





Symmetric Signcryption and E2EE Group Messaging in Keybase

Joseph Jaeger¹ , Akshaya Kumar¹ , and Igors Stepanovs² 

¹ School of Cybersecurity and Privacy, Georgia Institute of Technology,
Atlanta, GA, USA

{josephjaeger, akshayakumar}@gatech.edu

² Riga, Latvia

igors.stepanovs@gmail.com

<https://cc.gatech.edu/~josephjaeger/>, <https://cc.gatech.edu/~akumar805/>,
<https://igors.org/>

Abstract. We introduce a new cryptographic primitive called symmetric signcryption, which differs from traditional signcryption because the sender and recipient share a secret key. We prove that a natural composition of symmetric encryption and signatures achieves strong notions of security against attackers that can learn and control many keys. We then identify that the core encryption algorithm of the Keybase encrypted messaging protocol can be modeled as a symmetric signcryption scheme. We prove the security of this algorithm, though our proof requires assuming non-standard, brittle security properties of the underlying primitives.

1 Introduction

Keybase is a suite of encryption tools. It encompasses a public-key directory, an instant messenger, and a cloud storage service. Keybase was launched in 2014. In February 2020, it reported having accumulated more than 1.1M user accounts [27]. In May 2020, Keybase was acquired by Zoom. At the time, Zoom issued a public statement [36] saying that the Keybase’s team was meant to play a critical part in building scalable end-to-end encryption for Zoom. The acquisition appears to have put an end to an active development of new Keybase features, but as of February 2024 it keeps receiving regular maintenance updates.

Instant Messaging in Keybase. Keybase implements its own end-to-end encrypted instant messaging protocol. This protocol is designed to support large groups. One-on-one chats are treated as group chats and hence use the same protocol. The protocol also allows to send large files as encrypted attachments in chat. It is impossible to opt out of end-to-end encryption in Keybase. In this work we analyze the security of this protocol.

The Keybase client is open source [24], but the server is not. Our security analysis primarily relies on the source code. Keybase also provides the “Keybase Book” website [22] with excellent documentation that explains its cryptographic design. The only prior security analysis of Keybase was done by NCC Group in

2019 [31], which broadly looked at the security of the entire Keybase ecosystem. In comparison, we provide an in-depth analysis of a single component in Keybase.

Encrypted Group Chats. In this work we consider a setting in which an arbitrary number of users can form a group. All group members share a key for a symmetric encryption scheme. Each instant message within the group is encrypted with this key. Let us use g to denote the identity of a group and K_g to denote the key shared between the members of this group. In Keybase, every member of group g uses the same long-term key K_g to encrypt their outgoing chat messages. Each message is encrypted only once, simultaneously for all recipients. The resulting ciphertext is then broadcast to all members of the group.

The *Sender Keys* protocol [7, 28] can be seen as building on this basic design idea. In *Sender Keys*, every member of the group owns a distinct symmetric encryption key; they share it with other group members. Each outgoing message is encrypted with the sender’s own key, and the resulting ciphertext is broadcast to the group. Furthermore, each key is used to encrypt only a single message, and immediately afterwards a new key is derived to be used for the next encryption. So every group member tracks every other member’s current encryption key, decrypting each incoming ciphertext with the corresponding sender’s key and subsequently replacing it with an appropriately derived new key. Variants of the *Sender Keys* protocol are used in the *Signal* [28], *WhatsApp* [34], and *Matrix* [1, 2] messengers. In addition, the *Messaging Layer Security* (MLS) [8] protocol contains a component called *FS-GAEAD* [3] or *TreeDEM* [33] that similarly uses a sender’s key to encrypt and broadcast a message (but its overall design significantly differs from design of the *Sender Keys* protocol).

An encrypted group chat protocol should provide at least confidentiality and integrity of communication, with respect to an attacker that is not a member of the group. In part, this could be achieved by building the protocol from a symmetric encryption scheme that satisfies some notion of authenticated-encryption security. But care is needed to also prevent undesired message replays, reordering, or drops. These requirements are specific to a stateful protocol and do not necessarily follow from properties provided by the underlying stateless scheme.

Sender Authentication in Group Chats. Consider a group chat protocol that is built from a single symmetric encryption scheme and where every symmetric key is known to all group members. In such a protocol, group members are able to impersonate each other. This is true regardless of whether each group uses a single shared encryption key or has each member own a distinct encryption key. To prevent group members from impersonating each other, it is natural to use a digital signature scheme. Let us use u to denote the identity of a user and sk_u to denote this user’s signing key for a digital signature scheme.

What is a sound way to compose a symmetric encryption scheme with a digital signature scheme? Let us consider two sequential compositions of a signing algorithm *Sign* with an encryption algorithm *Encrypt*. We call the resulting schemes *Sign-then-Encrypt* and *Encrypt-then-Sign*, and we show them in Fig. 1. The *Sign-then-Encrypt* scheme first signs a message m to obtain its digital signature s and then encrypts (s, m) to obtain and return a symmetric ciphertext




Fig. 1. Warmup schemes obtained by composing digital signatures with symmetric encryption. **Left pane:** Sign-then-Encrypt. **Right pane:** Encrypt-then-Sign.




Fig. 2. A high-level representation of the SealPacket scheme in Keybase.

c . The Encrypt-then-Sign scheme first encrypts m as c and then computes a signature s over c ; it returns the pair (c, s) .

These compositions closely mirror those that are commonly used to build signcryption [5], which is a standard cryptographic primitive that combines digital signatures with public-key encryption [35], except we replace public-key encryption with symmetric encryption. It is well known that the corresponding compositions for signcryption are not secure in the multi-user setting, unless some effort is taken to bind together the message with the sender and recipient identities [5]. The standard advice is to always sign the recipient’s identity and always encrypt the sender’s identity. Our basic schemes in Fig. 1 would intuitively suffer from similar issues and benefit from similar countermeasures. However, the exact details would depend on what kind of security one expects from these schemes, so we defer this discussion.

The *Sender Keys* protocol [7] prescribes to sign a symmetric ciphertext; and this is indeed done by the *Signal*, *WhatsApp*, and *Matrix* messengers. The MLS protocol [8] protocol prescribes to encrypt a digital signature with a symmetric encryption scheme. So either protocol can be seen as using some variant of Encrypt-then-Sign or Sign-then-Encrypt as a subroutine. We will now discuss that Keybase can be seen as extending both of these basic schemes.

SealPacket: Sign-then-Encrypt in Keybase. Keybase uses a variant of the basic Sign-then-Encrypt scheme. It signs the symmetric encryption key along with the plaintext, meaning it signs (K_g, m) instead of just m . The resulting scheme is called SealPacket and is shown in Fig. 2. In the source code, the decision to sign K_g is explained as follows [26]:

simply using encryption and signing together isn’t good enough . . . the inner layer needs to assert something about the outer layer . . . a better approach is to mix the outer key into the inner crypto, so that it’s impossible to forget to check it . . . That means the inner signing layer needs to assert




Fig. 3. A high-level representation of the BoxMessage scheme in Keybase.

the encryption key ... We don't need to worry about whether the signature might leak the encryption key either, because the signature gets encrypted.

Keybase uses SealPacket to encrypt the following three types of plaintexts: (1) a metadata header that is automatically created and sent along with every chat message, (2) a file that is sent as an attachment in chat, and (3) an arbitrary string chosen by a chat bot (for secure server-side storage of bot data).

BoxMessage: Encrypt-then-Sign in Keybase. The BoxMessage scheme in Keybase is a variant of the basic Encrypt-then-Sign scheme. This scheme, unlike SealPacket, is used only for one purpose: to encrypt the body of a chat message. So we denote by c_{body} the symmetric ciphertext that is created in the inner (encryption) layer of BoxMessage. The BoxMessage scheme extends the basic Encrypt-then-Sign scheme in two ways and is shown in Fig. 3. First, it takes an associated-data field ad and signs (c_{body}, ad) instead of just c_{body} . Second, rather than use a signature scheme, BoxMessage uses SealPacket to sign (c_{body}, ad) .

Keybase uses the auxiliary-data field ad to authenticate a metadata header for the chat message. This header contains the group's identity and sender's identity among multiple other values. The data in ad is sent in plain over the network (along with c_{body}), meaning that SealPacket is not meant to provide confidentiality of ad . Indeed, the Keybase documentation explains that SealPacket is used to provide the confidentiality of the signature over (c_{body}, ad) [23]:

fields in the header aren't secret from the server, and it actually needs to know several of them ... The reason for sign-then-encrypting/signencrypting the header is instead to keep the signature itself private. Even though the server knows who's talking to whom, because it's delivering all the messages, it's better that it can't prove what it knows.

Interestingly, BoxMessage reuses the group's symmetric encryption key K_g between its calls to Encrypt and SealPacket. As mentioned above, SealPacket will itself first sign K_g and then run another instance of Encrypt with K_g as the key. In total, the same value of K_g is therefore used in 3 distinct contexts.

Symmetric Signcryption. We define *symmetric signcryption* as a new cryptographic primitive that combines symmetric encryption with digital signatures. We capture the setting where every user owns a signing key pair and in each group all users share a single symmetric encryption key. The encryption key is long-term, meaning it can be used an arbitrary number of times, simultaneously

by all members of the group. This will allow us to formalize and analyze the `SealPacket` and `BoxMessage` schemes.

We note that the use of sender-specific encryption keys in the *Sender Keys* [7] and MLS [8] protocols can also be captured by symmetric signcryption. Indeed, in either protocol all symmetric encryption keys can be seen as being independently sampled, and each individual key is used only once. This can be thought of as a collection of “one-time-use” symmetric signcryption schemes.

We adapt the standard syntax of (asymmetric) signcryption to suit our setting, defining algorithms `SigEnc` and `VerDec`. They both take explicit sender and group identities, nonces, and associated data.

Security of Symmetric Signcryption. We define two security notions for symmetric signcryption. The *out-group authenticated encryption* (OAE) security requires confidentiality and integrity of communication against an adversary that does not know the symmetric key of the group it attacks. This is required to hold even against an adversary that can assign to every user an arbitrary (possibly malformed) digital signature key pair. The *in-group unforgeability* (IUF) security requires unforgeability of messages sent by users whose signing keys are not known to the adversary. This is required to hold even against an adversary that can assign to every group an arbitrary (possibly malformed) symmetric key.

We show how to extend the basic Sign-then-Encrypt scheme, by carefully incorporating user and group identifiers, to achieve both of our security notions. We assume strong unforgeability of the underlying digital signature scheme and authenticated-encryption security of the underlying encryption scheme.

Implementation of `SealPacket` and `BoxMessage`. In the source code of the Keybase client, `SealPacket` is implemented in [26] and `BoxMessage` is implemented in [25]. These schemes are instantiated with the nonce-based authenticated encryption scheme XSalsa20-Poly1305 [15,16] and the digital signature scheme Ed25519 [17,18]. They also use SHA-512 and SHA-256 which does not significantly affect the design of either scheme so we omit discussing it here, but in the main body of the paper, we attempt to formalize both schemes precisely.

Keybase implements four versions of the `BoxMessage` scheme: V1, V2, V3, and V4. V1 is deprecated; the Keybase client allows to receive but not send messages that use V1. V2 is the default version that we formalize and analyze in this work. V3 is the same as V2, except it supports exploding messages; the body of an exploding message is encrypted using an ephemeral key instead of K_g . V4 is the same as V3, except it makes all group members use a dummy (zero) signing key and instead authenticate messages using pairwise MACs.

Provable Security Analysis. We model `BoxMessage` and `SealPacket` as symmetric signcryption schemes and provide formal reductions for their IUF and OAE security. Our analysis is done in a concrete security framework [11], and in a multi-key setting; we state precise bounds on the advantage of an attacker. The analysis of `BoxMessage` largely encompasses that of `SealPacket`, because `BoxMessage` uses `SealPacket` in a modular way. So we focus on the analysis of `BoxMessage` here. The main challenges arise from using K_g in 3 distinct contexts.

First, we aim to show it is hard to switch the context of the XSalsa20-Poly1305 ciphertexts c_{body} and c_{header} . Both are encrypted using the same key K_g , so there is a risk that an attacker could forge a valid encryption of some body-plaintext m from a known encryption of some header-plaintext (s, c_{body}, ad) , or vice versa. To rule out such attacks, we rely on an observation that every application-layer message m that is queried to be encrypted by `BoxMessage` is encoded in a specific way, whereas every header-plaintext (s, c_{body}, ad) is expected to start with a valid Ed25519 signature. Based on the specification of Ed25519 we show (in the ROM and GGM) that it is hard to cast the encoding used in m as a valid Ed25519 signature (with respect to any verification key of adversary’s choice).

Second, we need to show that XSalsa20-Poly1305 provides authenticated encryption even for certain messages derived from its secret key. This arises because in `SealPacket` the XSalsa20-Poly1305 key K_g is first signed with Ed25519 and then the resulting signature is encrypted using XSalsa20-Poly1305 under the same key. Here again we rely on the specification of Ed25519. An Ed25519 signature depends on two SHA-512 hash values of the message that is being signed, but it does not depend on the signed message beyond that. We use this (in the ROM) to eliminate the need to consider key-dependent messages and hence only require XSalsa20-Poly1305 to provide the standard notion of authenticated encryption.

We do not know any way to avoid the above analysis. The necessity to use non-standard security notions appears to be inherently implied by the design decisions made in `Keybase`. This could have been avoided (e.g. with out `Sign-then-Encrypt` scheme). Overall, our reductions (in the ROM and GGM) rely on the AEAD security of XSalsa20-Poly1305, collision resistance of SHA-256 and SHA-512, and strong unforgeability of Ed25519. We note that `Keybase` uses the version of Ed25519 that was recently shown to be SUF-CMA secure [10, 20].

Limitations of Our Work. Our analysis of `Keybase` is intentionally narrow in scope. We perform an in-depth, algorithmic analysis of specific chat components that can be modeled as symmetric signcryption. Other analysis is outside the scope of our work, such as whether these algorithms are secure against timing attacks and whether they provide protection against message replays, reordering, or drops when used within the broader stateful chat protocol. More broadly, our analysis does not explicitly cover many other applications of cryptography in `Keybase`, including other versions of `BoxMessage`, encryption of attachments or bot data, the initial key exchange used to agree on group keys, the public-key directory used to share user keys, and the cloud storage service. These applications are important for the overall security of `Keybase`, and have the potential to interplay with each other in subtle ways. For example, user signing keys are used for multiple tasks in `Keybase`. We believe appropriate context separation is used for these purposes (e.g. all messages signed in `SealPacket` start with “`Keybase-Chat-2`”). If not, subtle cross-application attacks may be possible.

Related Work. The Hybrid Public-Key Encryption (HPKE) standard is specified in RFC 9180 [9]. Alwen, Janneck, Kiltz, and Lipp [4] analyze the “pre-shared key” modes from RFC 9180. They cast the $\text{HPKE}_{\text{AuthPSK}}$ mode as an asymmetric

signcryption scheme that is augmented with a pre-shared symmetric key, and they define the corresponding security notions. They analyze the security that is achieved by $\text{HPKE}_{\text{AuthPSK}}$ depending on which combinations of keys are secure. Our definitions are similar in the sense that both works define a signcryption-type primitive that in addition uses a symmetric key. However, the algorithms in [4] use one more set of keys, and the definitions in [4] are stated in the two-user setting. In essence, our primitives are similar in form, but are tailored to be used as tools in different settings.

2 Preliminaries

We use standard pseudocode notation and assume familiarity with hash functions, random oracles, nonce-based encryption, and digital signatures. Collision resistance is defined by $\text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}_{\text{CR}}) = \Pr[\text{H}(x) = \text{H}(y) \wedge x \neq y : (x, y) \leftarrow_{\$} \mathcal{A}_{\text{CR}}]$.

2.1 Standard Security Notions in a Multi-key Setting

Key Management Oracles. Throughout this work we consider multi-key security notions. Adversaries in security games will be provided with three types of key management oracles. These oracles will allow (1) sampling new honest (i.e. challenge) keys, (2) exposing existing honest keys, and (3) adding corrupt keys of the adversary’s choice. When an honest key is exposed it becomes corrupt, but it was initially sampled from a correct key distribution. In contrast, when an adversary adds its own corrupt key, such a key could be maliciously crafted in an arbitrary way. In basic security notions an adversary cannot benefit from crafting corrupt keys, because no challenge queries are permitted with respect to such keys. This changes for more complex systems built from more than one keyed primitive when some security is required to hold even if some underlying secrets are exposed. The ability to use malicious keys was modeled in prior work on (asymmetric) signcryption [14], and will be needed in this work.

Our security model for symmetric signcryption in Sect. 3 will define two sets of key management oracles. The set of *user oracles* $\text{U} = \{\text{NEWHONUSER}, \text{EXPOSEUSER}, \text{NEWCORRUSER}\}$ will manage the keys for a digital signature scheme, whereas the set of *group oracles* $\text{G} = \{\text{NEWHONGROUP}, \text{EXPOSEGROUP}, \text{NEWCORRGROUP}\}$ will manage the keys for a nonce-based encryption scheme. We adopt the same terminology and notation across all of the multi-key security notions; each notion for an *asymmetric* primitive will define a set of user oracles U , and each notion for a *symmetric* primitive will define a set of group oracles G . For consistency, we include oracles for adding corrupt keys even when an adversary cannot benefit from using them. When simulating user oracles in a security reduction, we write SIMU to denote the set $\{\text{SIMNEWHONUSER}, \text{SIMEXPOSEUSER}, \text{SIMNEWCORRUSER}\}$ and do similarly for group oracles.

Nonce-Based Authenticated Encryption. Consider game $\mathcal{G}^{\text{AEAD}}$ of Fig. 4 for nonce-based encryption scheme NE and adversary $\mathcal{A}_{\text{AEAD}}$. The advantage of

$\mathcal{A}_{\text{AEAD}}$ in breaking the AEAD security of NE is defined as $\text{Adv}_{\text{NE}}^{\text{AEAD}}(\mathcal{A}_{\text{AEAD}}) = 2 \cdot \text{Pr}[\mathcal{G}_{\text{NE}}^{\text{AEAD}}(\mathcal{A}_{\text{AEAD}})] - 1$. The game samples a challenge bit b , and $\mathcal{A}_{\text{AEAD}}$ is required to guess it. Adversary $\mathcal{A}_{\text{AEAD}}$ is given the group oracles G , encryption oracle ENC , and decryption oracle DEC . Among the group oracles, NEWHONGROUP creates new groups with honestly generated NE keys, EXPOSEGROUP reveals the keys of existing groups, and NEWCORRGROUP instantiates new corrupt groups with NE keys of $\mathcal{A}_{\text{AEAD}}$'s choice. We require NE to be nonce-misuse resistant [30], meaning that no challenge message m is allowed to be queried across two distinct calls to ENC with respect to the same set of g, n, ad . A corrupt group key can only be used to call ENC with $m_0 = m_1$, and a group key cannot be exposed after it has been used in ENC with $m_0 \neq m_1$. The decryption oracle DEC takes g, n, c, ad as input and decrypts this to the corresponding plaintext m . Following the all-in-one style of [30, 32], it returns \perp if $b = 0$, and it returns m otherwise. This oracle never decrypts a ciphertext with an exposed group's key, and it never decrypts ciphertexts previously produced by ENC (with the same g, c, ad).

<p>Game $\mathcal{G}_{\text{NE}}^{\text{AEAD}}(\mathcal{A}_{\text{AEAD}})$</p> <p>$b \leftarrow \{0, 1\}$; $b' \leftarrow \mathcal{A}_{\text{AEAD}}^{\text{G,ENC,DEC}}$; Return $b = b'$</p> <p>$\text{ENC}(g, n, m_0, m_1, ad)$</p> <p>require $\text{K}[g] \neq \perp$ and $m_0 = m_1$</p> <p>require $\forall d \in \{0, 1\}, (g, n, m_d, ad) \notin N_d$</p> <p>If $m_0 \neq m_1$ then</p> <p> If $\text{group.is.corrupt}[g]$ then return \perp</p> <p> $\text{chal}[g] \leftarrow \text{true}$</p> <p> $c \leftarrow \text{NE.Enc}(\text{K}[g], n, m_b, ad)$</p> <p> $N_0 \leftarrow N_0 \cup \{(g, n, m_0, ad)\}$</p> <p> $N_1 \leftarrow N_1 \cup \{(g, n, m_1, ad)\}$</p> <p> $C \leftarrow C \cup \{(g, n, c, ad)\}$; Return c</p> <p>$\text{DEC}(g, n, c, ad)$</p> <p>require $\text{K}[g] \neq \perp$ and $\neg \text{group.is.corrupt}[g]$</p> <p>require $(g, n, c, ad) \notin C$</p> <p>$m \leftarrow \text{NE.Dec}(\text{K}[g], n, c, ad)$</p> <p>If $b = 0$ then return \perp else return m</p>		
<p>Game $\mathcal{G}_{\text{NE}}^{\text{KR}}(\mathcal{A}_{\text{KR}})$</p> <p>$\mathcal{A}_{\text{KR}}^{\text{G,ENC,DEC,GUESS}}$; Return win</p> <p>$\text{ENC}(g, n, m, ad)$</p> <p>require $\text{K}[g] \neq \perp$</p> <p>require $(g, n, m, ad) \notin N$</p> <p>$c \leftarrow \text{NE.Enc}(\text{K}[g], n, m, ad)$</p> <p>$N \leftarrow N \cup \{(g, n, m, ad)\}$; Return c</p> <p>$\text{DEC}(g, n, c, ad)$</p> <p>require $\text{K}[g] \neq \perp$</p> <p>$m \leftarrow \text{NE.Dec}(\text{K}[g], n, c, ad)$</p> <p>Return m</p> <p>$\text{GUESS}(K)$</p> <p>If $\exists g: (\text{K}[g] = K$ and $\neg \text{group.is.corrupt}[g])$ then</p> <p> win $\leftarrow \text{true}$</p>		
<p>$\text{NEWHONGROUP}(g)$</p> <p>require $\text{K}[g] = \perp$</p> <p>$\text{K}[g] \leftarrow \{0, 1\}^{\text{NE.kl}}$</p>	<p>$\text{EXPOSEGROUP}(g)$</p> <p>require $\text{K}[g] \neq \perp$ and $\neg \text{chal}[g]$</p> <p>$\text{group.is.corrupt}[g] \leftarrow \text{true}$</p> <p>Return $\text{K}[g]$</p>	<p>$\text{NEWCORRGROUP}(g, K)$</p> <p>require $\text{K}[g] = \perp$</p> <p>$\text{group.is.corrupt}[g] \leftarrow \text{true}$</p> <p>$\text{K}[g] \leftarrow K$</p>

Fig. 4. Left pane: Game defining authenticated-encryption security of a nonce-based encryption scheme NE. **Right pane:** Game defining key-recovery security of NE. **Bottom pane:** Group oracles $\mathsf{G} = \{\text{NEWHONGROUP}, \text{EXPOSEGROUP}, \text{NEWCORRGROUP}\}$ that are provided to an adversary in either game, except that the boxed code only appears in the AEAD security game.

Key-Recovery Security of NE. Consider game \mathcal{G}^{KR} of Fig. 4 for nonce-based encryption scheme NE and adversary \mathcal{A}_{KR} . The advantage of \mathcal{A}_{KR} in breaking the KR security of NE is defined as $\text{Adv}_{\text{NE}}^{\text{KR}}(\mathcal{A}_{\text{KR}}) = \Pr[\mathcal{G}_{\text{NE}}^{\text{KR}}(\mathcal{A}_{\text{KR}})]$.

Game $\mathcal{G}_{\text{DS}}^{\text{SUF CMA}}(\mathcal{A}_{\text{SUF CMA}})$		$\text{SIGN}(u, m)$
$(u, m, s) \leftarrow_{\$} \mathcal{A}_{\text{SUF CMA}}^{\text{U, SIGN}}$; If $\text{vk}[u] = \perp$ then return false		require $\text{sk}[u] \neq \perp$
$\text{win}_1 \leftarrow \neg \text{user_is_corrupt}[u]$; $\text{win}_2 \leftarrow ((u, m, s) \notin S)$		$s \leftarrow_{\$} \text{DS.Sig}(\text{sk}[u], m)$
$\text{win}_3 \leftarrow \text{DS.Ver}(\text{vk}[u], m, s)$		$S \leftarrow S \cup \{(u, m, s)\}$
Return win_1 and win_2 and win_3		Return s
$\text{NEWHONUSER}(u)$	$\text{EXPOSEUSER}(u)$	$\text{NEWCORRUSER}(u, \text{sk}, \text{vk})$
require $\text{sk}[u] = \text{vk}[u] = \perp$	require $\text{sk}[u] \neq \perp$	require $\text{sk}[u] = \text{vk}[u] = \perp$
$(\text{sk}[u], \text{vk}[u]) \leftarrow_{\$} \text{DS.Kg}$	$\text{user_is_corrupt}[u] \leftarrow \text{true}$	$\text{user_is_corrupt}[u] \leftarrow \text{true}$
Return $\text{vk}[u]$	Return $\text{sk}[u]$	$\text{sk}[u] \leftarrow \text{sk}$; $\text{vk}[u] \leftarrow \text{vk}$

Fig. 5. Game defining strong unforgeability of a digital signature scheme DS, where $U = \{\text{NEWHONUSER}, \text{EXPOSEUSER}, \text{NEWCORRUSER}\}$.

Strong Unforgeability of Digital Signatures. Consider game $\mathcal{G}^{\text{SUF CMA}}$ of Fig. 5 for signature scheme DS and adversary $\mathcal{A}_{\text{SUF CMA}}$. The advantage of $\mathcal{A}_{\text{SUF CMA}}$ in breaking the SUFCMA security of DS is defined as $\text{Adv}_{\text{DS}}^{\text{SUF CMA}}(\mathcal{A}_{\text{SUF CMA}}) = \Pr[\mathcal{G}_{\text{DS}}^{\text{SUF CMA}}(\mathcal{A}_{\text{SUF CMA}})]$.

3 Symmetric Signcryption

In this section we define syntax and security for multi-user symmetric signcryption. In symmetric signcryption, a user encrypts messages using their signing key and a symmetric key shared by a group of users. We want that nobody outside a group can learn what messages are being encrypted, and nobody at all can forge a message as having come from someone other than themselves.

Syntax. A symmetric signcryption scheme SS specifies algorithms SS.UserKg , SS.SigEnc , SS.VerDec , where SS.VerDec is deterministic. These algorithms use syntax $(\text{sk}, \text{vk}) \leftarrow_{\$} \text{SS.UserKg}$, $c \leftarrow_{\$} \text{SS.SigEnc}(g, K_g, u, \text{sk}_u, n, m, ad)$, and $m \leftarrow \text{SS.VerDec}(g, K_g, u, \text{vk}_u, n, c, ad)$. Associated to SS is a group-key length $\text{SS.gkl} \in \mathbb{N}$, a nonce space SS.NS , a plaintext space $\text{SS.MS} \subseteq \{0, 1\}^*$, and an associated-data space SS.AD . The user's key generation algorithm SS.UserKg returns a key pair (sk, vk) where sk is a signing key and vk is the corresponding verification key. The signcryption algorithm SS.SigEnc takes a group's identifier $g \in \{0, 1\}^*$ and its symmetric key $K_g \in \{0, 1\}^{\text{SS.gkl}}$, a sender's identifier $u \in \{0, 1\}^*$ and its signing key sk_u , a nonce $n \in \text{SS.NS}$, a plaintext $m \in \text{SS.MS}$, and associated data $ad \in \text{SS.AD}$; it returns a signcryption ciphertext c . The deterministic unsigncryption algorithm SS.VerDec takes $g, K_g, u, \text{vk}_u, n, c, ad$, where vk_u is the

verification key of the sender u ; it returns a plaintext $m \in \{0, 1\}^* \cup \{\perp\}$, where \perp indicates a failure to recover a plaintext. We say that SS is *deterministic* if SS.SigEnc is deterministic. Correctness is defined in the natural way.

3.1 In-Group Unforgeability

The strongest variant of in-group unforgeability requires that an attacker cannot modify anything about ciphertexts. We also capture weaker variants. For example, the `SealPacket` encryption algorithm in `Keybase` (as defined in Sect. 4) uses a signing key to bind its ciphertexts to a group’s symmetric key but not to a group’s identifier. So we parameterize our security definition in order to capture the type of authenticity that is as restrictive as possible except for allowing (what can be described as) cross-group forgeries.

IUF Game. Consider game \mathcal{G}^{IUF} of Fig. 6, defined for symmetric signcryption scheme SS , ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$, and adversary \mathcal{A}_{IUF} . The advantage of \mathcal{A}_{IUF} in breaking the IUF security of SS is defined as $\text{Adv}_{\text{SS}, \text{pred}_{\text{trivial}}^{\text{auth}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF}}) = \text{Pr}[\mathcal{G}_{\text{SS}, \text{pred}_{\text{trivial}}^{\text{auth}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF}})]$. Adversary \mathcal{A}_{IUF} is given access to user oracles U , group oracles G , encryption oracle SIGENC , and decryption oracle VERDEC . Its goal is to set the win flag by forging a ciphertext for an honest user.

Among user oracles, `NEWHONUSER` creates honest users with honestly generated signing keys, `NEWCORRUSER` creates corrupt users with malicious signing keys, and `EXPOSEUSER` exposes signing keys of existing users. Among group oracles, `NEWHONGROUP` creates honest groups with honestly sampled symmetric keys, `NEWCORRGROUP` creates corrupt groups with malicious symmetric keys, and `EXPOSEGROUP` exposes symmetric keys of existing groups. Oracles `NEWHONGROUP` and `NEWCORRGROUP` take as input a set `users` identifying the new group’s users; the encryption and decryption oracles then disallow queries that match a group to a non-member user. The user and group oracles use tables `user_is_corrupt` and `group_is_corrupt` in order to keep track of the users and groups whose keys are not secure, respectively. The IUF game never checks `group_is_corrupt`, deliberately giving the adversary full control over group keys.

The encryption oracle `SIGENC` takes (g, u, n, m, ad) and returns a ciphertext c that is produced by running $\text{SS.SigEnc}(g, \text{K}[g], u, \text{sk}[u], n, m, ad)$. Here note that the group and user keys $\text{K}[g]$ and $\text{sk}[u]$ are the only two inputs to SS.SigEnc that are not directly chosen by the adversary at the moment of querying the `SIGENC` oracle. At the end of each `SIGENC` query, the set C is updated to add the tuple $((g, u, n, m, ad), c)$ that can be interpreted as containing the input-output transcript of this query.

The decryption oracle `VERDEC` takes (g, u, n, c, ad) and returns the message m that is recovered by running $\text{SS.VerDec}(g, \text{K}[g], u, \text{vk}[u], n, c, ad)$. Keys $\text{K}[g], \text{sk}[g]$ are the only inputs to SS.VerDec not directly chosen by the adversary. If $m \neq \perp$, then the oracle determines if the current oracle query is a valid forgery and sets the win flag if so. In particular, `VERDEC` builds the tuple $z = ((g, u, n, m, ad), c)$ with all input and output values of the current decryption query. It checks z against the set C that contains the input-output behavior

<p>Game $\mathcal{G}_{SS, \text{pred}_{\text{trivial}}^{\text{auth}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF}})$</p> <p>$\mathcal{A}_{\text{IUF}}^{U, G, \text{SIGENC}, \text{VERDEC}}$</p> <p>Return win</p> <p><u>SIGENC(g, u, n, m, ad)</u></p> <p>require $K[g] \neq \perp$ require $sk[u] \neq \perp$ and $u \in \text{members}[g]$ $c \leftarrow \text{SS.SigEnc}(g, K[g], u, sk[u], n, m, ad)$ $C \leftarrow C \cup \{((g, u, n, m, ad), c)\}$ Return c</p> <p><u>VERDEC(g, u, n, c, ad)</u></p> <p>require $K[g] \neq \perp$ require $vk[u] \neq \perp$ and $u \in \text{members}[g]$ $m \leftarrow \text{SS.VerDec}(g, K[g], u, vk[u], n, c, ad)$ If $m = \perp$ then return \perp $z \leftarrow ((g, u, n, m, ad), c)$ If $\text{pred}_{\text{trivial}}^{\text{auth}}(z, C)$ then return m If $\neg \text{user_is_corrupt}[u]$ then win $\leftarrow \text{true}$ Return m</p>	<p>Game $\mathcal{G}_{SS, \text{pred}_{\text{trivial}}^{\text{sec}}, \text{func}_{\text{out}}^{\text{sec}}}^{\text{OAE}}(\mathcal{A}_{\text{OAE}})$</p> <p>$b \leftarrow \text{ss } \{0, 1\}$; $b' \leftarrow \text{ss } \mathcal{A}_{\text{OAE}}^{U, G, \text{SIGENC}, \text{VERDEC}}$</p> <p>Return $b = b'$</p> <p><u>SIGENC(g, u, n, m_0, m_1, ad)</u></p> <p>require $K[g] \neq \perp$ and $m_0 = m_1$ require $sk[u] \neq \perp$ and $u \in \text{members}[g]$ require $\forall d \in \{0, 1\}, (g, u, n, m_d, ad) \notin N_d$ If $m_0 \neq m_1$ then If $\text{group_is_corrupt}[g]$ then return \perp $\text{chal}[g] \leftarrow \text{true}$ $c \leftarrow \text{SS.SigEnc}(g, K[g], u, sk[u], n, m_b, ad)$ $N_0 \leftarrow N_0 \cup \{(g, u, n, m_0, ad)\}$ $N_1 \leftarrow N_1 \cup \{(g, u, n, m_1, ad)\}$ $C \leftarrow C \cup \{((g, u, n, m_b, ad), c)\}$ $Q \leftarrow Q \cup \{((g, u, n, m_0, m_1, ad), c)\}$ Return c</p> <p><u>VERDEC(g, u, n, c, ad)</u></p> <p>require $K[g] \neq \perp$ and $\neg \text{group_is_corrupt}[g]$ require $vk[u] \neq \perp$ and $u \in \text{members}[g]$ $m \leftarrow \text{SS.VerDec}(g, K[g], u, vk[u], n, c, ad)$ If $m = \perp$ then return \perp $z \leftarrow ((g, u, n, m, ad), c)$ If $\text{pred}_{\text{trivial}}^{\text{sec}}(z, C)$ then return $\text{func}_{\text{out}}^{\text{sec}}(z, Q)$ If $b = 0$ then return \perp else return m</p>
<p><u>NEWHONUSER(u)</u></p> <p>require $sk[u] = vk[u] = \perp$ $(sk[u], vk[u]) \leftarrow \text{SS.UserKg}$ Return $vk[u]$</p> <p><u>EXPOSEUSER(u)</u></p> <p>require $sk[u] \neq \perp$ $\text{user_is_corrupt}[u] \leftarrow \text{true}$ Return $sk[u]$</p> <p><u>NEWCORRUSER(u, sk, vk)</u></p> <p>require $sk[u] = vk[u] = \perp$ $\text{user_is_corrupt}[u] \leftarrow \text{true}$ $sk[u] \leftarrow sk$; $vk[u] \leftarrow vk$</p>	<p><u>NEWHONGROUP(g, users)</u></p> <p>require $K[g] = \perp$ $K[g] \leftarrow \text{ss } \{0, 1\}^{\text{SS.gkl}}$ $\text{members}[g] \leftarrow \text{users}$</p> <p><u>EXPOSEGROUP($g$)</u></p> <p>require $K[g] \neq \perp$ and $\neg \text{chal}[g]$ $\text{group_is_corrupt}[g] \leftarrow \text{true}$ Return $K[g]$</p> <p><u>NEWCORRGROUP(g, K, users)</u></p> <p>require $K[g] = \perp$ $\text{group_is_corrupt}[g] \leftarrow \text{true}$ $K[g] \leftarrow K$; $\text{members}[g] \leftarrow \text{users}$</p>

Fig. 6. Left pane: Game defining in-group unforgeability IUF of a symmetric signcryption scheme SS with respect to a ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$. **Right pane:** Game defining out-group authenticated-encryption security OAE of SS with respect to a ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{sec}}$ and an output-guarding function $\text{func}_{\text{out}}^{\text{sec}}$. **Bottom pane:** User oracles $U = \{\text{NEWHONUSER}, \text{EXPOSEUSER}, \text{NEWCORRUSER}\}$ and group oracles $G = \{\text{NEWHONGROUP}, \text{EXPOSEGROUP}, \text{NEWCORRGROUP}\}$ that are provided to an adversary in either game, except that the boxed code only appears in the OAE security game.

$\frac{}{\text{pred}_{\text{trivial}}^{\text{suf}}(z, C)}$ Return $z \in C$	$\frac{}{\text{pred}_{\text{trivial}}^{\text{suf-except-group}}(z, C)}$ $((g, u, n, m, ad), c) \leftarrow z$ Return $\exists g' : ((g', u, n, m, ad), c) \in C$
$\frac{}{\text{pred}_{\text{trivial}}^{\text{euf}}(z, C)}$ $((g, u, n, m, ad), c) \leftarrow z$ Return $\exists c' : ((g, u, n, m, ad), c') \in C$	$\frac{}{\text{pred}_{\text{trivial}}^{\text{suf-except-user}}(z, C)}$ $((g, u, n, m, ad), c) \leftarrow z$ Return $\exists u' : ((g, u', n, m, ad), c) \in C$

Fig. 7. Sample ciphertext-triviality predicates which capture rules for deciding if a successfully decrypted VERDEC query was trivially obtainable or forgeable.

of all the prior encryption queries. If z is determined to be trivially obtainable from the information in C , then VERDEC exits early (with m as its output value); otherwise it sets the win flag. This check is performed by the *ciphertext-triviality predicate* $\text{pred}_{\text{trivial}}^{\text{auth}}$. We will describe the syntax and the sample variants of $\text{pred}_{\text{trivial}}^{\text{auth}}$ below.

Ciphertext-Triviality Predicates. The IUF security game is parameterized by ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$ (we will also parameterize the OAE game with $\text{pred}_{\text{trivial}}^{\text{auth}}$). Predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$ takes a tuple $z = ((g, u, n, m, ad), c)$ and a set C as input, where C contains tuples of the same format. Here z describes the input-output values of the current query to VERDEC oracle and each element of C contains an input-output transcript of a prior SIGENC oracle query. Predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$ returns true if z is considered to be trivially forgeable based on the information in C and false otherwise.

In Fig. 7 we define several ciphertext-triviality predicates. Predicate $\text{pred}_{\text{trivial}}^{\text{suf}}$ checks if $z \in C$, capturing the strongest possible level of authenticity. This requires that only prior outputs of SIGENC can be successfully queried to the VERDEC oracle; any other successful decryption query causes the adversary to win the IUF game. This predicate can be thought of as making the IUF game capture the “strong” unforgeability of ciphertexts in our group setting. One could capture existential unforgeability by considering the predicate $\text{pred}_{\text{trivial}}^{\text{euf}}$ that does not allow the adversary to win by merely producing new ciphertexts that decrypt to some tuple (g, u, n, m, ad) previously queried to SIGENC. Predicates $\text{pred}_{\text{trivial}}^{\text{suf-except-group}}$ and $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ capture the authenticity of schemes where a ciphertext encrypting (g, u, n, m, ad) is not bound to the group’s identifier or to the user’s identifier, respectively. We use $\text{pred}_{\text{trivial}}^{\text{suf}}$, $\text{pred}_{\text{trivial}}^{\text{suf-except-group}}$ and $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ in our security analysis of Keybase. In this work, we do not use $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ with the IUF game – we need it for OAE.

3.2 Out-Group Authenticated Encryption

The strongest version of the out-group AE security requires that an attacker outside a chat group can neither learn any information about the exchanged messages, nor modify the exchanged ciphertexts in any way. We also capture

weaker variants of this security notion. For example, the SealPacket encryption algorithm (as defined in Sect. 4) does not use a group’s symmetric key to explicitly bind its ciphertexts to a user’s signing key or a user’s identifier when used in isolation. So we capture a variant of out-group AE security that is as restrictive as possible except for allowing an attacker to violate the sender’s authenticity within any particular group.

OAE Game. Consider game \mathcal{G}^{OAE} of Fig. 6 for symmetric signcryption scheme SS , ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{auth}}$, output-guarding function $\text{func}_{\text{out}}^{\text{sec}}$, and adversary \mathcal{A}_{OAE} . The advantage in breaking the OAE security of SS is defined as $\text{Adv}_{\text{SS}, \text{pred}_{\text{trivial}}^{\text{auth}}, \text{func}_{\text{out}}^{\text{sec}}}^{\text{OAE}}(\mathcal{A}_{\text{OAE}}) = \Pr[\mathcal{G}_{\text{SS}, \text{pred}_{\text{trivial}}^{\text{auth}}, \text{func}_{\text{out}}^{\text{sec}}}^{\text{OAE}}(\mathcal{A}_{\text{OAE}})]$. Adversary \mathcal{A}_{OAE} is given access to user and group oracles U and G and to the encryption and decryption oracles SIGENC and VERDEC . The goal of the adversary is to guess the challenge bit b . Our security game is defined in the all-in-one style of [29, 32], where an adversary can learn the challenge bit by forging a ciphertext to its decryption oracle.

The user and group oracles in the OAE game are defined as in the IUF game, except it does not allow calling the EXPOSEGROUP oracle to expose the key of a group that was previously used for a left-or-right challenge-encryption query (as explained below). The OAE game never checks the contents of `user_is_corrupt`, deliberately giving the adversary full control over user keys.

The encryption oracle SIGENC takes (g, u, n, m_0, m_1, ad) and returns a ciphertext c by running $\text{SS.SigEnc}(g, \text{K}[g], u, \text{sk}[u], n, m_b, ad)$. The group and user keys $\text{K}[g]$ and $\text{sk}[u]$ are the only inputs to SS.SigEnc not directly chosen by the adversary querying the SIGENC oracle (and the encrypted message m_b depends on the challenge bit). The SIGENC query requires that $|m_0| = |m_1|$ and will only use insecure group keys for non-challenge encryptions (i.e. for $m_0 = m_1$). This SIGENC oracle captures nonce-misuse resistance [30], using the sets N_d to prevent trivial wins. At the end of SIGENC queries, the set C is updated to add the tuple $((g, u, n, m_b, ad), c)$, and the set Q is updated to add the tuple $((g, u, n, m_0, m_1, ad), c)$. Here the Q set can contain the input-output “transcript” of SIGENC queries from the adversary’s point of view, whereas the set C is more informative because it contains the message that was actually encrypted. We will explain the purpose of these sets below.

The decryption oracle VERDEC takes (g, u, n, c, ad) and returns the message m output by $\text{SS.VerDec}(g, \text{K}[g], u, \text{vk}[u], n, c, ad)$. Keys $\text{K}[g], \text{sk}[g]$ are the only inputs to SS.VerDec not directly chosen by the adversary querying the VERDEC oracle. The VERDEC oracle disallows queries with corrupt group keys; if an adversary knows a group’s key then it can decrypt ciphertexts for the group on its own. If SS.VerDec recovers a non- \perp message m and the end of the VERDEC oracle is reached, then the challenge bit is meant to be revealed through returning m if $b = 1$ and \perp otherwise. However, this intuition is not precise; it depends on how VERDEC responds to queries that are identified as being trivially forgeable. Similarly to how trivial forgeries were handled in the IUF game, here VERDEC builds $z = ((g, u, n, m, ad), c)$ and uses a ciphertext-triviality predicate $\text{pred}_{\text{trivial}}^{\text{sec}}$ to check z against the set C from SIGENC . If z is considered *not* trivially obtainable from the information in C , then VERDEC proceeds to its last instruction

that returns \perp or m depending on the challenge bit. Otherwise, VERDEC should return an output that does not depend on the challenge bit to prevent trivial wins. Such an output is produced by the *output-guarding function* $\text{func}_{\text{out}}^{\text{sec}}$, i.e. VERDEC returns the output of $\text{func}_{\text{out}}^{\text{sec}}(z, Q)$. We now describe the syntax and variants of $\text{func}_{\text{out}}^{\text{sec}}$.

Output-Guarding Functions. The OAE game can be parameterized by different choices of an output-guarding function $\text{func}_{\text{out}}^{\text{sec}}$. We define $\text{func}_{\text{out}}^{\text{sec}}$ to take a tuple $z = ((g, u, n, m, ad), c)$ and a set Q as input, where Q contains tuples with the format $((g, u, n, m_0, m_1, ad), c)$. Here z describes the input-output values of a single query to the VERDEC oracle, and each element of Q specifies the input-output of a prior SIGENC oracle query. At a high level, z contains the message m that was recovered during an ongoing VERDEC call, and m is the only value in z, Q not necessarily known by the adversary. One might want to define VERDEC to return m whenever the input is identified as a trivial forgery, but m could potentially trivially reveal the challenge bit. So one could roughly think of $\text{func}_{\text{out}}^{\text{sec}}$ as the function that should enable VERDEC to return m when possible. However, it should determine – from z and Q – if m would trivially help the adversary win and then “guard” VERDEC against returning this m .

$$\begin{array}{l|l} \text{func}_{\text{out}}^{\perp}(z, Q) & \text{func}_{\text{out}}^{\text{silence-with-}m_1}[\text{pred}_{\text{trivial}}](z, Q) \\ \text{Return } \perp & \text{For each } ((g, u, n, m_0, m_1, ad), c) \in Q \text{ do} \\ & \text{If } m_0 \neq m_1 \text{ then} \\ & \quad \text{If } \text{pred}_{\text{trivial}}(z, \{(g, u, n, m_0, ad), c\}) \text{ then return } m_1 \\ & \quad \text{If } \text{pred}_{\text{trivial}}(z, \{(g, u, n, m_1, ad), c\}) \text{ then return } m_1 \\ & \quad ((g, u, n, m, ad), c) \leftarrow z; \text{ Return } m \end{array}$$

Fig. 8. Sample output-guarding functions $\text{func}_{\text{out}}^{\perp}$ and $\text{func}_{\text{out}}^{\text{silence-with-}m_1}[\text{pred}_{\text{trivial}}]$. Function $\text{func}_{\text{out}}^{\text{silence-with-}m_1}$ is parameterized by a ciphertext-triviality predicate $\text{pred}_{\text{trivial}}$.

In Fig. 8 we define two output-guarding functions. The function $\text{func}_{\text{out}}^{\perp}$ always returns \perp . This provides no useful information to the adversary and so captures a comparatively weaker security notion. The function $\text{func}_{\text{out}}^{\text{silence-with-}m_1}[\text{pred}_{\text{trivial}}]$ is parameterized by an arbitrary ciphertext-triviality predicate $\text{pred}_{\text{trivial}}$ and captures the following logic. For every element in Q that describes a challenge encryption (i.e. $m_0 \neq m_1$) performed by SIGENC, this function checks whether z is trivially forgeable based on the information that the adversary could have learned from the corresponding response. This is checked if z would be trivially forgeable for *both* choices of $b \in \{0, 1\}$ or only for *only one* choice of b . The output-guarding function returns m_1 when this condition passes. If no element of Q triggered the above, then the output-guarding function returns the m contained in z , i.e. the actual message recovered in VERDEC.

The Use of $\text{func}_{\text{out}}^{\text{silence-with-}m_1}[\text{pred}_{\text{trivial}}]$ in Our Work. We target $\text{func}_{\text{out}}^{\text{silence-with-}m_1}[\text{pred}_{\text{trivial}}]$ as the output-guarding function that provides the strongest possible security guarantees for the schemes that we analyze in this work. For every $\text{pred}_{\text{trivial}}$ we use,

```

BoxMessage.SigEnc( $g, K_g, u, sk_u, n, m_{\text{body}}, ad$ ) // where  $n = (n_{\text{body}}, n_{\text{header}})$ 
 $c_{\text{body}} \leftarrow \text{XSalsa20-Poly1305.Enc}(K_g, n_{\text{body}}, m_{\text{body}})$ 
 $h_{\text{body}} \leftarrow \text{SHA-256}(n_{\text{body}} \parallel c_{\text{body}})$ 
 $m_{\text{header}} \leftarrow \langle ad, u, g, h_{\text{body}} \rangle$ 
 $c_{\text{header}} \leftarrow \text{SealPacket.SigEnc}(g, K_g, u, sk_u, n_{\text{header}}, m_{\text{header}}, \varepsilon)$ 
Return ( $c_{\text{body}}, c_{\text{header}}$ )

SealPacket.SigEnc( $g, K_g, u, sk_u, n, m, ad$ ) // where  $ad = \varepsilon$ 
 $h \leftarrow \text{SHA-512}(m)$ ;  $m_s \leftarrow \text{"Keybase-Chat-2"} \parallel \langle K_g, n, h \rangle$ 
 $s \leftarrow \text{Ed25519.Sig}(sk_u, m_s)$ ;  $m_e \leftarrow s \parallel m$ 
 $c \leftarrow \text{XSalsa20-Poly1305.Enc}(K_g, n, m_e)$ 
Return  $c$ 

```

Fig. 9. The `BoxMessage` and `SealPacket` algorithms used in Keybase for encrypting chat messages from a user to a group. Here g is the group’s identifier, K_g is the symmetric key shared by all group members, u is the sender’s identifier, and sk_u is the sender’s signing key.

$\text{pred}_{\text{trivial}}(z, \{(g, u, n, m^*, ad), c\})$ can only be true when z contains m^* . So for elements of Q with $m_0 \neq m_1$ only one of the two if conditions can pass, meaning it is necessary to silence the output. Otherwise the adversary can trivially win the game by building z, Q and evaluating $\text{pred}_{\text{trivial}}$ to distinguish between $b = 0$ or $b = 1$. (This attack assumes the adversary can always compute $\text{pred}_{\text{trivial}}(z, C)$ for SS, in spite of not knowing the challenge bit b that is needed to explicitly build C . This is true in all of our proofs.)

3.3 Symmetric Signcryption from Encryption and Signatures

In the full version, we introduce a provably secure version of Sign-then-Encrypt (StE). Its signcryption algorithm signs $s \leftarrow_s \text{DS.Sig}(sk_u, \langle g, n, m, ad \rangle)$ then outputs ciphertext $c \leftarrow \text{NE.Enc}(K_g, n, s \parallel m, \langle u, ad \rangle)$. We prove bounds of the form $\text{Adv}_{\text{StE}, \text{pred}_{\text{trivial}}^{\text{SUF}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF}}) \leq \text{Adv}_{\text{DS}}^{\text{SUF}}(\mathcal{A}_{\text{SUF}})$ and $\text{Adv}_{\text{StE}, \text{pred}_{\text{trivial}}^{\text{SUF}}, \text{func}_{\text{out}}^{\perp}}^{\text{OAE}}(\mathcal{A}_{\text{OAE}}) \leq \text{Adv}_{\text{NE}}^{\text{AEAD}}(\mathcal{A}_{\text{AEAD}})$.

4 Keybase Chat Encryption as Symmetric Signcryption

We analyze the security of the cryptographic algorithm `BoxMessage` that Keybase uses to encrypt and authenticate chat messages from a sender to a group. `BoxMessage` combines multiple cryptographic primitives to offer end-to-end encrypted messaging. In particular it uses XSalsa20-Poly1305, SHA-256, SHA-512, and Ed25519 as building blocks. Within `BoxMessage`, the `SealPacket` subroutine encrypts and authenticates message headers. We show the pseudocode for these algorithms in Fig. 9. We omit the decryption algorithms

`BoxMessage.VerDec` and `SealPacket.VerDec` from Fig. 9 as Keybase’s implementation of these algorithms follows naturally from the corresponding `SigEnc` algorithms. We define the `VerDec` algorithms explicitly in our formalization of `BoxMessage` and `SealPacket`.

To formalize the security of `BoxMessage`, it is crucial to first identify the formal primitive underlying this algorithm and the security goals it aims to achieve. None of the existing primitives in literature seem to aptly model this object, but it is naturally captured by the symmetric signcryption primitive that we defined in Sect. 3. Similarly, `SealPacket` can also be modeled as a symmetric signcryption scheme from which `BoxMessage` is built. In this section, we present modular constructions that cast `BoxMessage` and `SealPacket` as symmetric signcryption schemes. We first provide a general overview of the two algorithms.

The `BoxMessage` Chat-Encryption Algorithm. The `BoxMessage.SigEnc` algorithm accepts the following inputs – group’s identifier g , symmetric group key K_g , sender identifier u , sender signing key sk_u , nonce $n = (n_{\text{body}}, n_{\text{header}})$, message m_{body} , and associated data ad . It performs the following steps. First it calls `XSalsa20-Poly1305.Enc` to encrypt m_{body} using key K_g and nonce n_{body} , and obtains the ciphertext c_{body} . It builds header plaintext m_{header} as $\langle ad, u, g, h_{\text{body}} \rangle$ (a unique encoding of ad , u , g , and hash $h_{\text{body}} = \text{SHA-256}(n_{\text{body}} \parallel c_{\text{body}})$). It then invokes `SealPacket.SigEnc` to encrypt m_{header} using sk_u , K_g , and n_{header} , and obtains the ciphertext c_{header} . Finally, it returns $(c_{\text{body}}, c_{\text{header}})$. To decrypt ciphertext $(c_{\text{body}}, c_{\text{header}})$, the algorithm `BoxMessage.VerDec` (not shown) ensures that c_{header} decrypts into the header plaintext m_{header} that is equal to the unique string $\langle ad, u, g, h_{\text{body}} \rangle$ composed from the inputs of `BoxMessage.VerDec`. In Keybase, the sender identifier u is their username and the group identifier g is constructed canonically from the usernames of the group members.

The `SealPacket` Header-Encryption Algorithm. The `SealPacket` algorithm accepts the same inputs as `BoxMessage`, except it does not take associated data ad as input. We capture this by setting `SealPacket.AD` = $\{\varepsilon\}$, meaning $ad = \varepsilon$ is always true. When `SealPacket.SigEnc` is called from `BoxMessage.SigEnc`, it encrypts chat headers. To encrypt m with nonce n , it starts by hashing m to obtain $h = \text{SHA-512}(m)$. Then it builds an input m_s to the Ed25519 signature scheme by concatenating the prefix string “Keybase-Chat-2” with the unique encoding $\langle K_g, n, h \rangle$ of K_g , n , and h . It invokes `Ed25519.Sig` to produce a signature s over m_s using the signing key sk_u . Finally it calls `XSalsa20-Poly1305.Enc` to encrypt $m_e = s \parallel m$ using the key K_g and nonce n , and obtains the ciphertext c which is returned. To decryption ciphertext c , the `SealPacket.VerDec` algorithm (not shown) first recovers m_e from c and then parses m_e to obtain $s \parallel m$. Note that m_e can be unambiguously parsed into $s \parallel m$ because Ed25519 produces fixed-length signatures. Then `SealPacket.VerDec` reconstructs m_s and ensures that s verifies as a valid signature for m_s under the sender’s public key vk_u . We study the security of `SealPacket` in the context of the `BoxMessage` algorithm, but this is not the only context in which Keybase uses `SealPacket`. It is also used independently for the encryption of long strings and attachments. In the full version we detail other uses of `SealPacket` in Keybase.

Analysis Challenges. The descriptions of `BoxMessage` and `SealPacket` that we have given so far already present the following challenges in their analysis.

Key Reuse in BoxMessage. The same symmetric key K_g is used in `BoxMessage` and `SealPacket`. This violates the principle of key separation, which says that one should always use distinct keys for distinct algorithms and modes of operation. Without context separation, this potentially allows an attacker to forward ciphertexts produced by one algorithm to another. There is no explicit context separation, so our analysis will “extract” separation by making assumptions of Ed25519 and using low-level details of how messages are encoded.

Cyclic Key Dependency in SealPacket. The message m_s signed in `SealPacket` is derived from the symmetric group key K_g which is also used to encrypt the signature. This produces what is known as an “encryption cycle”, a generalization of encrypting one’s own key [19]. Standard AEAD security does not guarantee security when messages being encrypted depend on the key used for encryption. We use an extension of AEAD security allowing key-dependent messages and prove (in the random oracle model) that XSalsa20-Poly1305 achieves it for the particular key-dependent messages required.

Lack of Group/User Binding in SealPacket. By looking at the `SealPacket` algorithm in Fig. 9 we can see that the inputs u and g are never used by the algorithm. This means that a `SealPacket` ciphertext does not, in general, bind to the group’s or user’s identifiers. This could potentially allow a malicious user to impersonate another group member. When `SealPacket` is used within `BoxMessage`, it is always invoked on a message that contains the group’s and the user’s identifier, so the lack of group/user binding in `SealPacket` is not consequential there.

Nonce Repetition in Keybase. XSalsa20-Poly1305 is not secure when nonces repeat so our security analysis disallows nonce repetition between `BoxMessage` and/or `SealPacket`. The Keybase implementation uses uniformly random nonces, making collisions highly unlikely. Moreover, our results show that `BoxMessage` is robust to accidental non-uniformity in randomness as long nonces do not repeat. The XSalsa20-Poly1305 authenticated encryption scheme combines the XSalsa20 stream cipher and the Poly1305 *one-time* message authentication code. The stream is derived from the key and nonce and is used for keying Poly1305, so if nonces repeat then privacy and integrity may both be broken.

Message Encryption Scheme BM. Our modular symmetric signcryption construction BM models the `BoxMessage` chat-encryption algorithm as follows.

Construction 1. Let $\mathcal{M} \subseteq \{0, 1\}^*$. Let NE be a nonce-based encryption scheme. Let H be a hash function. Let SP be a deterministic symmetric signcryption scheme. Then $\text{BM} = \text{BOX-MESSAGE-SS}[\mathcal{M}, \text{NE}, \text{H}, \text{SP}]$ is the deterministic symmetric signcryption scheme as defined in Fig. 10, with message space $\text{BM.MS} = \mathcal{M}$ and associated-data space $\text{BM.AD} = \{0, 1\}^*$. We require the following. The group key taken by BM is used as the key for both NE and SP, so $\text{BM.gkl} = \text{NE.kl} = \text{SP.gkl}$. The nonce taken by BM is a pair containing a separate nonce for each of NE and SP, so $\text{BM.NS} = \{0, 1\}^{\text{NE.nl}} \times \text{SP.NS}$.

```

BM.UserKg
( $sk, vk$ )  $\leftarrow$   $\$$  SP.UserKg ; Return ( $sk, vk$ )

BM.SigEnc( $g, K_g, u, sk_u, n, m_{\text{body}}, ad$ ) //  $K_g \in \{0, 1\}^{256}$ 
( $n_{\text{body}}, n_{\text{header}}$ )  $\leftarrow$   $n$  //  $n_{\text{body}}, n_{\text{header}} \in \{0, 1\}^{192}$ 
// Encrypt the message body
 $c_{\text{body}} \leftarrow$  NE.Enc( $K_g, n_{\text{body}}, m_{\text{body}}$ ) // NE = XSalsa20-Poly1305
// Create and encrypt the message header
 $h_{\text{body}} \leftarrow$  H( $n_{\text{body}} \parallel c_{\text{body}}$ ) ;  $m_{\text{header}} \leftarrow \langle ad, u, g, h_{\text{body}} \rangle$  // H = SHA-256
 $c_{\text{header}} \leftarrow$  SP.SigEnc( $g, K_g, u, sk_u, n_{\text{header}}, m_{\text{header}}, \varepsilon$ ) // SP = SEAL-PACKET-SS
Return ( $c_{\text{body}}, c_{\text{header}}$ )

BM.VerDec( $g, K_g, u, vk_u, n, c, ad$ )
( $n_{\text{body}}, n_{\text{header}}$ )  $\leftarrow$   $n$  ; ( $c_{\text{body}}, c_{\text{header}}$ )  $\leftarrow$   $c$ 
// Recover and verify the message header
 $m_{\text{header}} \leftarrow$  SP.VerDec( $g, K_g, u, vk_u, n_{\text{header}}, c_{\text{header}}, \varepsilon$ )
 $h_{\text{body}} \leftarrow$  H( $n_{\text{body}} \parallel c_{\text{body}}$ )
If  $m_{\text{header}} \neq \langle ad, u, g, h_{\text{body}} \rangle$  then return  $\perp$ 
// Recover and return the message body
 $m_{\text{body}} \leftarrow$  NE.Dec( $K_g, n_{\text{body}}, c_{\text{body}}$ ) ; Return  $m_{\text{body}}$ 

```

Fig. 10. Symmetric signcryption scheme BM = BOX-MESSAGE-SS[\mathcal{M} , NE, H, SP]. The right-aligned comments provide a guideline for modeling Keybase.

```

SP.UserKg
( $sk, vk$ )  $\leftarrow$   $\$$  DS.Kg ; Return ( $sk, vk$ )

SP.SigEnc( $g, K_g, u, sk_u, n, m, ad$ ) //  $K_g \in \{0, 1\}^{256}, n \in \{0, 1\}^{192}, ad = \varepsilon$ 
 $h \leftarrow$  H( $m$ ) // H = SHA-512
 $m_s \leftarrow$  “Keybase-Chat-2”  $\parallel \langle K_g, n, h \rangle$ 
 $s \leftarrow$  DS.Sig( $sk_u, m_s$ ) ;  $m_e \leftarrow s \parallel m$  // DS = Ed25519
 $c \leftarrow$  NE.Enc( $K_g, n, m_e$ ) ; Return  $c$  // NE = XSalsa20-Poly1305

SP.VerDec( $g, K_g, u, vk_u, n, c, ad$ ) //  $ad = \varepsilon$ 
 $m_e \leftarrow$  NE.Dec( $K_g, n, c$ )
If  $m_e = \perp$  then return  $\perp$ 
 $s \parallel m \leftarrow$   $m_e$  // s.t.  $|s| = \text{DS.sl}, |m| \geq 0$ 
 $h \leftarrow$  H( $m$ )
 $m_s \leftarrow$  “Keybase-Chat-2”  $\parallel \langle K_g, n, h \rangle$ 
If  $\neg \text{DS.Ver}(vk_u, m_s, s)$  then return  $\perp$  else return  $m$ 

```

Fig. 11. Symmetric signcryption scheme SP = SEAL-PACKET-SS[H, DS, NE]. The right-aligned comments provide a guideline for modeling Keybase.

Header Encryption Scheme SP. Our modular symmetric signcryption construction SP models the header-encryption algorithm SealPacket as follows.

Construction 2. Let H be a hash function. Let DS be a deterministic digital signature scheme. Let NE be a nonce-based encryption scheme. Then $SP = \text{SEAL-PACKET-SS}[H, DS, NE]$ is the symmetric signcryption scheme as defined in Fig. 11, with group-key length $SP.\text{gkl} = NE.\text{kl}$, nonce space $SP.\text{NS} = \{0, 1\}^{NE.\text{nl}}$, message space $SP.\text{MS} = \{0, 1\}^*$, and associated-data space $SP.\text{AD} = \{\varepsilon\}$.

5 Security Analysis of Keybase Chat Encryption

In this section we analyze the security of the symmetric signcryption schemes BM and SP defined in Sect. 4. In Sect. 5.1, we show the in-group unforgeability of BM and SP . In Sects. 5.2 and 5.3, we show the out-group AE security of BM and SP . This requires us to introduce two weaker variants of the OAE security notion, one each for BM and SP , by relaxing the level of nonce-misuse requirements of the OAE game defined in Fig. 6. The SP analysis requires two new security notions, \mathcal{M} -sparsity for digital signature schemes and *authenticated encryption for key-dependent messages* for nonce-based encryption schemes.

5.1 In-Group Unforgeability of BoxMessage and SealPacket

In-Group Unforgeability of BoxMessage. In-group unforgeability of $BM = \text{BOX-MESSAGE-SS}[\mathcal{M}, NE, H, SP]$ reduces to the security of SP and H . A BM ciphertext is a pair $(c_{\text{body}}, c_{\text{header}})$ comprising an NE ciphertext c_{body} and an SP ciphertext c_{header} , which encrypts $\langle ad, u, g, h_{\text{body}} \rangle$. The adversary's objective is to forge a BM ciphertext by either forging c_{body} or c_{header} . The adversary can use a corrupt group key K_g , so c_{body} ciphertexts are easily forged. However, this does not suffice to produce a BM forgery because c_{header} encrypts the hash of c_{body} . Therefore, it would need to forge a corresponding c_{header} ciphertext. The IUF security of SP prevents the adversary from forging c_{header} ciphertexts. As a result, the adversary can only reuse honestly generated c_{header} from its prior queries to SIGENC in its forgery attempts. Since an honest c_{header} effectively commits to ad , u , g , h_{body} , and n_{header} , using an old c_{header} to construct a new BM ciphertext requires finding a new NE nonce-ciphertext pair that hashes to the same h_{body} under H . Collision resistance of H prevents this. The formal proof of Theorem 1 is in the full version.

Theorem 1. Let $BM = \text{BOX-MESSAGE-SS}[\mathcal{M}, NE, H, SP]$ be the symmetric signcryption scheme built from some \mathcal{M}, NE, H, SP as specified in Construction 1. Let $\text{pred}_{\text{trivial}}^{\text{suf}}$ and $\text{pred}_{\text{trivial}}^{\text{suf-except-group}}$ be the ciphertext-triviality predicates as defined in Fig. 7. Let $\mathcal{A}_{\text{IUF-of-BM}}$ be any adversary against the IUF security of BM with respect to $\text{pred}_{\text{trivial}}^{\text{suf}}$. Then we can build adversaries $\mathcal{A}_{\text{IUF-of-SP}}$ and \mathcal{A}_{CR} such that

$$\text{Adv}_{\text{BM}, \text{pred}_{\text{trivial}}^{\text{suf}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF-of-BM}}) \leq \text{Adv}_{\text{SP}, \text{pred}_{\text{trivial}}^{\text{suf-except-group}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF-of-SP}}) + \text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}_{\text{CR}}).$$

In-Group Unforgeability of SealPacket. In-group unforgeability of the symmetric signcryption scheme $SP = \text{SEAL-PACKET-SS}[H, DS, NE]$ reduces to the




Fig. 12. A summary of the reductions that we provide for the wOAE security of SP and the bwOAE security of BM.

security of DS and H. We parameterize the IUF security of SP to aim for a relaxed version of strong unforgeability because SP ciphertexts do not directly depend on the group’s identifier g (even though it depends on the group key K_g).

An SP ciphertext encrypts $s \parallel m$ under K_g . The adversary can use a corrupt K_g , but forging an SP ciphertext still requires the signature s . So the adversary must either forge a new signature or reuse an honest signature from a prior SIGENC query. The SUFCMA security of DS prevents the former. An honest signature s is computed over “Keybase-Chat-2” $\parallel \langle K_g, n, h \rangle$ where h is the hash of the message m . Hence reusing an honest signature could use a new SP ciphertext that encrypts $s \parallel m$ with K_g, n , but the tidiness of NE prevents this. So reusing an honest signature requires finding a new message that hashes to the same h under H. Collision resistance of H prevents this. The formal proof of Theorem 2 is in the full version.

Theorem 2. *Let $SP = \text{SEAL-PACKET-SS}[H, DS, NE]$ be the symmetric sign-encryption scheme built from some H, DS, and NE as specified in Construction 2. Let $\text{pred}_{\text{trivial}}^{\text{suf-except-group}}$ be the ciphertext-triviality predicate as defined in Fig. 7. Let $\mathcal{A}_{\text{IUF-of-SP}}$ be any adversary against the IUF security of SP with respect to $\text{pred}_{\text{trivial}}^{\text{suf-except-group}}$. Then we can build adversaries $\mathcal{A}_{\text{SUFCMA}}$ and \mathcal{A}_{CR} such that*

$$\text{Adv}_{\text{SP, pred}_{\text{trivial}}^{\text{suf-except-group}}}^{\text{IUF}}(\mathcal{A}_{\text{IUF-of-SP}}) \leq \text{Adv}_{\text{DS}}^{\text{SUFCMA}}(\mathcal{A}_{\text{SUFCMA}}) + \text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}_{\text{CR}}).$$

5.2 Out-Group AE Security of BoxMessage

Out-group AE security of $BM = \text{BOX-MESSAGE-SS}[\mathcal{M}, NE, H, SP]$ reduces to the security of its underlying primitives as summarized by the rightmost arrows of Fig. 12. At a high level, we show that BM achieves a variant of OAE security (bwOAE) if SP achieves another variant of OAE security (wOAE) and H is collision-resistant. Because $NE = \text{XSalsa20-Poly1305}$ in Keybase (which is not nonce-misuse resistant), both variants disallow nonce repetition.

Intuition. An BM ciphertext is a pair $(c_{\text{body}}, c_{\text{header}})$ consisting of an NE ciphertext c_{body} and an SP ciphertext c_{header} . One way the adversary could learn the challenge bit is by querying its VERDEC oracle on a forged BM ciphertext that decrypts successfully. In order to accomplish that, the adversary must either forge the underlying SP ciphertext c_{header} or reuse an honestly generated c_{header} .

The former is prevented by the out-group AE security of SP. The latter is prevented by the collision resistance of H because of the following. An honestly generated c_{header} effectively commits to ad , u , g , h_{body} , and n_{header} . In order to reuse c_{header} , an adversary must find a new NE nonce-ciphertext pair that hashes to h_{body} , hence producing a collision. It follows that the VERDEC oracle is essentially useless to the adversary; it can only serve to decrypt non-challenge ciphertexts that were previously returned by SIGENC. So it remains to show that the adversary cannot learn the challenge bit solely based on the BM ciphertexts that it receives from SIGENC. For any ciphertext $(c_{\text{body}}, c_{\text{header}})$ returned by SIGENC, the SP ciphertext c_{header} encrypts a hash of c_{body} but otherwise does not depend on the challenge bit. So the adversary gains no advantage from observing c_{header} . Finally, the AEAD security of NE guarantees that c_{body} does not reveal the challenge bit.

Because the header encryption scheme SP and the body encryption scheme NE use the same symmetric key K_g , we require integrity of SP ciphertexts produced using K_g hold even when the adversary can obtain other NE encryptions under the same key. Similarly, the NE ciphertexts generated using the symmetric key K_g should be indistinguishable even when the adversary can obtain SP encryptions and decryptions under the same key. We introduce a variant of the OAE game in Definition 3 to capture these joint requirements.

Restrictions on Nonce Misuse in BM and SP. We now define new variants of out-group AE security for our analysis of Keybase. The BM and SP schemes in Keybase are not nonce-misuse resistant so we modify the OAE game to disallow nonce repetition. We start with wOAE security for SP.

Definition 1. *Let SS be a symmetric signcryption scheme. Consider the OAE security game for SS of Fig. 6 (w.r.t. any $\text{pred}_{\text{trivial}}^{\text{sec}}$, $\text{func}_{\text{out}}^{\text{sec}}$). We define a new variant of this game as follows. The instruction preventing nonce misuse*

require $\forall d \in \{0, 1\}, (g, u, n, m, ad) \notin N_d$ is replaced with *require $(g, n) \notin N$.*

In addition, the instructions updating the nonce set

$N_0 \leftarrow N_0 \cup \{(g, u, n, m_0, ad)\}$ *are replaced with* $N \leftarrow N \cup \{(g, n)\}.$
 $N_1 \leftarrow N_1 \cup \{(g, u, n, m_1, ad)\}$

We denote the resulting game (and security notion) by wOAE. It is a weak variant of OAE that does not require nonce-misuse resistance. We define an adversary's advantage in breaking the wOAE security of SS in the natural way.

Now we define bwOAE security for BM. The nonce of BM is a pair of two separate nonces $n = (n_{\text{body}}, n_{\text{header}})$. The bwOAE security game independently applies the group-nonce uniqueness condition introduced in Definition 1 to each of (g, n_{body}) and (g, n_{header}) , and it also requires that $n_{\text{body}} \neq n_{\text{header}}$. This is a necessary because BM calls NE.Enc on (g, n_{body}) , and SP calls NE.Enc on (g, n_{header}) . In Keybase both NE schemes are XSalsa20-Poly1305 using the same key.

Definition 2. *Let \mathcal{X}, \mathcal{Y} be any sets. Let SS be a symmetric signcryption scheme with the nonce space $\text{SS.NS} = \mathcal{X} \times \mathcal{Y}$. Consider the OAE security game for SS of Fig. 6 (w.r.t. any $\text{pred}_{\text{trivial}}^{\text{sec}}$, $\text{func}_{\text{out}}^{\text{sec}}$). We define a new variant of this game as follows. The instruction preventing nonce misuse*

require $\forall d, (g, u, n, m_d, ad) \notin N_d$ is replaced with

$$\begin{aligned} & (n_{\text{body}}, n_{\text{header}}) \leftarrow n \\ & \text{If } n_{\text{body}} = n_{\text{header}} \text{ then return } \perp \\ & \text{If } (g, n_{\text{body}}) \in N \text{ then return } \perp \\ & \text{If } (g, n_{\text{header}}) \in N \text{ then return } \perp. \end{aligned}$$

In addition, the instructions updating the nonce set

$$\begin{aligned} N_0 & \leftarrow N_0 \cup \{(g, u, n, m_0, ad)\} & N & \leftarrow N \cup \{(g, n_{\text{header}})\} \\ N_1 & \leftarrow N_1 \cup \{(g, u, n, m_1, ad)\} & N & \leftarrow N \cup \{(g, n_{\text{body}})\}. \end{aligned}$$

We denote the resulting game (and security notion) by **bwOAE**. Beyond being defined for **SS** with a bipartite nonce space, this variant of **OAE** is weak in that it does not require nonce-misuse resistance. We define an adversary's advantage in breaking the **bwOAE** security of **SS** in the natural way.

The Joint Security Required of SP and NE. Here we define the security notion required from **SP** when it is used in the presence of arbitrary **NE** encryptions under the same symmetric group keys that are used by **SP**. We call this notion **wOAE**[**ENC**[\mathcal{M} , **NE**]]. It is a parameterized version of the **wOAE** game defined in Definition 1. We use it for our analysis of the **bwOAE** security of **BM**.

At the start of this section we discussed that the security reduction for **BM** intuitively requires that it is hard to forge an **SP** ciphertext (without knowing the corresponding group key K_g) in the presence of **NE** encryptions. Our definition of **wOAE**[**ENC**[\mathcal{M} , **NE**]] captures this by providing the adversary access to an **NE** encryption oracle **ENC** in addition to the **SIGENC** and **VERDEC** oracles in the out-group **AE** security game of **SP**. We stress that proving the security of **BM** does not, in principle, require us to provide the **SIGENC** oracle to the adversary. We choose to require this stronger level of security from **SP** because of the following reasons. On the one hand, in Sect. 4 we explained why it is beneficial to prove that **SP** satisfies a strong security notion, going beyond what is required by **BM**. On the other hand, this stronger security notion that we require from **SP** will not come at the cost of introducing additional assumptions or achieving looser concrete-security bounds in our analysis of **BM**.

Definition 3. Let **SS** be a symmetric signcryption scheme. Let $\mathcal{M} \subseteq \{0, 1\}^*$. Let **NE** be a nonce-based encryption scheme. Consider the **wOAE** security game for **SS** as defined in Definition 1 (w.r.t. any $\text{pred}_{\text{trivial}}^{\text{sec}}, \text{func}_{\text{out}}^{\text{sec}}$). We define a variant of this game by adding an oracle that is defined as follows.

$$\begin{aligned} & \text{ENC}[\mathcal{M}, \text{NE}](g, n_{\text{body}}, m_{\text{body},0}, m_{\text{body},1}) \\ & \text{require } K[g] \neq \perp \text{ and } |m_{\text{body},0}| = |m_{\text{body},1}| \\ & \text{require } (g, n_{\text{body}}) \notin N \text{ and } m_{\text{body},0}, m_{\text{body},1} \in \mathcal{M} \\ & \text{If } m_{\text{body},0} \neq m_{\text{body},1} \text{ then} \\ & \quad \text{If } \text{group_is_corrupt}[g] \text{ then return } \perp \\ & \quad \text{chal}[g] \leftarrow \text{true} \\ & c_{\text{body}} \leftarrow \text{NE.Enc}(K[g], n_{\text{body}}, m_{\text{body},b}) \\ & N \leftarrow N \cup \{(g, n_{\text{body}})\}; \text{ Return } c_{\text{body}} \end{aligned}$$

It shares set N , bit b , and the tables **K**, **group_is_corrupt**, **chal** with the rest of the security game. We denote the resulting game (and security notion) by **wOAE**[**ENC**[\mathcal{M} , **NE**]]. It simultaneously requires out-group **AE** security of **SS**

(without nonce repetition) and an IND-style security of NE. We define an adversary's advantage in breaking this security notion in the natural way.

Note that we require the messages that the adversary queries to the ENC oracle to be in \mathcal{M} . Intuitively, in our security analysis of BM, an adversary will only be able to obtain NE encryptions of messages in the message space of BM. So in the security reduction for BM we will use $\mathcal{M} = \text{BM.MS}$.

Out-Group AE Security of BoxMessage. We prove bwOAE security of BM. The formal proof of Theorem 3 is in the full version.

Theorem 3. *Let $\text{BM} = \text{BOX-MESSAGE-SS}[\mathcal{M}, \text{NE}, \text{H}, \text{SP}]$ be the symmetric signcryption scheme built from some $\mathcal{M}, \text{NE}, \text{H}, \text{SP}$ as specified in Construction 1. Let $\text{pred}_{\text{trivial}}^{\text{suf}}$ and $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ be the ciphertext-triviality predicates as defined in Fig. 7. Let $\text{func}_{\text{out}}^{\perp}$ be the output-guarding functions as defined in Fig. 8. Let $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ be the security notion as defined in Definition 3. Let $\mathcal{A}_{\text{bwOAE-of-BM}}$ be any adversary against the bwOAE security of BM with respect to $\text{pred}_{\text{trivial}}^{\text{suf}}$ and $\text{func}_{\text{out}}^{\text{sec}}$. Then we build adversaries $\mathcal{A}_{\text{wOAE-of-SP}}$ and \mathcal{A}_{CR} such that*

$$\begin{aligned} \text{Adv}_{\text{BM}, \text{pred}_{\text{trivial}}^{\text{suf}}, \text{func}_{\text{out}}^{\perp}}^{\text{bwOAE}}(\mathcal{A}_{\text{bwOAE-of-BM}}) &\leq \text{Adv}_{\text{SP}, \text{pred}_{\text{trivial}}^{\text{suf-except-user}}, \text{func}_{\text{out}}^{\perp}}^{\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]}(\mathcal{A}_{\text{wOAE-of-SP}}) \\ &\quad + \text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}_{\text{CR}}). \end{aligned}$$

Note that we prove the security of BM with respect to $\text{pred}_{\text{trivial}}^{\text{suf}}$ and $\text{func}_{\text{out}}^{\perp}$. As discussed in Sect. 3, $\text{pred}_{\text{trivial}}^{\text{suf}}$ essentially requires BM to have ciphertext integrity. Our result relies on the security of SP with respect to $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ and $\text{func}_{\text{out}}^{\perp}$. Recall that $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ basically requires SP to have ciphertext integrity, except it allows for an honest ciphertext to be successfully decrypted even with respect to a wrong user identifier; the latter is not considered a “valid” forgery. This does not translate to an attack against BM because it only uses SP to encrypt header messages $m_{\text{header}} = \langle ad, u, g, h_{\text{body}} \rangle$ that contain u , and the BM.VerDec algorithm verifies that the group identifier it received as input matches the one that was parsed from m_{header} .

5.3 Out-Group AE Security of SealPacket

Out-group AE security of $\text{SP} = \text{SEAL-PACKET-SS}[\text{H}, \text{DS}, \text{NE}]$ reduces to the security NE and DS (see Fig. 12). In particular, $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ security holds if NE provides authenticated encryption for key-dependent messages and DS produces \mathcal{M} -sparse signatures. We introduce these security notions below.

Intuition. Recall that in the $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ game, the adversary is provided with (un)signcryption oracles SIGENC and VERDEC for SP, and an encryption oracle ENC for NE. Each of these returns output based on a challenge bit that is shared between them. The adversary can use three approaches to learn the challenge bit. It can (a) attempt SP forgeries by calling its SP decryption

oracle VERDEC; (b) make left-or-right queries to its NE encryption oracle ENC; (c) make left-or-right queries to its SP encryption oracle SIGENC.

The adversary is allowed to expose users' signing keys so it could attempt to forge an SP ciphertext using an exposed DS signing key and its ENC oracle. The adversary would then query the resulting ciphertext to its VERDEC oracle in an attempt to trivially win the game. We show that the adversary is unable to accomplish this. The ENC oracle is defined to only produce encryptions of the messages from the set \mathcal{M} . In the implementation of Keybase, the messages from \mathcal{M} have a specific encoding; we will rely on this property in our proof. In contrast, any ciphertext successfully decrypted by VERDEC must encrypt a message of the form $m_e = s \parallel m$ where s is a valid DS signature. So the adversary needs to find a signature s that is consistent with the message encoding that is permitted by ENC. The \mathcal{M} -sparseness of DS signatures, which we formalize below, prevents this. It follows that the VERDEC oracle does not help the adversary to win the game by querying ciphertexts that were previously returned by ENC.¹

Now we can reimagine the ENC and SIGENC oracles as producing NE encryptions of key-dependent messages. The SIGENC oracle requires messages to be derived as a specific function of the symmetric group key K_g . The ENC oracle can be thought of as allowing to messages that are derived from "constant" functions, meaning the chosen messages do not depend on K_g . We can also view the VERDEC oracle as an NE decryption oracle that prevents the adversary from trivially winning the game by merely querying the ciphertexts it previously obtained from either ENC or SIGENC. We define the AE security of NE for key-dependent messages and show that the adversary can only win the $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ game against SP if it can win the KDMAE game against NE.

Reliance on the Message Encoding in Keybase. We mentioned in the intuition that we rely on the encoding of messages in \mathcal{M} in our proof. We emphasize that avoiding this dependency is non-trivial. The cyclic key dependency within SP and the key reuse between BM and SP pose significant challenges when considering the possibility of an alternate proof.

\mathcal{M} -sparse Signatures. Consider game $\mathcal{G}^{\text{SPARSE}}$ of Fig. 13, defined for a digital signature scheme DS, a set $\mathcal{M} \subseteq \{0, 1\}^*$, and an adversary $\mathcal{A}_{\text{SPARSE}}$. The advantage of $\mathcal{A}_{\text{SPARSE}}$ in breaking the \mathcal{M} -SPARSE security of DS is defined as $\text{Adv}_{\text{DS}, \mathcal{M}}^{\text{SPARSE}}(\mathcal{A}_{\text{SPARSE}}) = \text{Pr}[\mathcal{G}_{\text{DS}, \mathcal{M}}^{\text{SPARSE}}(\mathcal{A}_{\text{SPARSE}})]$. Intuitively, this game captures the inability of an adversary to produce a signature that conforms to the message space \mathcal{M} even though the adversary chooses the public key used to verify the signature. More formally, the adversary wins if it is able to return (vk, m, s, γ) such that s verifies as a signature over the message m under the verification key vk and $s \parallel \gamma \in \mathcal{M}$. We stress that the adversary is allowed to choose an arbitrary – possibly malformed – verification key. The adversary is not required to know the corresponding signing key, and such a key may in fact not exist.

¹ The $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ game itself also prevents the adversary from trivially winning by querying VERDEC on a ciphertext that was previously returned by SIGENC.

We verify our intuition about the \mathcal{M} -sparsity of the Ed25519 signature scheme underlying SP in the full version. Ed25519 is a deterministic signature scheme introduced by Bernstein, Duif, Lange, Schwabe, and Yang in [17]. It is obtained by applying the commitment-variant of the Fiat-Shamir transform to an identification scheme. Therefore a signature produced by Ed25519 consists of the commitment and response of the identification scheme. The adversary can only win the SPARSE game of Ed25519 if it is able to produce an accepting conversation transcript for the identification scheme such that the corresponding commitment conforms to \mathcal{M} . Commitments in the identification scheme underlying Ed25519 are elements of a prime-order group. We prove that finding such a commitment is only possible if the adversary is able to find a group element and its discrete logarithm such that the group element is in \mathcal{M} which we show is hard in the generic group model.

The Message Space \mathcal{M} . Keybase uses the MessagePack serialization format [21] to encode plaintext messages. Plaintext messages are represented using a custom data structure in Keybase. So the serialized MessagePack encoding of a plaintext is a byte sequence that not only stores the plaintext itself but also some metadata about the data structure that represents it. For messages encrypted by BM, the metadata about the data structure happens to be located in the first 17 bytes of the encoding. This means that the encoding of every plaintext encrypted by BM contains a fixed 17-byte prefix. Let this 17-byte prefix be pre . Then we define the message space of BM by $\text{BM.MS} = \{\text{pre} \parallel \nu \mid \nu \in \{0, 1\}^*\}$.

Message-Deriving Functions. Let ϕ be any function that takes a symmetric key K as input and uses it to derive and return some message m . We call ϕ a *message-deriving* function and will consider some classes (i.e. sets) Φ of message-deriving functions. We require that the length of an output returned by ϕ must not depend on its input; we denote the output length of ϕ by $\|\phi\|$.

AE Security of NE for Key-Dependent Messages. Consider game $\mathcal{G}^{\text{KDMAE}}$ of Fig. 14, defined for a nonce-based encryption scheme NE, a class of message-deriving functions Φ , and an adversary $\mathcal{A}_{\text{KDMAE}}$. The advantage of $\mathcal{A}_{\text{KDMAE}}$ in breaking the Φ -KDMAE security of NE is defined as $\text{Adv}_{\text{NE}, \Phi}^{\text{KDMAE}}(\mathcal{A}) = 2 \cdot \Pr[\mathcal{G}_{\text{NE}, \Phi}^{\text{KDMAE}}(\mathcal{A})] - 1$. This game can be thought of as a modification of the AEAD security game for NE (Fig. 4) which does not require nonce-misuse resistance. The core difference is that the ENC oracle now takes message-deriving functions $\phi_0, \phi_1 \in \Phi$ as input. The challenge message is derived as $m_b \leftarrow \phi_b(\text{K}[g])$ for $b \in \{0, 1\}$, where $\text{K}[g]$ is the symmetric group key associated to g . Trivial attacks are prevented by requiring that ϕ_0, ϕ_1 have the same output length and that $\phi_0 = \phi_1$ whenever ENC is called for a corrupt group. Our definition is based on prior work [6, 12, 13, 19]. There are strong impossibility results [12] regarding




Fig. 13. Game defining \mathcal{M} -sparsity of a digital signature scheme DS for a set \mathcal{M} .

<u>Game $\mathcal{G}_{\text{NE}, \Phi}^{\text{KDMAE}}(\mathcal{A}_{\text{KDMAE}})$</u> $b \leftarrow \$ \{0, 1\}$; $b' \leftarrow \$ \mathcal{A}_{\text{KDMAE}}^{\text{G, ENC, DEC}}$ Return $b = b'$	<u>$\text{ENC}(g, n, \phi_0, \phi_1)$</u> require $K[g] \neq \perp$ and $(g, n) \notin N$ require $\phi_0, \phi_1 \in \Phi$ and $\ \phi_0\ = \ \phi_1\ $ If $\phi_0 \neq \phi_1$ then If $\text{group_is_corrupt}[g]$ then return \perp $\text{chal}[g] \leftarrow \text{true}$ $m_b \leftarrow \phi_b(K[g])$; $c \leftarrow \text{NE.Enc}(K[g], n, m_b)$ $N \leftarrow N \cup \{(g, n)\}$; $C \leftarrow C \cup \{(g, n, c)\}$ Return c	
<u>$\text{DEC}(g, n, c)$</u> require $K[g] \neq \perp$ and $\neg \text{group_is_corrupt}[g]$ require $(g, n, c) \notin C$ $m \leftarrow \text{NE.Dec}(K[g], n, c)$ If $b = 0$ then return \perp else return m	<u>$\text{NEWHONGROUP}(g)$</u> require $K[g] = \perp$ $K[g] \leftarrow \$ \{0, 1\}^{\text{NE.kl}}$	<u>$\text{EXPOSEGROUP}(g)$</u> require $K[g] \neq \perp$ and $\neg \text{chal}[g]$ $\text{group_is_corrupt}[g] \leftarrow \text{true}$ Return $K[g]$
		<u>$\text{NEWCORRGROUP}(g, K)$</u> require $K[g] = \perp$ $\text{group_is_corrupt}[g] \leftarrow \text{true}$ $K[g] \leftarrow K$

Fig. 14. Game defining authenticated-encryption security of NE for Φ -key-dependent messages, where Φ is a class of message-deriving functions and $G = \{\text{NEWHONGROUP}, \text{EXPOSEGROUP}, \text{NEWCORRGROUP}\}$.

the existence of schemes that are secure with respect to very large classes of message-deriving functions Φ . We sidestep these results by considering a very narrow and simple class Φ_{SP} that we define below.

The Class of Message-Deriving Functions Φ_{SP} Earlier we discussed that in the $\text{wOAE}[\text{ENC}[\mathcal{M}, \text{NE}]]$ security game for SP, the SIGENC and ENC oracles can be thought of as returning an NE ciphertext that encrypts an output of some message-deriving function. We now define the class Φ_{SP} containing all message-deriving functions that are used by either SIGENC or ENC .

Construction 3. Let NE be a nonce-based encryption scheme. Let H be a hash function. Let DS be a digital signature scheme. Let SIGENC-DER and ENC-DER be the parameterized message-deriving functions that are defined as follows, each taking an NE key $K \in \{0, 1\}^{\text{NE.kl}}$ as input.

<u>$\text{SIGENC-DER}[\text{NE}, \text{H}, \text{DS}, m, n, sk](K)$</u> $h \leftarrow \text{H}(m)$; $m_s \leftarrow \text{"Keybase-Chat-2"} \parallel \langle K, n, h \rangle$ $s \leftarrow \text{DS.Sig}(sk, m_s)$; $m_e \leftarrow s \parallel m$; Return m_e	<u>$\text{ENC-DER}[m](K)$</u> Return m
--	--

Then $\Phi_{\text{SP}} = \text{MSG-DER-FUNC}[\text{NE}, \text{H}, \text{DS}]$ is the class of all message-deriving functions of these forms.

Note that SIGENC-DER only uses K as a part of the message m_s signed by DS.Sig . Keybase instantiates DS with Ed25519 which computes two SHA-512 hashes of m_s (mixed with other inputs). The resulting signature does not depend on m_s in any other way. Using this observation and an indistinguishability result of Bellare, Davis, and Di [10] (for SHA-512 with output reduced modulo a prime) we capture SIGENC-DER as a special class of message-deriving functions for which we can prove security in the random oracle model.

KDMAE Security for Messages Derived from a Hashed Key. Let H be a hash function. Let Φ be a class of message-deriving functions such that each $\phi \in \Phi$ on input K is only allowed to derive messages from the hash value $H(K)$, and never directly from K . We will roughly show that every AEAD-secure nonce-based encryption scheme NE is also Φ -KDMAE-secure, provided that H is modeled as a random oracle. We formalize this class of functions as follows.

Definition 4. We say Φ derives messages from a hashed key if there exists a set Γ and a function H (modeled as a random oracle) such that $\Phi = \{\phi_\gamma \mid \phi_\gamma(\cdot) = \gamma(H(\cdot)), \gamma \in \Gamma\}$.

In the full version we show how to capture Φ_{SP} as satisfying this definition. Thereby, the following result will give us Φ_{SP} -KDMAE security.

Proposition 1. Let NE be a nonce-based encryption scheme. Let Φ be a class of message-deriving functions that derives messages from a hash key. Let \mathcal{A}_{KDMAE} be an adversary against the Φ -KDMAE security of NE making $q_{\text{NEWHONGROUP}}$ queries to its NEWHONGROUP oracle. Then we can build adversaries \mathcal{A}_{KR} and \mathcal{A}_{AEAD} such that (in the random oracle model)

$$\text{Adv}_{NE, \Phi}^{\text{KDMAE}}(\mathcal{A}_{KDMAE}) \leq 2 \cdot \text{Adv}_{NE}^{\text{KR}}(\mathcal{A}_{KR}) + \text{Adv}_{NE}^{\text{AEAD}}(\mathcal{A}_{AEAD}) + \frac{q_{\text{NEWHONGROUP}}^2}{2^{\text{NE.kl}}}.$$

The constructed adversaries will not repeat (g, n) across ENC queries, so non-nonce-misuse resistant NE suffices. To prove this, we first assert that a Φ -KDMAE adversary \mathcal{A}_{KDMAE} can never directly query the random oracle on any of the (non-exposed) honest keys; otherwise, we could use \mathcal{A}_{KDMAE} in order to break the key-recovery security of AEAD. But then \mathcal{A}_{KDMAE} cannot distinguish between messages derived from $H(K[g])$ or from some $H^*(g)$. Here H is the actual random oracle and H^* is a simulated random oracle whose output depends on a group's identifier g instead of this group's key $K[g]$. We switch from using $H(K[g])$ to $H^*(g)$, thus breaking the dependency of each challenge message on the corresponding NE key. The AEAD security of NE then guarantees that \mathcal{A}_{KDMAE} cannot guess the challenge bit. The formal proof of Proposition 1 is in the full version.

Out-Group AE Security of SealPacket. We prove wOAE security of SP . The formal proof of Theorem 4 is in the full version.

Theorem 4. Let $\mathcal{M} \subseteq \{0, 1\}^*$. Let $SP = \text{SEAL-PACKET-SS}[H, DS, NE]$ be the symmetric signcryption scheme built from some H, DS, NE as specified in Construction 2. Let $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ be the ciphertext-triviality predicate as defined in Fig. 7. Let $\text{func}_{\text{out}}^\perp$ be the output-guarding function as defined in Fig. 8. Let $\text{wOAE}[\text{ENC}[\mathcal{M}, NE]]$ be the security notion as defined in Definition 3. Let $\Phi_{SP} = \text{MSG-DER-FUNC}[NE, H, DS]$ be the class of message-deriving functions defined in Construction 3. Let $\mathcal{A}_{\text{wOAE-of-SP}}$ be an adversary against the $\text{wOAE}[\text{ENC}[\mathcal{M}, NE]]$ security of SP with respect to $\text{pred}_{\text{trivial}}^{\text{suf-except-user}}$ and $\text{func}_{\text{out}}^\perp$. Then we can build adversaries \mathcal{A}_{KDMAE} and $\mathcal{A}_{\text{SPARSE}}$ such that

$$\text{Adv}_{SP, \text{pred}_{\text{trivial}}^{\text{suf-except-user}}, \text{func}_{\text{out}}^\perp}^{\text{wOAE}[\text{ENC}[\mathcal{M}, NE]]}(\mathcal{A}_{\text{wOAE-of-SP}}) \leq \text{Adv}_{NE, \Phi_{SP}}^{\text{KDMAE}}(\mathcal{A}_{NE}) + 2 \cdot \text{Adv}_{DS, \mathcal{M}}^{\text{SPARSE}}(\mathcal{A}_{\text{SPARSE}}).$$

The OAE security results in Theorems 3 and 4 used the weaker output guarding function $\text{func}_{\text{out}}^{\perp}$. In the full version of this paper, we show that for $\text{SS} \in \{\text{BM}, \text{SP}\}$, the OAE security of SS with respect to $\text{func}_{\text{out}}^{\perp}$ implies its OAE security with respect to the stronger output guarding function $\text{func}_{\text{out}}^{\text{sec}}$.

6 Conclusions

Combining Theorem 1 with Theorem 2 and Theorem 3 with Theorem 4 establishes the in-group unforgeability and out-group authenticated encryption security of Keybase’s `BoxMessage` algorithm. These results rely on some standard security assumptions (unforgeability of Ed25519 and collision resistance of SHA-256) as well as some non-standard assumptions (key-dependent message security of XSalsa20-Poly1305 and sparsity of Ed25519). These non-standard assumptions arose, respectively, from the key cycle in `SealPacket` and the key reuse without explicit context separation `BoxMessage`. While we were able to justify these assumptions, we consider them brittle as they are not well studied, their justifications required ideal models, and (in the case of sparsity) they required properties of the specific messaging encoding format used by Keybase.

The comparative simplicity of our Sign-then-Encrypt construction speaks to the value of formalizing the syntax and security of symmetric signcryption. Explicit goals allow designing schemes in parallel with writing proofs to identify precisely what is needed.

References

1. Albrecht, M., Dowling, B., Jones, D.: Device-oriented group messaging: a formal cryptographic analysis of matrix’core. In: IEEE S&P 2024 (2023)
2. Albrecht, M.R., Celi, S., Dowling, B., Jones, D.: Practically-exploitable cryptographic vulnerabilities in matrix. In: 2023 IEEE Symposium on Security and Privacy (SP), pp. 1419–1436. IEEE Computer Society (2022)
3. Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Modular design of secure group messaging protocols and the security of MLS. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021, pp. 1463–1483. ACM Press, November 2021. <https://doi.org/10.1145/3460120.3484820>
4. Alwen, J., Janneck, J., Kiltz, E., Lipp, B.: The pre-shared key modes of HPKE. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology - ASIACRYPT 2023. Springer, Heidelberg (2023). https://doi.org/10.1007/978-981-99-8736-8_11
5. An, J.H., Dodis, Y., Rabin, T.: On the security of joint signature and encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 83–107. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_6
6. Backes, M., Pfitzmann, B., Scedrov, A.: Key-dependent message security under active attacks - BRSIM/UC-soundness of symbolic encryption with key cycles. In: Sabelfeld, A. (ed.) CSF 2007 Computer Security Foundations Symposium, pp. 112–124. IEEE Computer Society Press (2007). <https://doi.org/10.1109/CSF.2007.23>
7. Balbás, D., Collins, D., Gajland, P.: WhatsApp with sender keys? Analysis, improvements and security proofs. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology - ASIACRYPT 2023, pp. 307–341. Springer, Heidelberg (2023). https://doi.org/10.1007/978-981-99-8733-7_10

8. Barnes, R., Beurdouche, B., Robert, R., Millican, J., Omara, E., Cohn-Gordon, K.: The Messaging Layer Security (MLS) Protocol. RFC 9420, July 2023. <https://doi.org/10.17487/RFC9420>
9. Barnes, R., Bhargavan, K., Lipp, B., Wood, C.A.: Hybrid Public Key Encryption. RFC 9180, February 2022. <https://doi.org/10.17487/RFC9180>
10. Bellare, M., Davis, H., Di, Z.: Hardening signature schemes via derive-then-derandomize: stronger security proofs for EdDSA. In: Boldyreva, A., Kolesnikov, V. (eds.) PKC 2023, Part I. LNCS, vol. 13940, pp. 223–250. Springer, Heidelberg (2023). https://doi.org/10.1007/978-3-031-31368-4_9
11. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: 38th FOCS, pp. 394–403. IEEE Computer Society Press, October 1997. <https://doi.org/10.1109/SFCS.1997.646128>
12. Bellare, M., Keelveedhi, S.: Authenticated and misuse-resistant encryption of key-dependent data. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 610–629. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_35
13. Bellare, M., Meiklejohn, S., Thomson, S.: Key-versatile signatures and applications: RKA, KDM and joint Enc/Sig. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 496–513. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_28
14. Bellare, M., Stepanovs, I.: Security under message-derived keys: Signcryption in iMessage. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 507–537. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-45727-3_17
15. Bernstein, D.J.: The Poly1305-AES message-authentication code. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 32–49. Springer, Heidelberg (2005). https://doi.org/10.1007/11502760_3
16. Bernstein, D.J.: The Salsa20 family of stream ciphers. In: Robshaw, M., Billet, O. (eds.) New Stream Cipher Designs. LNCS, vol. 4986, pp. 84–97. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68351-3_8
17. Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: High-speed high-security signatures. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 124–142. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23951-9_9
18. Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.Y.: High-speed high-security signatures. *J. Cryptogr. Eng.* **2**(2), 77–89 (2012). <https://doi.org/10.1007/s13389-012-0027-1>
19. Black, J., Rogaway, P., Shrimpton, T.: Encryption-scheme security in the presence of key-dependent messages. In: Nyberg, K., Heys, H.M. (eds.) SAC 2002. LNCS, vol. 2595, pp. 62–75. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36492-7_6
20. Brendel, J., Cremers, C., Jackson, D., Zhao, M.: The provable security of Ed25519: theory and practice. In: 2021 IEEE Symposium on Security and Privacy, pp. 1659–1676. IEEE Computer Society Press, May 2021. <https://doi.org/10.1109/SP40001.2021.00042>
21. Furuhashi, S.: Messagepack. <https://msgpack.org/>
22. Keybase: Keybase Book. <https://book.keybase.io/>
23. Keybase: Keybase Book—Chat—Crypto. <https://github.com/keybase/book-content/blob/master/D-docs/04-chat/01-crypto.md?plain=1#L89-L93>
24. Keybase: Keybase client. <https://github.com/keybase/client>
25. Keybase: Keybase client—boxer.go—BoxMessage. <https://github.com/keybase/client/blob/v6.2.2/go/chat/boxer.go/#L1564-L1566>

26. Keybase: Keybase client—codec.go—Design Notes. <https://github.com/keybase/client/blob/v6.2.2/go/chat/signencrypt/codec.go/#L95-L110>
27. Keybase: Keybase stats. <https://web.archive.org/web/20200207065125/https://keybase.io/>. Accessed 28 Feb 2024
28. Marlinspike, M.: Private group messaging, May 2014. <https://signal.org/blog/private-groups/>
29. Rogaway, P.: Nonce-based symmetric encryption. In: Roy, B.K., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 348–359. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-25937-4_22
30. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_23
31. Ryan, K., Pornin, T., Fitzgerald, S.: Keybase protocol security review, February 2019. https://keybase.io/docs-assets/blog/NCC_Group_Keybase_KB2018_Public_Report_2019-02-27_v1.3.pdf
32. Shrimpton, T.: A characterization of authenticated-encryption as a form of chosen-ciphertext security. Cryptology ePrint Archive, Report 2004/272 (2004). <https://eprint.iacr.org/2004/272>
33. Wallez, T., Protzenko, J., Beurdouche, B., Bhargavan, K.: TreeSync: authenticated group management for messaging layer security. In: 32nd USENIX Security Symposium, pp. 1217–1233. USENIX Association, Anaheim, CA, August 2023
34. WhatsApp: Whatsapp encryption overview: Technical white paper, September 2023. <https://www.whatsapp.com/security/WhatsApp-Security-Whitepaper.pdf>
35. Zheng, Y.: Digital signcryption or how to achieve $\text{cost}(\text{signature} \ \& \ \text{encryption}) \ll \text{cost}(\text{signature}) + \text{cost}(\text{encryption})$. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 165–179. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0052234>
36. Zoom: Zoom acquires keybase and announces goal of developing the most broadly used enterprise end-to-end encryption offering, May 2020. <https://blog.zoom.us/zoom-acquires-keybase-and-announces-goal-of-developing-the-most-broadly-used-enterprise-end-to-end-encryption-offering/>