathcal{A}$  would only be able to set  $\text{bad}_0$  if it guessed one of the honest keys. Using a union bound, we can therefore upper bound the righthand term by  $m q / |\mathcal{K}|$ , since there are at most  $m$  malicious keys and at most  $q$  honest keys.

<u>S.Init():</u>	<u>S.Enc(<math>N, AD, \mathcal{L}, \sigma_s</math>):</u>	<u>S.Dec(<math>\mathbb{K}, N, AD, T_k, C, \sigma_s</math>):</u>
Return (0, [·])	( $id, T$ ) $\leftarrow \sigma_s$	For ( $honest, data$ ) $\in \mathbb{K}$ :
<u>S.Kg(<math>kid, \sigma_s</math>):</u>	( $j, \ell$ ) $\leftarrow \mathcal{L}$	If $honest = false$ :
( $id, T$ ) $\leftarrow \sigma_s$	$C \leftarrow_{\$} \{0, 1\}^{clen(\ell)}$	( $kid, K$ ) $\leftarrow data$
$T[id] \leftarrow kid$ ; $id \leftarrow id + 1$	$T_k \leftarrow T[j]$	If $kid = T_k$ :
Return ( $kid, (id, T)$ )	Return ( $T_k, C, \sigma_s$ )	$M \leftarrow AEAD.Dec(K, N, AD    T_k, C)$
		If $M \neq \perp$ : Return ( $M, \sigma_s$ )
		Return ( $\perp, \sigma_s$ )

Fig. 13: Simulator S for the proof of Theorem 2.

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $bad_1$  is set. In  $G_1$  the flag  $bad_1$  is set during a call to DEC when a malicious key successfully decrypts a ciphertext that was output by ENC.  $G_1$  returns the plaintext  $M$  that the malicious key decrypts the ciphertext to, while  $G_2$  ignores this result and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } bad_1].$$

We now bound the rightmost term by providing the following FROB adversary  $\mathcal{B}$  such that

$$\Pr[G_2 \text{ sets } bad_1] \leq \mathbf{Adv}_{AEAD}^{\text{frob}}(\mathcal{B}).$$

$\mathcal{B}$  runs  $\mathcal{A}$  and simulates game  $G_2$  for  $\mathcal{A}$ . If the flag  $bad_1$  is set to **true**, then  $\mathcal{B}$  halts and returns the malicious key  $K$  used to decrypt the ciphertext, the honest key  $K'$  stored in  $T$  that output  $N, AD, T_k, C$  during a call to ENC, nonce  $N$ , associated data  $AD || T_k$ , and ciphertext  $C$ . Notice that  $K$  and  $K'$  must be distinct, since both games do not allow for a malicious key to be one of the honestly generated keys. Therefore,  $\mathcal{B}$  wins game FROB when  $bad_1$  is set.

**Claim 4:** Game  $G_3$  differs from  $G_2$  in that during a query to the encryption oracle  $G_3$  chooses a random string as the ciphertext. Furthermore, during a query to the decryption oracle, if an honest key was called to decrypt a ciphertext not output by ENC, then that key is not used to decrypt. Notice that for an honest key, it only needs to check that the associated plaintext is not  $\perp$ , and not that  $kid$  matches  $T_k$ , because  $M \neq \perp$  would already imply that  $kid = T_k$ .

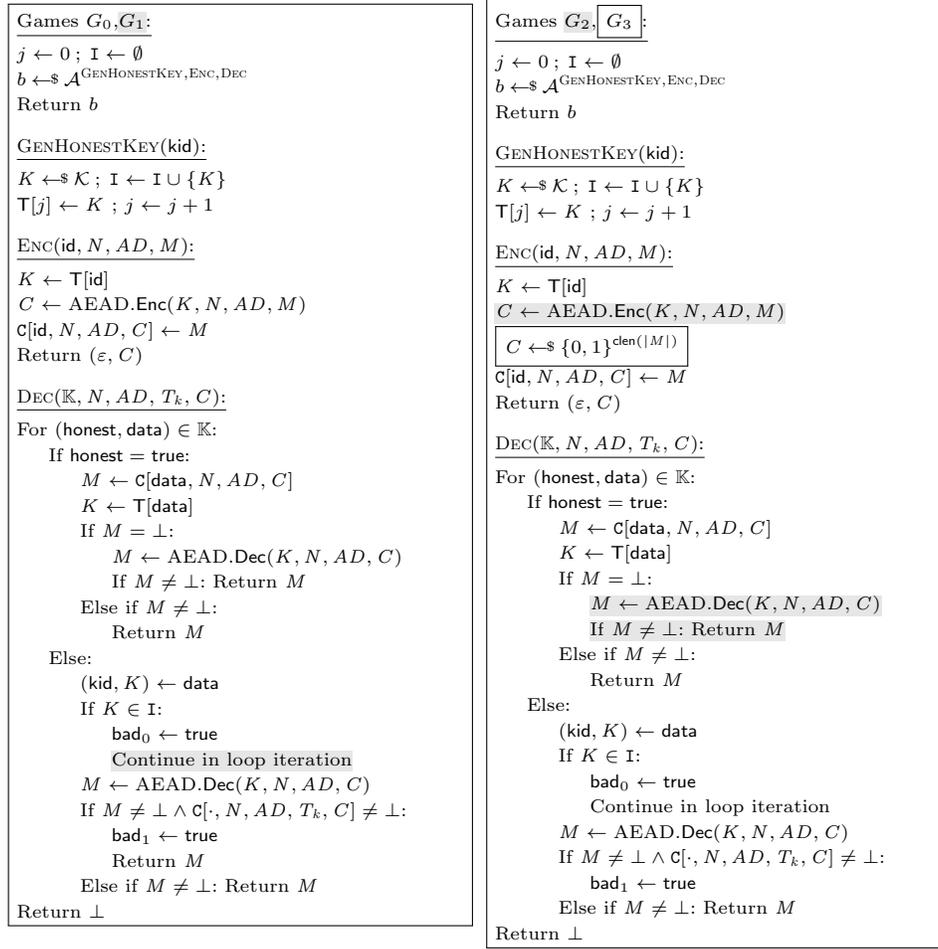
We can bound the advantage of  $\mathcal{A}$  in distinguishing between these games by reducing to the MU-nAE security of the AEAD scheme; call this reduction  $\mathcal{C}$ . Adversary  $\mathcal{C}$  receives oracles GENKEY, ENC, and DEC from game MU-nAE<sub>AEAD</sub> and simulates oracles GENKEYSIM, ENCSIM, and DECSIM from games  $G_2$  and  $G_3$  for adversary  $\mathcal{A}$ . In GENKEYSIM,  $\mathcal{C}$  calls GENKEY to generate the key and then assigns the key label in the same way as in games  $G_2, G_3$ . In ENCSIM,  $\mathcal{C}$  calls ENC to encrypt the queried plaintext with the specified key and then returns this with the key label as the key tag.

In DECSIM,  $\mathcal{C}$  iterates through the key vector  $\mathbb{K}$ . If a key is honest and its associated identifier matches the key tag, then the decryption oracle DEC is called on the key and ciphertext to decrypt. If decryption is successful, then the resulting plaintext is returned. Otherwise, if a key is malicious and its identifier matches the key tag, then  $\mathcal{C}$  decrypts the ciphertext itself using the key. Again, if decryption is successful, then the resulting plaintext is returned. If no key is successful, then  $\perp$  is returned. Finally,  $\mathcal{C}$  returns what  $\mathcal{A}$  returns.

Notice that when  $\mathcal{C}$  is playing game MU-nAE1, then a call to DEC actually decrypts the ciphertext, so that in this case  $\mathcal{C}$  simulates game  $G_2$  for  $\mathcal{A}$ . However, when  $\mathcal{C}$  is playing MU-nAE0, then DEC only returns  $\perp$ . In this case,  $\mathcal{C}$  simulates game  $G_3$ , which does not attempt to decrypt in this case but simply moves on to the next key in  $\mathbb{K}$ . Since  $\mathcal{C}$  simulates both games  $G_2, G_3$  for  $\mathcal{A}$ , whenever  $\mathcal{A}$  is able to distinguish between the two games,  $\mathcal{C}$  is able to distinguish between MU-nAE1, MU-nAE0, thereby proving our claim.

**Claim 5:** We can see by inspecting the pseudocode of  $G_3$  that it is almost a syntactic rewriting of game KI-nAE0 with simulator S. During a query to DEC, however,  $G_3$  decrypts in order of the keys given in  $\mathbb{K}$ , while game KI-nAE0 with S first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_3$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_3$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims.  $\square$

Fig. 14: Games  $G_0$  through  $G_3$  for the proof of Theorem 3.

## D Proof of Theorem 3

*Proof.* We construct a KI-nAE-KA simulator as shown below. We will argue that such a simulator satisfies the KI-nAE-KA definition using  $\mathbf{L}_{\text{Enc}}^{\text{anon}}$  for scheme TD.

Assume by contradiction that  $\mathcal{S}$  does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build FROB adversary  $\mathcal{B}$  against AEAD and MU-nAE adversary  $\mathcal{C}$  against AEAD. The proof follows from the following claims:

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr[G_0] - \Pr[G_1]| \leq mq/|\mathcal{K}|$
3.  $|\Pr[G_1] - \Pr[G_2]| \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B})$
4.  $|\Pr[G_2] - \Pr[G_3]| \leq \mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{C})$
5. Game  $G_3$  is equivalent to KI-nAE0 with simulator  $\mathcal{S}$

**Claim 1:**  $G_0$  is a syntactic rewriting of KI-nAE1 for scheme TD. The encryption oracle ENC additionally keeps track of all unique key-message pairs that output the ciphertext  $(N, AD, C)$  for the key with internal game identifier id. Since all key labels are  $\varepsilon$ , the game does not keep track of them.

The decryption oracle DEC iterates through the key vector to find the first key, if any, that can decrypt the ciphertext. If the key is honest but there is no entry for the ciphertext in the table, then the game uses AEAD to decrypt. If there is an entry, then the game returns the plaintext stored that has a key matching the key used to decrypt, which by the correctness of AEAD should match the decryption output. Finally, if the key is malicious, then the game again uses AEAD to decrypt.

<u>S.Init():</u>	<u>S.Enc(<math>N, AD, \mathcal{L}, \sigma_s</math>):</u>	<u>S.Dec(<math>\mathbb{K}, N, AD, T_k, C, \sigma_s</math>):</u>
Return $\emptyset$	$\ell \leftarrow \mathcal{L}$	For $(\text{honest}, \text{data}) \in \mathbb{K}$ :
<u>S.Kg(<math>\text{kid}, \sigma_s</math>):</u>	$C \leftarrow_{\mathcal{S}} \{0, 1\}^{\text{clen}(\ell)}$	If $\text{honest} = \text{false}$ :
Return $(\varepsilon, \sigma_s)$	Return $(\varepsilon, C, \sigma_s)$	$(\text{kid}, K) \leftarrow \text{data}$
		$M \leftarrow \text{AEAD.Dec}(K, N, AD, C)$
		If $M \neq \perp$ : Return $(M, \sigma_s)$
		Return $(\perp, \sigma_s)$

Fig. 15: Simulator  $\mathcal{S}$  for the proof of Theorem 3.

**Claim 2:** Game  $G_1$  is identical to  $G_0$  until the flag  $\text{bad}_0$  is set. In  $G_0$  the flag  $\text{bad}_0$  is set when during a call to DEC the malicious AEAD encryption key is actually an honest key previously generated during a call to GENHONESTKEY. If  $\text{bad}_0$  is set in  $G_1$ , DEC ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_0 \Rightarrow 1] - \Pr[G_1 \Rightarrow 1]| \leq \Pr[G_1 \text{ sets } \text{bad}_0].$$

At best  $\mathcal{A}$  would only be able to set  $\text{bad}_0$  if it guessed one of the honest keys. Using a union bound, we can therefore upper bound the righthand term by  $mq/|\mathcal{K}|$ , since there are at most  $m$  malicious keys and at most  $q$  honest keys.

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $\text{bad}_1$  is set. In  $G_1$  the flag  $\text{bad}_1$  is set during a call to DEC when a malicious key successfully decrypts a ciphertext that was output by ENC.  $G_1$  returns the plaintext  $M$  that the malicious key decrypts the ciphertext to, while  $G_2$  ignores this result and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } \text{bad}_1].$$

We now bound the rightmost term by providing the following FROB adversary  $\mathcal{B}$  such that

$$\Pr[G_2 \text{ sets } \text{bad}_1] \leq \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{B}).$$

$\mathcal{B}$  runs  $\mathcal{A}$  and simulates game  $G_2$  for  $\mathcal{A}$ . If the flag  $\text{bad}_1$  is set to true, then  $\mathcal{B}$  halts and returns the malicious key  $K$  used to decrypt the ciphertext, the honest key  $K'$  stored in  $\mathbb{T}$  that output  $N, AD, T_k, C$  during a call to ENC, nonce  $N$ , associated data  $AD \| T_k$ , and ciphertext  $C$ . Notice that  $K$  and  $K'$  must be distinct, since both games do not allow for a malicious key to be one of the honestly generated keys. Therefore,  $\mathcal{B}$  wins game FROB when  $\text{bad}_1$  is set.

**Claim 4:** Game  $G_3$  differs from  $G_2$  in that during a query to the encryption oracle  $G_3$  chooses a random string as the ciphertext. Furthermore, during a query to the decryption oracle, if an honest key was called to decrypt a ciphertext not output by ENC, then that key is not used to decrypt. Notice that for an honest key, it only needs to check that the associated plaintext is not  $\perp$ , and not that kid matches  $T_k$ , because  $M \neq \perp$  would already imply that  $\text{kid} = T_k$ .

We can bound the advantage of  $\mathcal{A}$  in distinguishing between these games by reducing to the MU-nAE security of the AEAD scheme; call this reduction  $\mathcal{C}$ . Adversary  $\mathcal{C}$  receives oracles GENKEY, ENC, and DEC from game MU-nAE<sub>AEAD</sub> and simulates oracles GENKEYSIM, ENCSIM, and DECSIM from games  $G_2$  and  $G_3$  for adversary  $\mathcal{A}$ . In GENKEYSIM,  $\mathcal{C}$  calls GENKEY to generate the key and then assigns the key label in the same way as in games  $G_2, G_3$ . In ENCSIM,  $\mathcal{C}$  calls ENC to encrypt the queried plaintext with the specified key and then returns this with the key label as the key tag.

In DECSIM,  $\mathcal{C}$  iterates through the key vector  $\mathbb{K}$ . If a key is honest and its associated identifier matches the key tag, then the decryption oracle DEC is called on the key and ciphertext to decrypt. If decryption is successful, then the resulting plaintext is returned. Otherwise, if a key is malicious and its identifier matches the key tag, then  $\mathcal{C}$  decrypts the ciphertext itself using the key. Again, if decryption is successful, then the resulting plaintext is returned. If no key is successful, then  $\perp$  is returned. Finally,  $\mathcal{C}$  returns the bit that  $\mathcal{A}$  returns.

Notice that when  $\mathcal{C}$  is playing game MU-nAE1, then a call to DEC actually decrypts the ciphertext, so that in this case  $\mathcal{C}$  simulates game  $G_2$  for  $\mathcal{A}$ . However, when  $\mathcal{C}$  is playing MU-nAE0, then DEC only returns

⊥. In this case,  $\mathcal{C}$  simulates game  $G_3$ , which does not attempt to decrypt in this case but simply moves on to the next key in  $\mathbb{K}$ . Since  $\mathcal{C}$  simulates both games  $G_2, G_3$  for  $\mathcal{A}$ , whenever  $\mathcal{A}$  is able to distinguish between the two games,  $\mathcal{C}$  is able to distinguish between MU-nAE1, MU-nAE0, thereby proving our claim.

**Claim 5:** We can see by inspecting the pseudocode of  $G_3$  that it is almost a syntactic rewriting of game KI-nAE0 with simulator  $S$ . During a query to DEC, however,  $G_3$  decrypts in order of the keys given in  $\mathbb{K}$ , while game KI-nAE0 with  $S$  first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_3$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_3$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims.

□

## E Proofs of Theorem 5 and Theorem 7

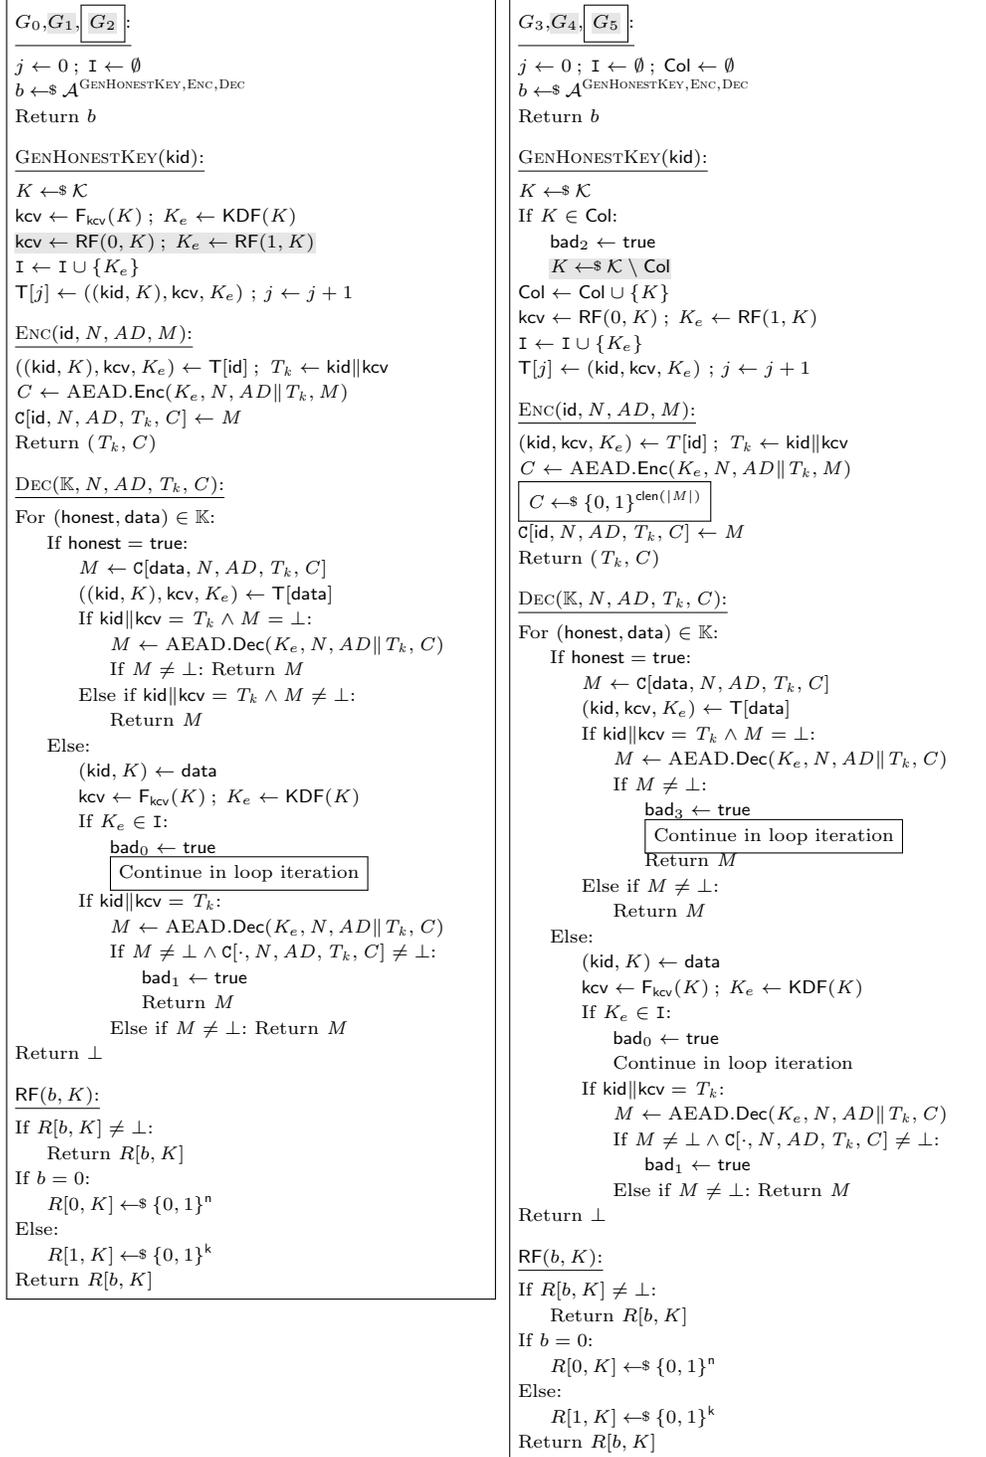


Fig. 16: Games  $G_0$  through  $G_5$  for the proof of Theorem 5. Games  $G_1$  and  $G_4$  feature the pseudocode that is high-lighted, while games  $G_2$  and  $G_5$  feature the pseudocode that is highlighted and the pseudocode that is boxed.

<p><b>S.Init():</b> Return (0, [])</p> <p><b>S.Kg(kid, <math>\sigma_s</math>):</b> (id, <math>T</math>) <math>\leftarrow \sigma_s</math>; <math>\text{kcv} \leftarrow \mathcal{S} \{0, 1\}^n</math> <math>T[\text{id}] \leftarrow \text{kid} \parallel \text{kcv}</math>; <math>\text{id} \leftarrow \text{id} + 1</math> Return (id, <math>T</math>)</p>	<p><b>S.Enc(<math>N, AD, \mathcal{L}, \sigma_s</math>):</b> (id, <math>T</math>) <math>\leftarrow \sigma_s</math>; (<math>j, \ell</math>) <math>\leftarrow \mathcal{L}</math> <math>C \leftarrow \mathcal{S} \{0, 1\}^{\text{clen}(\ell)}</math>; <math>T_k \leftarrow T[j]</math> Return (<math>T_k, C, \sigma_s</math>)</p>	<p><b>S.Dec(<math>\mathbb{K}, N, AD, T_k, C, \sigma_s</math>):</b> For (honest, data) <math>\in \mathbb{K}</math>:   If honest = false:     (kid*, <math>K^*</math>) <math>\leftarrow</math> data     <math>\text{kcv}^* \leftarrow F_{\text{kcv}}(K^*)</math>; <math>K_e \leftarrow \text{KDF}(K^*)</math>     If kid* <math>\parallel</math> <math>\text{kcv}^* = T_k</math>:       <math>M \leftarrow \text{AEAD.Dec}(K_e, N, AD \parallel T_k, C)</math>       If <math>M \neq \perp</math>: Return (<math>M, \sigma_s</math>) Return (<math>\perp, \sigma_s</math>)</p>
---	---	--

Fig. 17: Simulator S for the proofs of Theorem 5 and Theorem 7.

## E.1 Proof of Theorem 5

*Proof.* We construct the KI-nAE simulator S as shown in Figure 17. We will argue that such a simulator satisfies the KI-nAE definition using  $L_{\text{Enc}}^{\text{id}}$  for scheme KCV.

Assume by contradiction that S does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build the adversaries from the theorem statement using a series of game hops. We provide the pseudocode for games  $G_0$  through  $G_5$  in Figure 16. The proof follows from the following claims:

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr[G_0] - \Pr[G_1]| \leq \mathbf{Adv}_{F_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B})$
3.  $|\Pr[G_1] - \Pr[G_2]| \leq q_k \cdot \mathbf{Adv}_{\text{KDF}}^{\text{pre}}(\mathcal{C})$
4.  $|\Pr[G_2] - \Pr[G_3]| \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{D})$
5.  $|\Pr[G_3] - \Pr[G_4]| \leq \frac{q_k^2}{2^{\kappa+1}}$
6.  $|\Pr[G_4] - \Pr[G_5]| \leq \mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{E})$
7. Game  $G_5$  is equivalent to KI-nAE0 with simulator S

**Claim 1:**  $G_0$  is a syntactic rewriting of KI-nAE1 for scheme KCV. Since key check values and AEAD encryption keys remain static for each secret key, these values are computed in GENHONESTKEY and stored. The encryption oracle ENC additionally keeps track of all unique key-message pairs that output the ciphertext ( $N, AD, T_k, C$ ) for the key with internal game identifier id.

The decryption oracle DEC iterates through the key vector to find the first key, if any, that produces a matching key tag. If the key is honest but there is no entry for the ciphertext and key tag in the table, then the game uses AEAD to decrypt. If there is an entry, then the game returns the plaintext stored that has a key matching the key used to decrypt, which by the correctness of AEAD should match the decryption output. Finally, if the key is malicious, then the game again uses AEAD to decrypt when the key tag matches the key check value.

**Claim 2:** Game  $G_1$  is identical to  $G_0$  except that all calls to  $F_{\text{kcv}}$  and KDF for honestly generated keys are replaced with calls to random function RF. RF takes as input a bit  $b$ , where  $b = 0$  represents a call to  $F_{\text{kcv}}$  and  $b = 1$  represents a call to KDF and value  $K$ . For any new input ( $b, K$ ), the function chooses a value uniformly at random from either  $\{0, 1\}^n$  (when  $b = 0$ ) or  $\{0, 1\}^k$  (when  $b = 1$ ), then stores this value in table  $R$ . When RF is called on a repeat input, the value chosen previously and stored in  $R$  is returned.

We can bound the advantage of  $\mathcal{A}$  in distinguishing between  $G_0$  and  $G_1$  by reducing to the MU-PRF security of  $F_{\text{kcv}}$  and KDF; call this reduction  $\mathcal{B}$ . The adversary  $\mathcal{B}$  simulates games  $G_0$  and  $G_1$  by calling GENKEY to generate the AEAD-KI key, Func0 to generate  $\text{kcv}$ , and Func1 to generate  $K_e$ . The adversary  $\mathcal{B}$  then returns whatever  $\mathcal{A}$  does. When  $\mathcal{B}$  is in the REAL or IDEAL game, then  $\mathcal{B}$  correctly simulates game  $G_0$  or  $G_1$ , respectively, for  $\mathcal{A}$ .

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $\text{bad}_0$  is set. In  $G_1$  the flag  $\text{bad}_0$  is set when during a call to DEC the AEAD encryption key  $K_e$  output by some malicious key  $K$  was already output for some honest key in GENHONESTKEY. If  $\text{bad}_0$  is set in  $G_2$ , DEC ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } \text{bad}_0].$$

We now bound the rightmost term by providing the following  $\text{PRE}_{\text{KDF}}$  adversary  $\mathcal{C}$  such that

$$\Pr [G_2 \text{ sets } \text{bad}_0] \leq q_k \cdot \text{Adv}_{\text{KDF}}^{\text{pre}}(\mathcal{C}).$$

$\mathcal{C}$  gets as input a random range point  $y$  in  $\{0, 1\}^k$  and runs  $\mathcal{A}$ , simulating game  $G_2$  for  $\mathcal{A}$ .  $\mathcal{C}$  chooses at random to replace  $K_e$  with  $y$  during one of  $\mathcal{A}$ 's queries to  $\text{GENHONESTKEY}$ . When  $\text{bad}_0$  is set,  $\mathcal{C}$  immediately halts and returns the malicious key  $K$  as a pre-image for  $y$ . We abuse notation here and can assume that the second input is the empty string. Notice that since there are  $q_k$  honest keys generated, there is probability  $1/q_k$  that the value for which  $\mathcal{A}$  found a pre-image is  $y$ . Therefore, we can lower bound the advantage of  $\mathcal{C}$  by  $\Pr [G_2 \text{ sets } \text{bad}_0] \cdot 1/q_k$ .

**Claim 4:** Game  $G_3$  is identical to  $G_2$  until the flag  $\text{bad}_1$  is set. In  $G_2$  the flag  $\text{bad}_1$  is set during a call to  $\text{DEC}$  when a malicious key successfully decrypts a ciphertext that was output by  $\text{ENC}$ .  $G_2$  returns the plaintext  $M$  that the malicious key decrypts the ciphertext to, while  $G_3$  ignores this result and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr [G_2 \Rightarrow 1] - \Pr [G_3 \Rightarrow 1]| \leq \Pr [G_3 \text{ sets } \text{bad}_1].$$

We now bound the rightmost term by providing the following  $\text{FROB}$  adversary  $\mathcal{D}$  such that

$$\Pr [G_3 \text{ sets } \text{bad}_1] \leq \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{D}).$$

$\mathcal{D}$  runs  $\mathcal{A}$  and simulates game  $G_3$  for  $\mathcal{A}$ . If the flag  $\text{bad}_1$  is set to `true`, then  $\mathcal{D}$  halts and returns the malicious key  $K_e$  used to decrypt the ciphertext, the honest key  $K'_e$  stored in  $\mathbb{T}$  that output  $N, AD, T_k, C$  during a call to  $\text{ENC}$ , nonce  $N$ , associated data  $AD \| T_k$ , and ciphertext  $C$ . Notice that  $K_e$  and  $K'_e$  must be distinct, since both games do not allow for a malicious key to output an AEAD encryption key that was already output by an honest key. Therefore,  $\mathcal{D}$  wins game  $\text{FROB}$  when  $\text{bad}_1$  is set.

**Claim 5:** Game  $G_4$  differs from  $G_3$  in that during a query to  $\text{GENHONESTKEY}$  if the key  $K$  chosen at random from  $\mathcal{K}$  has already been chosen on a previous query, then a new fresh value is chosen that is guaranteed to not be a repeat value. Recall that  $q_k$  is the number of queries  $\mathcal{A}$  makes to  $\text{GENHONESTKEY}$  and  $\mathcal{K} = 2^\kappa$  for this scheme. We can bound  $\mathcal{A}$ 's distinguishing advantage between games  $G_3$  and  $G_4$  by the probability of a collision occurring, which is bounded by the birthday bound  $q_k^2/2^{\kappa+1}$ .

**Claim 6:** Game  $G_5$  differs from  $G_4$  in that during a query to the encryption oracle it chooses a random string as the ciphertext and during a query to the decryption oracle if an honest key was called to decrypt a ciphertext not output by  $\text{ENC}$  then that key is not used to decrypt. We can bound the advantage of  $\mathcal{A}$  in distinguishing between these games by reducing to the  $\text{MU-nAE}$  security of the AEAD scheme; call this reduction  $\mathcal{E}$ . The adversary  $\mathcal{E}$  uses its own encryption oracle to simulate queries by  $\mathcal{A}$  to  $\text{ENC}$  and its decryption oracle to simulate queries by  $\mathcal{A}$  to  $\text{DEC}$ . When  $\mathcal{E}$  is playing game  $\text{MU-nAE1}$  or  $\text{MU-nAE0}$ , then  $\mathcal{E}$  correctly simulates game  $G_5$  or  $G_4$ , respectively, for  $\mathcal{A}$ .

**Claim 7:** We can see by inspecting the pseudocode of  $G_5$  that it is almost a syntactic rewriting of game  $\text{KI-nAE0}$  with simulator  $\mathbb{S}$ . During a query to  $\text{DEC}$ , however,  $G_5$  decrypts in order of the keys given in  $\mathbb{K}$ , while game  $\text{KI-nAE0}$  with  $\mathbb{S}$  first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_5$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_5$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims. □

## E.2 Proof of Theorem 7

<p><math>G_0, G_1, G_2</math>:</p> <p><math>j \leftarrow 0; I \leftarrow \emptyset</math>  <math>b \leftarrow \mathcal{A}^{\text{GENHONESTKEY, ENC, DEC}}</math>  Return <math>b</math></p> <p><u>GENHONESTKEY(kid):</u>  <math>K \leftarrow \mathcal{K}</math>  If <math>K \in \text{Col}</math>:    <math>\text{bad}_1 \leftarrow \text{true}</math>  <math>\text{Col} \leftarrow \text{Col} \cup \{K\}</math>  <math>\text{kcv} \leftarrow F_{\text{kcv}}(K); K_e \leftarrow \text{KDF}(K)</math>  <math>\text{kcv} \leftarrow \text{RF}(0, K); K_e \leftarrow \text{RF}(1, K)</math>  <math>I \leftarrow I \cup \{\text{kcv}\}</math>  <math>T[j] \leftarrow ((\text{kid}, K), \text{kcv}, K_e); j \leftarrow j + 1</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>((\text{kid}, K), \text{kcv}, K_e) \leftarrow T[\text{id}]; T_k \leftarrow \text{kid} \parallel \text{kcv}</math>  <math>C \leftarrow \text{AEAD.ENC}(K_e, N, AD \parallel T_k, M)</math>  <math>C[\text{id}, N, AD, T_k, C] \leftarrow M</math>  Return <math>(T_k, C)</math></p> <p><u>DEC(<math>\mathbb{K}, N, AD, T_k, C</math>):</u>  For (honest, data) <math>\in \mathbb{K}</math>:    If honest = true:      <math>M \leftarrow C[\text{data}, N, AD, T_k, C]</math>      <math>((\text{kid}, K), \text{kcv}, K_e) \leftarrow T[\text{data}]</math>      If <math>\text{kid} \parallel \text{kcv} = T_k \wedge M = \perp</math>:        <math>M \leftarrow \text{AEAD.DEC}(K_e, N, AD \parallel T_k, C)</math>        If <math>M \neq \perp</math>: Return <math>M</math>      Else if <math>\text{kid} \parallel \text{kcv} = T_k \wedge M \neq \perp</math>:        Return <math>M</math>    Else:      <math>(\text{kid}, K) \leftarrow \text{data}</math>      <math>\text{kcv} \leftarrow F_{\text{kcv}}(K); K_e \leftarrow \text{KDF}(K)</math>      If <math>\text{kcv} \in I</math>:        <math>\text{bad}_0 \leftarrow \text{true}</math>        Continue in loop iteration      If <math>\text{kid} \parallel \text{kcv} = T_k</math>:        <math>M \leftarrow \text{AEAD.DEC}(K_e, N, AD \parallel T_k, C)</math>        If <math>M \neq \perp \wedge C[\cdot, N, AD, T_k, C] \neq \perp</math>:          <math>\text{bad}_1 \leftarrow \text{true}</math>          Return <math>M</math>        Else if <math>M \neq \perp</math>: Return <math>M</math>    Return <math>\perp</math></p> <p><u>RF(<math>b, K</math>):</u>  If <math>R[b, K] \neq \perp</math>:    Return <math>R[b, K]</math>  If <math>b = 0</math>:    <math>R[0, K] \leftarrow \{0, 1\}^n</math>  Else:    <math>R[1, K] \leftarrow \{0, 1\}^k</math>  Return <math>R[b, K]</math></p>	<p><math>G_3, G_4</math>:</p> <p><math>j \leftarrow 0; I \leftarrow \emptyset; \text{Col} \leftarrow \emptyset</math>  <math>b \leftarrow \mathcal{A}^{\text{GENHONESTKEY, ENC, DEC}}</math>  Return <math>b</math></p> <p><u>GENHONESTKEY(kid):</u>  <math>K \leftarrow \mathcal{K}</math>  If <math>K \in \text{Col}</math>:    <math>\text{bad}_1 \leftarrow \text{true}</math>    <math>K \leftarrow \mathcal{K} \setminus \text{Col}</math>  <math>\text{Col} \leftarrow \text{Col} \cup \{K\}</math>  <math>\text{kcv} \leftarrow \text{RF}(0, K); K_e \leftarrow \text{RF}(1, K)</math>  <math>I \leftarrow I \cup \{\text{kcv}\}</math>  <math>T[j] \leftarrow (\text{kid}, \text{kcv}, K_e); j \leftarrow j + 1</math></p> <p><u>ENC(id, N, AD, M):</u>  <math>(\text{kid}, \text{kcv}, K_e) \leftarrow T[\text{id}]; T_k \leftarrow \text{kid} \parallel \text{kcv}</math>  <math>C \leftarrow \text{AEAD.ENC}(K_e, N, AD \parallel T_k, M)</math>  <math>C \leftarrow \{0, 1\}^{\text{clen}(M)}</math>  <math>C[\text{id}, N, AD, T_k, C] \leftarrow M</math>  Return <math>(T_k, C)</math></p> <p><u>DEC(<math>\mathbb{K}, N, AD, T_k, C</math>):</u>  For (honest, data) <math>\in \mathbb{K}</math>:    If honest = true:      <math>M \leftarrow C[\text{data}, N, AD, T_k, C]</math>      <math>(\text{kid}, \text{kcv}, K_e) \leftarrow T[\text{data}]</math>      If <math>\text{kid} \parallel \text{kcv} = T_k \wedge M = \perp</math>:        <math>M \leftarrow \text{AEAD.DEC}(K_e, N, AD \parallel T_k, C)</math>        If <math>M \neq \perp</math>:          <math>\text{bad}_2 \leftarrow \text{true}</math>          Continue in loop iteration        Return <math>M</math>      Else if <math>M \neq \perp</math>:        Return <math>M</math>    Else:      <math>(\text{kid}, K) \leftarrow \text{data}</math>      <math>\text{kcv} \leftarrow F_{\text{kcv}}(K); K_e \leftarrow \text{KDF}(K)</math>      If <math>\text{kcv} \in I</math>:        <math>\text{bad}_0 \leftarrow \text{true}</math>        Continue in loop iteration      If <math>\text{kid} \parallel \text{kcv} = T_k</math>:        <math>M \leftarrow \text{AEAD.DEC}(K_e, N, AD \parallel T_k, C)</math>        If <math>M \neq \perp</math>: Return <math>M</math>    Return <math>\perp</math></p> <p><u>RF(<math>b, K</math>):</u>  If <math>R[b, K] \neq \perp</math>:    Return <math>R[b, K]</math>  If <math>b = 0</math>:    <math>R[0, K] \leftarrow \{0, 1\}^n</math>  Else:    <math>R[1, K] \leftarrow \{0, 1\}^k</math>  Return <math>R[b, K]</math></p>
---	--

Fig. 18: Games  $G_0$  through  $G_4$  for the proof of Theorem 7. Games  $G_1$  and  $G_4$  feature the pseudocode that is highlighted, while game  $G_2$  features the pseudocode that is highlighted and the pseudocode that is boxed.

*Proof.* This proof uses the KI-nAE simulator  $S$  shown in Figure 17. We will argue that such a simulator satisfies the KI-nAE-KA definition using  $L_{\text{Enc}}^{\text{anon}}$  for scheme KCV.

Assume by contradiction that  $\mathsf{S}$  does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build the adversaries from the theorem statement using a series of game hops. We provide the pseudocode for games  $G_0$  through  $G_4$  in Figure 18. The proof follows from the following claims:

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr[G_0] - \Pr[G_1]| \leq \mathbf{Adv}_{\mathsf{F}_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B})$
3.  $|\Pr[G_1] - \Pr[G_2]| \leq q_k \cdot \mathbf{Adv}_{\mathsf{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C})$
4.  $|\Pr[G_2] - \Pr[G_3]| \leq \frac{q_k^2}{2^{\kappa+1}}$
5.  $|\Pr[G_3] - \Pr[G_4]| \leq \mathbf{Adv}_{\mathsf{AEAD}}^{\text{mu-nae}}(\mathcal{D})$
6. Game  $G_4$  is equivalent to KI-nAE0 with simulator  $\mathsf{S}$

Since most of the games and claims presented in this proof are nearly identical to those presented in the proof of Theorem 5, for brevity we will only prove Claims 3 and 6.

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $\text{bad}_0$  is set. In  $G_1$  the flag  $\text{bad}_0$  is set when during a call to  $\text{DEC}$  the key check value  $\text{kcv}$  output by some malicious key  $K$  was already output for some honest key during a call to  $\text{GENHONESTKEY}$ . If  $\text{bad}_0$  is set in  $G_2$ ,  $\text{DEC}$  ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } \text{bad}_0].$$

We now bound the rightmost term by providing the following  $\text{PRE}_{\mathsf{F}_{\text{kcv}}}$  adversary  $\mathcal{C}$  such that

$$\Pr[G_2 \text{ sets } \text{bad}_0] \leq q_k \cdot \mathbf{Adv}_{\mathsf{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C}).$$

$\mathcal{C}$  gets as input a random range point  $y$  in  $\{0, 1\}^n$  and runs  $\mathcal{A}$ , simulating game  $G_2$  for  $\mathcal{A}$ .  $\mathcal{C}$  chooses at random to replace some key check value  $\text{kcv}$  with  $y$  during one of  $\mathcal{A}$ 's queries to  $\text{GENHONESTKEY}$ . When  $\text{bad}_0$  is set,  $\mathcal{C}$  immediately halts and returns the malicious key  $K$  as a pre-image for  $y$ . Notice that since there are  $q_k$  such key check values generated, there is probability  $1/q_k$  that the value for which  $\mathcal{A}$  found a pre-image is  $y$ . Therefore, we can lower bound the advantage of  $\mathcal{C}$  by  $\Pr[G_2 \text{ sets } \text{bad}_0] \cdot 1/q_k$ .

**Claim 6:** We can see by inspecting the pseudocode of  $G_4$  that it is almost a syntactic rewriting of game KI-nAE0 with simulator  $\mathsf{S}$ . During a query to  $\text{DEC}$ , however,  $G_4$  decrypts in order of the keys given in  $\mathbb{K}$ , while game KI-nAE0 with  $\mathsf{S}$  first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_4$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_4$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims. □

## F Proofs of Theorem 9 and Theorem 11

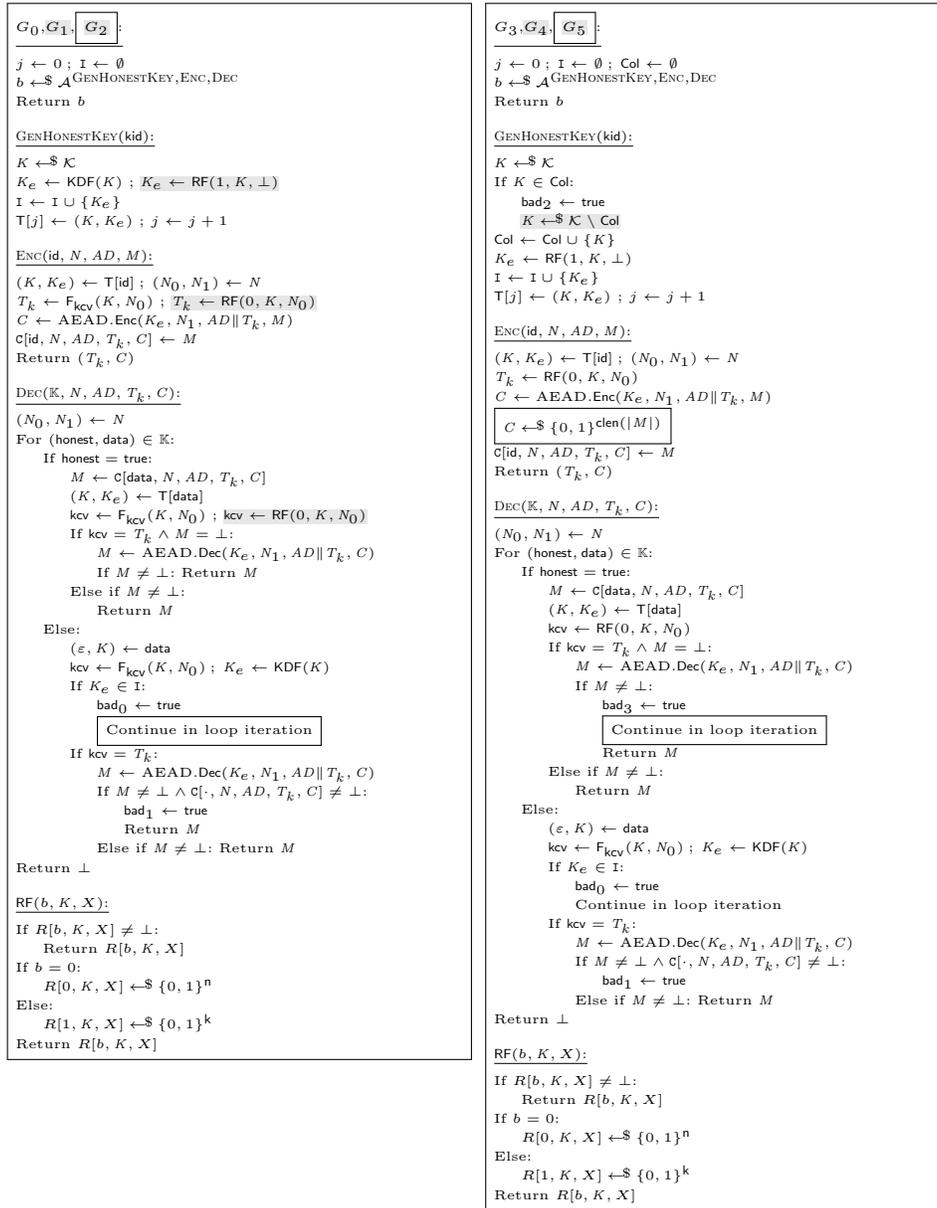


Fig. 19: Games  $G_0$  through  $G_5$  for the proof of Theorem 9. Games  $G_1$  and  $G_4$  feature the pseudocode that is highlighted, while games  $G_2$  and  $G_5$  feature the pseudocode that is highlighted and the pseudocode that is boxed.

### F.1 Proof of Theorem 9

*Proof.* We construct a KI-nAE-KA simulator as shown in Figure 20. We will argue that such a simulator satisfies the KI-nAE-KA definition using  $\mathcal{L}_{\text{Enc}}^{\text{anon}}$  for scheme nKCV.

Assume by contradiction that  $\mathcal{S}$  does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build the adversaries from the theorem statement using a series of game hops. We provide the pseudocode for games  $G_0$  through  $G_5$  in Figure 19. The proof follows from the following claims:

<u><math>\mathcal{S}.\text{Init}()</math></u>	<u><math>\mathcal{S}.\text{Enc}(N, AD, \mathcal{L}, \sigma_s)</math></u>	<u><math>\mathcal{S}.\text{Dec}(\mathbb{K}, N, AD, T_k, C, \sigma_s)</math></u>
$T \leftarrow []$	$\ell \leftarrow \mathcal{L}$	$(N_0, N_1) \leftarrow N$
Return $T$	$C \leftarrow_{\$} \{0, 1\}^{\text{len}(\ell)}$	For $(\text{honest}, \text{data}) \in \mathcal{L}$ :
<u><math>\mathcal{S}.\text{Kg}(\text{id}, \sigma_s)</math></u>	$T_k \leftarrow_{\$} \{0, 1\}^n$	If $\text{honest} = \text{false}$ :
Return $(\varepsilon, \sigma_s)$	Return $(T_k, C, \sigma_s)$	$(\varepsilon, K) \leftarrow \text{data}$
		$\text{kcv} \leftarrow F_{\text{kcv}}(K, N_0); K_e \leftarrow \text{KDF}(K)$
		If $\text{kcv} = T_k$ :
		$M \leftarrow \text{AEAD}.\text{Dec}(K_e, N_1, AD    T_k, C)$
		If $M \neq \perp$ : Return $(M, \sigma_s)$
		Return $(\perp, \sigma_s)$

Fig. 20: The simulator  $\mathcal{S}$  used in the proofs of Theorem 9 and Theorem 11.

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr[G_0] - \Pr[G_1]| \leq \text{Adv}_{F_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B})$
3.  $|\Pr[G_1] - \Pr[G_2]| \leq q_k \cdot \text{Adv}_{\text{KDF}}^{\text{pre}}(\mathcal{C})$
4.  $|\Pr[G_2] - \Pr[G_3]| \leq \text{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{D})$
5.  $|\Pr[G_3] - \Pr[G_4]| \leq \frac{q_k^2}{2^{\kappa+1}}$
6.  $|\Pr[G_4] - \Pr[G_5]| \leq \text{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{E})$
7. Game  $G_5$  is equivalent to KI-nAE0 with simulator  $\mathcal{S}$

**Claim 1:**  $G_0$  is a syntactic rewriting of KI-nAE1 for scheme nKCV. The encryption oracle ENC keeps track of the plaintext that outputs the ciphertext  $(N, AD, T_k, C)$  for the key with internal game identifier id.

The decryption oracle DEC iterates through the key vector to find the first key, if any, that produces a matching key tag. If the key is honest but there is no entry for the ciphertext and key tag in the table, then the game uses AEAD to decrypt. If there is an entry, then the game returns the plaintext stored, which by the correctness of AEAD should match the decryption output. Finally, if the key is malicious, then the game derives the key check value kcv and AEAD encryption key  $K_e$  from the specified key, checks whether the key tag matches kcv, and, if so, uses AEAD to decrypt with  $K_e$ .

**Claim 2:** Game  $G_1$  is identical to  $G_0$  except that all calls to  $F_{\text{kcv}}$  and KDF for honestly generated keys are replaced with calls to random function RF. RF takes as input a bit  $b$ , where  $b = 0$  represents a call to  $F_{\text{kcv}}$  and  $b = 1$  represents a call to KDF and values  $K, X$ . Since KDF never takes a second input, we replace these calls with  $\text{RF}(1, \cdot, \perp)$ . For any new input  $(b, K, X)$ , the function chooses a value uniformly at random from either  $\{0, 1\}^n$  (when  $b = 0$ ) or  $\{0, 1\}^k$  (when  $b = 1$ ), then stores this value in table  $R$ . When RF is called on a repeat input, the value chosen previously and stored in  $R$  is returned.

We can bound the advantage of  $\mathcal{A}$  in distinguishing between  $G_0$  and  $G_1$  by reducing to the MU-PRF security of  $F_{\text{kcv}}$  and KDF; call this reduction  $\mathcal{B}$ . The adversary  $\mathcal{B}$  simulates games  $G_0$  and  $G_1$  by calling GENKEY to generate the AEAD-KI key, Func0 to generate kcv, and Func1 to generate  $K_e$ . The adversary  $\mathcal{B}$  then returns whatever  $\mathcal{A}$  does. When  $\mathcal{B}$  is in the REAL or IDEAL game, then  $\mathcal{B}$  correctly simulates game  $G_0$  or  $G_1$ , respectively, for  $\mathcal{A}$ .

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $\text{bad}_0$  is set. In  $G_1$  the flag  $\text{bad}_0$  is set when during a call to DEC the AEAD encryption key  $K_e$  output by some malicious key  $K$  was already output for some honest key in GENHONESTKEY. If  $\text{bad}_0$  is set in  $G_2$ , DEC ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } \text{bad}_0].$$

We now bound the rightmost term by providing the following  $\text{PRE}_{\text{KDF}}$  adversary  $\mathcal{C}$  such that

$$\Pr[G_2 \text{ sets } \text{bad}_0] \leq q_k \cdot \text{Adv}_{\text{KDF}}^{\text{pre}}(\mathcal{C}).$$

$\mathcal{C}$  gets as input a random range point  $y$  in  $\{0, 1\}^k$  and runs  $\mathcal{A}$ , simulating game  $G_2$  for  $\mathcal{A}$ .  $\mathcal{C}$  chooses at random to replace  $K_e$  with  $y$  during one of  $\mathcal{A}$ 's queries to GENHONESTKEY. When  $\text{bad}_0$  is set,  $\mathcal{C}$  immediately halts and returns the malicious key  $K$  as a pre-image for  $y$ . We abuse notation here and can assume that the second input is the empty string. Notice that since there are  $q_k$  honest keys generated, there is probability

$1/q_k$  that the value for which  $\mathcal{A}$  found a pre-image is  $y$ . Therefore, we can lower bound the advantage of  $\mathcal{C}$  by  $\Pr[G_2 \text{ sets } \mathbf{bad}_0] \cdot 1/q_k$ .

**Claim 4:** Game  $G_3$  is identical to  $G_2$  until the flag  $\mathbf{bad}_1$  is set. In  $G_2$  the flag  $\mathbf{bad}_1$  is set during a call to DEC when a malicious key successfully decrypts a ciphertext that was output by ENC.  $G_2$  returns the plaintext  $M$  that the malicious key decrypts the ciphertext to, while  $G_3$  ignores this result and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_2 \Rightarrow 1] - \Pr[G_3 \Rightarrow 1]| \leq \Pr[G_3 \text{ sets } \mathbf{bad}_1].$$

We now bound the rightmost term by providing the following FROB adversary  $\mathcal{D}$  such that

$$\Pr[G_3 \text{ sets } \mathbf{bad}_1] \leq \mathbf{Adv}_{\text{AEAD}}^{\text{frob}}(\mathcal{D}).$$

$\mathcal{D}$  runs  $\mathcal{A}$  and simulates game  $G_3$  for  $\mathcal{A}$ . If the flag  $\mathbf{bad}_1$  is set to true, then  $\mathcal{D}$  halts and returns the malicious key  $K_e$  used to decrypt the ciphertext, the honest key  $K'_e$  stored in  $\mathbb{T}$  that output  $N, AD, T_k, C$  during a call to ENC, nonce  $N$ , associated data  $AD || T_k$ , and ciphertext  $C$ . Notice that  $K_e$  and  $K'_e$  must be distinct, since both games do not allow for a malicious key to output an AEAD encryption key that was already output by an honest key. Therefore,  $\mathcal{D}$  wins game FROB when  $\mathbf{bad}_1$  is set.

**Claim 5:** Game  $G_4$  differs from  $G_3$  in that during a query to GENHONESTKEY if the key  $K$  chosen at random from  $\mathcal{K}$  has already been chosen on a previous query, then a new fresh value is chosen that is guaranteed to not be a repeat value. Recall that  $q_k$  is the number of queries  $\mathcal{A}$  makes to GENHONESTKEY and  $\mathcal{K} = 2^\kappa$  for this scheme. We can bound  $\mathcal{A}$ 's distinguishing advantage between games  $G_3$  and  $G_4$  by the probability of a collision occurring, which is bounded by the birthday bound  $q_k^2/2^{\kappa+1}$ .

**Claim 6:** Game  $G_5$  differs from  $G_4$  in that during a query to the encryption oracle it chooses a random string as the ciphertext and during a query to the decryption oracle if an honest key was called to decrypt a ciphertext not output by ENC then that key is not used to decrypt. We can bound the advantage of  $\mathcal{A}$  in distinguishing between these games by reducing to the MU-nAE security of the AEAD scheme; call this reduction  $\mathcal{E}$ . The adversary  $\mathcal{E}$  uses its own encryption oracle to simulate queries by  $\mathcal{A}$  to ENC and its decryption oracle to simulate queries by  $\mathcal{A}$  to DEC. When  $\mathcal{E}$  is playing game MU-nAE1 or MU-nAE0, then  $\mathcal{E}$  correctly simulates game  $G_5$  or  $G_4$ , respectively, for  $\mathcal{A}$ .

**Claim 7:** We can see by inspecting the pseudocode of  $G_5$  that it is almost a syntactic rewriting of game KI-nAE0 with simulator  $\mathbb{S}$ . During a query to DEC, however,  $G_5$  decrypts in order of the keys given in  $\mathbb{K}$ , while game KI-nAE0 with  $\mathbb{S}$  first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_5$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_5$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims. □

## F.2 Proof of Theorem 11

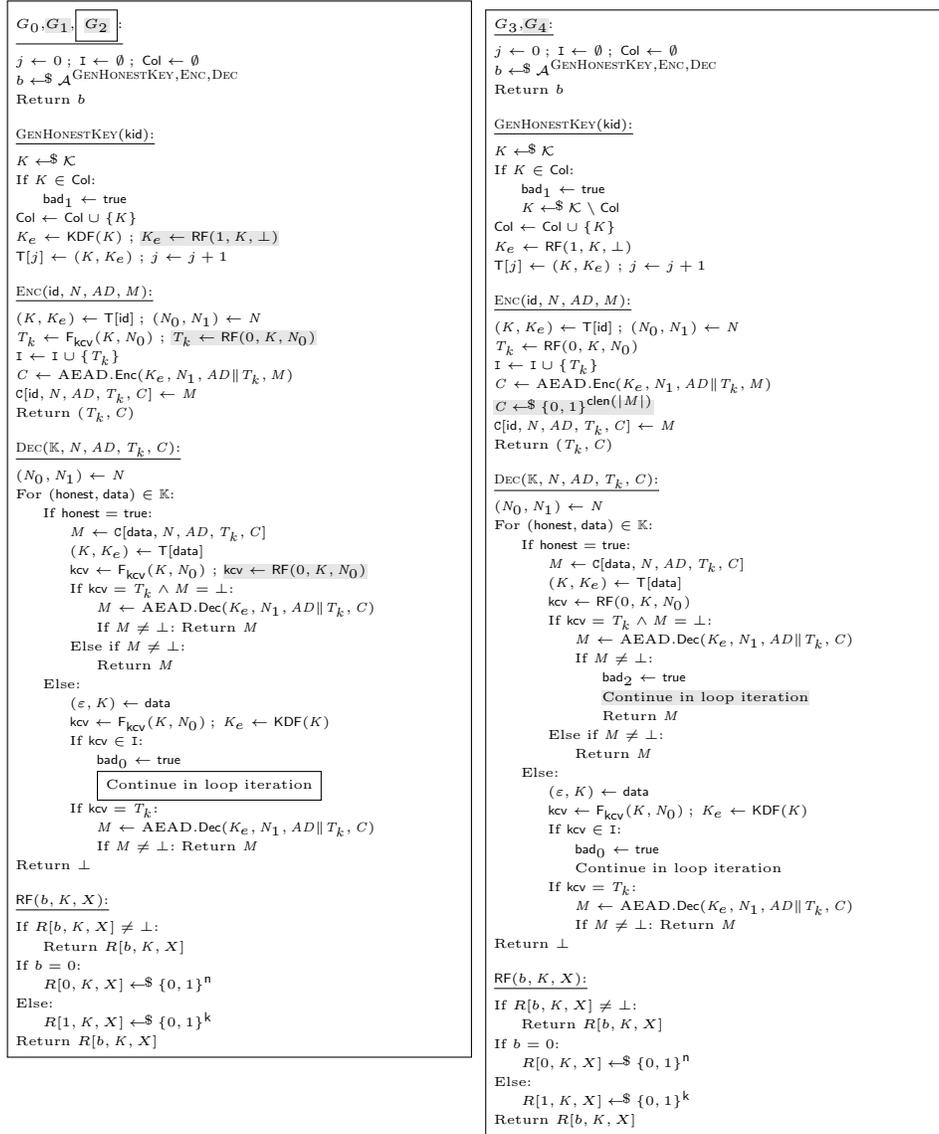


Fig. 21: Games  $G_0$  through  $G_4$  for the proof of Theorem 11. Games  $G_1$  and  $G_4$  feature the pseudocode that is highlighted, while game  $G_2$  features the pseudocode that is highlighted and the pseudocode that is boxed.

*Proof.* We construct the KI-nAE-KA simulator as shown in Figure 20. We will argue that such a simulator satisfies the KI-nAE-KA definition using  $\mathcal{L}_{\text{Enc}}^{\text{anon}}$  for scheme nKCV.

Assume by contradiction that  $\mathcal{S}$  does not satisfy the definition, so there exists an efficient adversary  $\mathcal{A}$  that can distinguish between KI-nAE0 and KI-nAE1. We will use  $\mathcal{A}$  to build the adversaries from the theorem statement using a series of game hops. We provide the pseudocode for games  $G_0$  through  $G_4$  in Figure 21. The proof follows from the following claims:

1. Game  $G_0$  is equivalent to KI-nAE1
2.  $|\Pr[G_0] - \Pr[G_1]| \leq \mathbf{Adv}_{\text{F}_{\text{kcv}}, \text{KDF}}^{\text{mu-prf}}(\mathcal{B})$
3.  $|\Pr[G_1] - \Pr[G_2]| \leq q_e \cdot \mathbf{Adv}_{\text{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C})$

4.  $|\Pr[G_2] - \Pr[G_3]| \leq \frac{q_k^2}{2^{\kappa-1}}$
5.  $|\Pr[G_3] - \Pr[G_4]| \leq \mathbf{Adv}_{\text{AEAD}}^{\text{mu-nae}}(\mathcal{D})$
6. Game  $G_4$  is equivalent to KI-nAE0 with simulator  $\mathcal{S}$

Since most of the games and claims presented in this proof are nearly identical to those presented in the proof of Theorem 9, for brevity we will only prove Claims 3 and 6.

**Claim 3:** Game  $G_2$  is identical to  $G_1$  until the flag  $\mathbf{bad}_0$  is set. In  $G_1$  the flag  $\mathbf{bad}_0$  is set when during a call to DEC the key check value  $\mathbf{kcv}$  output by some malicious key  $K$  and nonce  $N_0$  was already output for some honest key during a call to ENC. If  $\mathbf{bad}_0$  is set in  $G_2$ , DEC ignores this malicious key and keeps iterating through  $\mathbb{K}$ . We can therefore use the fundamental lemma of game-playing [11] to show that

$$|\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1]| \leq \Pr[G_2 \text{ sets } \mathbf{bad}_0].$$

We now bound the rightmost term by providing the following  $\text{PRE}_{\text{F}_{\text{kcv}}}$  adversary  $\mathcal{C}$  such that

$$\Pr[G_2 \text{ sets } \mathbf{bad}_0] \leq q_e \cdot \mathbf{Adv}_{\text{F}_{\text{kcv}}}^{\text{pre}}(\mathcal{C}).$$

$\mathcal{C}$  gets as input a random range point  $y$  in  $\{0, 1\}^n$  and runs  $\mathcal{A}$ , simulating game  $G_2$  for  $\mathcal{A}$ .  $\mathcal{C}$  chooses at random to replace some key tag  $T_k$  with  $y$  during one of  $\mathcal{A}$ 's queries to ENC. When  $\mathbf{bad}_0$  is set,  $\mathcal{C}$  immediately halts and returns the malicious key  $K$  and nonce  $N_0$  as a pre-image for  $y$ . Notice that since there are  $q_e$  such key tags generated, there is probability  $1/q_e$  that the value for which  $\mathcal{A}$  found a pre-image is  $y$ . Therefore, we can lower bound the advantage of  $\mathcal{C}$  by  $\Pr[G_2 \text{ sets } \mathbf{bad}_0] \cdot 1/q_e$ .

**Claim 6:** We can see by inspecting the pseudocode of  $G_4$  that it is almost a syntactic rewriting of game KI-nAE0 with simulator  $\mathcal{S}$ . During a query to DEC, however,  $G_4$  decrypts in order of the keys given in  $\mathbb{K}$ , while game KI-nAE0 with  $\mathcal{S}$  first iterates through all the honest keys and then iterates through the malicious keys if no honest key successfully decrypts.  $G_4$  guarantees that no malicious key can decrypt an honestly generated ciphertext and that no honest key can decrypt a malicious ciphertext. Therefore, in  $G_4$  iterating through  $\mathbb{K}$  in order is identical to iterating through the honest keys first and then the malicious keys, proving our claim.

The theorem statement thus follows from our claims. □