Provably Secure Public-Key Encryption for Length-Preserving Chaumian Mixes

Bodo Möller

Technische Universität Darmstadt, Fachbereich Informatik
`moeller@cdc.informatik.tu-darmstadt.de`

Abstract. Mix chains as proposed by Chaum allow sending untraceable electronic e-mail without requiring trust in a single authority: messages are recursively public-key encrypted to multiple intermediates (mixes), each of which forwards the message after removing one layer of encryption. To conceal as much information as possible when using variable (source routed) chains, all messages passed to mixes should be of the same length; thus, message length should not decrease when a mix transforms an input message into the corresponding output message directed at the next mix in the chain. Chaum described an implementation for such length-preserving mixes, but it is not secure against active attacks. We show how to build practical cryptographically secure length-preserving mixes. The conventional definition of security against chosen ciphertext attacks is not applicable to length-preserving mixes; we give an appropriate definition and show that our construction achieves provable security.

1   Introduction

Chaum's *mix* concept [6] is intended to allow users to send untraceable electronic e-mail without having to trust a single authority. The idea is to use a number of intermediates such that it suffices for just one of these to be trustworthy in order to achieve untraceability. (The sender does not have to decide which particular intermediate he is willing to trust, he just must be convinced that at least one in a given list will behave as expected.) These intermediates, the *mixes,* are remailers accepting public-key encrypted input. Messages must be of a fixed size (shorter messages can be padded, longer messages can be split into multiple parts). To send a message, it is routed through a chain of mixes $M_1, \ldots, M_n$: the sender obtains the public key of each mix; then he recursively encrypts the message (including the address of the final recipient) yielding $E_{M_1}\big(E_{M_2}(\ldots E_{M_n}(payload)\ldots)\big)$ where E_{M_i} denotes encryption with M_i's public key, and sends the resulting ciphertext to mix M_1. (The public-key cryptosystem will typically be hybrid, i.e. involve use of symmetric-key cryptography for bulk data encryption.) Each mix removes the corresponding layer of encryption and forwards the decrypted message to the next mix; thus mix M_n will finally recover *payload*.

Each mix is expected to collect a large batch of messages before forwarding the decryption results. The messages in the batch must be reordered (mixed) to prevent message tracing. It is important to prevent replay attacks – a mix must not process the same message twice, or active adversaries would be able to trace messages by duplicating them at submission to cause multiple delivery. (Time-stamps can be used to limit the timespan for which mixes have to remember which messages they have already processed; see [6] and [7].)

Usually it is desireable to allow *source routing,* i.e. let senders choose mix chains on a per-message basis. This increases the flexibility of the whole scheme: senders can make use of new mixes that go into operation, and they can avoid mixes that appear not to work properly; in particular, they can avoid mixes that suppress messages (be it intentionally or because of technical problems), which might be noticed when sending probe messages to oneself over mix chains. For source routing, in the recursively encrypted message, each layer must contain the address of the next mix in the chain so that each mix knows where to forward the message. A problem with the straightforward implementation of source routing is that messages will shrink as they proceed through the chain, not only because of the forwarding address for each layer that must be included, but also because public-key encryption increases message sizes. For optimal untraceability, we need *length-preserving mixes:* the messages that mixes receive should essentially look like the resulting messages that they forward to other mixes.

A construction of length-preserving mixes is given in [6] (a variant of this is used in the fielded system Mixmaster [15]): mix messages consist of a fixed number of slots of a fixed size. The first slot is public-key encrypted so that it can be read by the first mix in the chain. Besides control information directed at the respective mix (such as the address of the next mix in the chain or, in case of the last mix, the address of the final recipient), decryption yields a symmetric key that the mix uses to decrypt all the other slots. Then slots are shifted by one position: the decryption of the second slot becomes the new first slot, and so on. A new final slot is added consisting of random (garbage) data to obtain a mix message of the desired fixed length, which can the be forwarded to the next mix. On the way through the chain, each mix message consists of a number of slots with valid data followed by enough slots with garbage to fill all available slots; mixes are oblivious of how many slots contain valid data and how many are random. Each mix in the chain, when decrypting and shifting the slots, will "decrypt" the garbage slots already present and add a new final slot, thus increasing the number of random slots by one. We note that while the transformation of an incoming mix message to the corresponding outgoing mix message is *length-preserving,* the decryption step itself is actually *length-expanding* because some of the data obtained by decryption is control data directed at the current mix.

The problem with this method for obtaining a hybrid public-key cryptosystem with length-expanding decryption is that it is not secure against active attacks: assume that an adversary controls at least one mix, and that all senders submit well-formed messages. Now when the victim submits a message, the ad-

versary can mark it by modifying one of the slots. This mark will persist as the message is forwarded through the mix chain: assuming CBC or CFB mode block cipher encryption, certain blocks of the corresponding slot will essentially contain garbage due to the decryption and slot shifting performed by each mix (whereas unmarked messages will not show such defects). If the final mix is controlled by the adversary, the adversary may be able to notice the modification and thus break untraceability.

To rule out such attacks, the hybrid public-key cryptosystem should provide security against adaptive chosen ciphertext attacks (CCA security): that is, it should be secure against an adversary who can request the decryption of arbitrary ciphertexts.

In this paper, we present a secure and practical hybrid construction for length-preserving mixes, which is also more flexible than the slot-based approach. The structure of each layer of encryption resembles that of the public-key encryption scheme DHAES/DHIES ([1], [2]). The standard notion of CCA security for public-key cryptosystems is not applicable to our hybrid public-key cryptosystem with length-expanding decryption because the encryption operation must be defined differently for our application: encryption cannot be treated as an atomic black-box operation that takes a plaintext and returns a ciphertext – in this model, we could not have length-expanding decryption. Rather, our encryption process first is input part of the desired *ciphertext* and determines a corresponding part of what will be the decryption result (it is this step that provides length expansion); then it is input the payload plaintext and finally outputs the complete ciphertext, which includes the portion requested in the first input.

We describe the hybrid construction in section 2. Section 3 discusses appropriate security notions and gives provable security results for the construction.

This paper shows essentially how Mixmaster ([15], [16]) should have been specified to be cryptographically secure against active attacks. (Note that an entirely different model of operation for mix networks is assumed by Ohkubo and Abe [17] and Jakobsson and Juels [11]: these constructions assume the existence of a shared authenticated bulletin board, whereas in the present paper we are interested in an open system that can be used with point-to-point communication by e-mail or similar means.) We use many ideas and techniques described by Cramer and Shoup in [8], but adapt them to fit the new notions of security needed in the context of length-preserving mixes.

1.1   Notation

Strings are binary, i.e. elements of $\{0,1\}^*$. The concatenation of strings s and t is denoted $s \parallel t$. The length of string s is $|s|$. For $|s| \geq w$, $\mathrm{prefix}_w(s)$ denotes the string consisting of the leftmost w bits of s; thus, $\mathrm{prefix}_{|s|}(s \parallel t) = s$. Messages (plaintexts, ciphertexts) are strings.

Algorithms are usually probabilistic (randomized).

2 A Construction for Length-Preserving Mixes

Our construction allows much flexibility in the choice of cryptographic schemes, even in a single chain. The single parameter that must be fixed for all mixes is the desired mix message length ℓ. Also it may be desireable to define a maximum length for the actual message payload, i.e. the part of the plaintext as recovered by the final mix that can be chosen by the sender: as a message proceeds through the mix chain, more and more of the data will be pseudo-random gibberish; the length of the useful part of the final plaintext reveals that chains leaving less than this amount cannot have been used. The length n of the mix chain need not be fixed.

For purposes of exposition, we number the mixes $\mathsf{M}_1, \ldots, \mathsf{M}_n$ according to their position in the chain chosen by the sender (note that the same mix might appear multiple times in one chain). For each mix M_i, the following must be defined and, with the exception of the secret key $\mathsf{SK}_{\mathsf{M}_i}$, known to senders who want to use the mix:

– A *key encapsulation mechanism*

$$\mathsf{KEM}_{\mathsf{M}_i}$$

and a key pair

$$(\mathsf{PK}_{\mathsf{M}_i}, \mathsf{SK}_{\mathsf{M}_i})$$

consisting of a *public key* $\mathsf{PK}_{\mathsf{M}_i}$ and a *secret key* $\mathsf{SK}_{\mathsf{M}_i}$ for this key encapsulation mechanism. A key encapsulation mechanism, similarly to a public-key encryption mechanism, provides an algorithm

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}$$

using the public key and an algorithm

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Decrypt}$$

using the secret key. Associated with $\mathsf{KEM}_{\mathsf{M}_i}$ is a *key generation algorithm*

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{KeyGen}$$

that returns appropriate key pairs $(\mathsf{PK}, \mathsf{SK})$. (Note that our formalization uses no explicit security parameter, so desired key sizes are implicit to $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{KeyGen}$.) In contrast with public-key encryption, the $\mathsf{Encrypt}$ algorithm of a key encapsulation mechanism takes no input apart from the public key: $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK}_{\mathsf{M}_i})$ generates a ciphertext $\mathfrak{K}$ of a fixed length

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen}$$

corresponding to a (pseudo-)random string K of a fixed length

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{OutLen}$$

and outputs the pair $(K, \mathfrak{K})$. Evaluation of $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Decrypt}(\mathsf{SK}_{\mathsf{M}_i}, \mathfrak{K})$ will return said string K if the key pair $(\mathsf{PK}_{\mathsf{M}_i}, \mathsf{SK}_{\mathsf{M}_i})$ has been produced by $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{KeyGen}$. In this sense, ciphertext $\mathfrak{K}$ encapsulates the random message K (which can be used as a key for symmetric-key cryptography). On

arbitrary inputs $\mathfrak{K}$, the computation $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Decrypt}(\mathsf{SK}_{\mathsf{M}_i}, \mathfrak{K})$ may either return some string K or fail (return the special value invalid).

One example of a key encapsulation mechanism is the Diffie-Hellman key exchange [9] with static keys for one party and ephemeral keys for the other party where the concatenation of the ephemeral public key with the common secret group element is hashed to obtain K (cf. [1], [19]); $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen}$ can be kept particularly small for elliptic curve Diffie-Hellman using just x coordinates of points (see [14]). Specifications for various key encapsulation mechanisms can be found in [19].

– A *one-time message authentication code*

$$\mathsf{MAC}_{\mathsf{M}_i}$$

with key length

$$\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen}$$

and output length

$$\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}.$$

A one-time message authentication code specifies an efficient deterministic algorithm that takes as input a key K of length $\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen}$ and a bit string s and returns a string $\mathsf{MAC}_{\mathsf{M}_i}(K, s)$ of length $\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}$. In our construction, $\mathsf{MAC}_{\mathsf{M}_i}$ will only be used for strings s of fixed length $\ell - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}$.

Candidate one-time message authentication codes are UHASH [12][1] and HMAC [3].

– A *pseudo-random bit string generator*

$$\mathsf{STREAM}_{\mathsf{M}_i}$$

taking as input a key K of length

$$\mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{KeyLen}$$

and deterministically generating an output string $\mathsf{STREAM}_{\mathsf{M}_i}(K)$ of length

$$\mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{OutLen}.$$

This will be used for an XOR-based stream cipher. A convenient example implementation is the so-called *counter mode* of a block cipher (the output sequence is the prefix of appropriate length of $E_K(0) \parallel E_K(1) \parallel E_K(2) \parallel \ldots$ where E_K denotes block cipher encryption using key K; see [4] and [13]).

– An integer

$$\mathsf{PlainLen}_{\mathsf{M}_i}$$

specifying the length of the prefix of each decrypted message that is considered control data directed to mix $\mathsf{M}_{\mathsf{M}_i}$ and will not be forwarded. (This is the amount of message expansion: the decrypted message minus the prefix must be of size ℓ because that is what will be sent to the next mix.)

---

[1] UHASH is the combination of a *key derivation function* with an *almost strongly universal hash function* [20] and is the core of UMAC as specified in [12]. Note that the security arguments provided for the earlier version of UMAC described in [5] are based on a different approach.

The parameters must fulfil the following conditions:

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} + \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} + \mathsf{PlainLen}_{\mathsf{M}_i} < \ell \qquad (1)$$

$$\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{OutLen} = \mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{KeyLen} + \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen} \qquad (2)$$

$$\mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{OutLen} = \mathsf{PlainLen}_{\mathsf{M}_i} + \ell \qquad (3)$$

2.1 Encryption

We now describe encryption for sending a message through a chain $\mathsf{M}_1, \ldots, \mathsf{M}_n$. Let *payload* be the message of length

$$|payload| = \ell - \sum_{1 \le i \le n} (\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} + \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} + \mathsf{PlainLen}_{\mathsf{M}_i}) \quad (4)$$

(messages shorter than this maximum should be randomly padded on the right). For each i, let $plain_i$ be the control message of length $\mathsf{PlainLen}_{\mathsf{M}_i}$ directed to the respective mix. The encryption algorithm for these arguments is denoted

$$chain\_encrypt_{\mathsf{M}_1,\ldots,\mathsf{M}_n}(plain_1, \ldots, plain_n; payload)$$

or

$$chain\_encrypt_{\mathsf{M}_1,\ldots,\mathsf{M}_n}(plain_1, \ldots, plain_n; payload; \lambda)$$

where λ is the empty string. The algorithm is defined recursively. Let $1 \le i \le n$, and let C_i be a string of length

$$|C_i| = \sum_{1 \le k < i} (\mathsf{KEM}_{\mathsf{M}_k}.\mathsf{CipherLen} + \mathsf{MAC}_{\mathsf{M}_k}.\mathsf{OutLen} + \mathsf{PlainLen}_{\mathsf{M}_k}) \qquad (5)$$

(thus specifically $|C_1| = 0$). Then algorithm

$$chain\_encrypt_{\mathsf{M}_i,\ldots,\mathsf{M}_n}(plain_i, \ldots, plain_n; payload; C_i)$$

works as follows:

1. Use $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK}_{\mathsf{M}_i})$ to generate a pair $(K_i, \mathfrak{K}_i)$.
2. Split K_i in the form

$$K_i = K_{i,\mathsf{MAC}} \parallel K_{i,\mathsf{STREAM}}$$

   such that $|K_{i,\mathsf{MAC}}| = \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen}$ (so $|K_{i,\mathsf{STREAM}}| = \mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{KeyLen}$ by (2)).
3. Compute $\mathsf{STREAM}_{\mathsf{M}_i}(K_{i,\mathsf{STREAM}})$ and split this string in the form

$$stream_{i,\mathsf{L}} \parallel stream_{i,\mathsf{R}}$$

   such that the left part $stream_{i,\mathsf{L}}$ is of length

$$\ell - |C_i| - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}.$$

4. If $i = n$ (last mix), then by (4) and (5) it follows that $|stream_{n,\mathsf{L}}| = \mathsf{PlainLen}_n + |payload|$. In this case, set

$$\mathfrak{C}_n = \big(stream_{n,\mathsf{L}} \oplus (plain_n \parallel payload)\big) \parallel C_n.$$

   Otherwise, let

$$C_{i+1} = stream_{i,\mathsf{R}} \oplus (C_i \parallel 0^{\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen}+\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}+\mathsf{PlainLen}_{\mathsf{M}_i}}),$$

by recursion compute

$$x_i = chain\_encrypt_{\mathsf{M}_{i+1},\ldots,\mathsf{M}_n}(plain_{i+1}, \ldots, plain_n; payload; C_{i+1}),$$

and let

$$\mathfrak{C}_i = \left(stream_{i,\mathsf{L}} \oplus (plain_i \parallel \mathrm{prefix}_{|stream_{i,\mathsf{L}}|-\mathsf{PlainLen}_{\mathsf{M}_i}}(x_i))\right) \parallel C_i.$$

5. Compute $\mathfrak{M}_i = \mathsf{MAC}_i(K_{i,\mathsf{MAC}}, \mathfrak{C}_i)$.
6. Return the ciphertext $\mathfrak{K}_i \parallel \mathfrak{M}_i \parallel \mathfrak{C}_i$, which is of length ℓ.

Appendix A.1 illustrates a ciphertext generated by $chain\_encrypt_{\mathsf{M}_1,\ldots,\mathsf{M}_n}$.

2.2 Decryption

The decryption algorithm for mix M_i ($1 \leq i \leq n$) works as follows, given a length-ℓ ciphertext $\mathfrak{K} \parallel \mathfrak{M} \parallel \mathfrak{C}$ (split into its three components according to parameters $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} = |\mathfrak{K}|$ and $\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} = |\mathfrak{M}|$). We denote it

$$mix\_decrypt_{\mathsf{M}_i}(\mathfrak{K} \parallel \mathfrak{M} \parallel \mathfrak{C}).$$

Remember that M_i in general is not aware of the mix chain used by the sender or even of its own position i in the chain.

1. Compute $K = \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Decrypt}(\mathsf{SK}_{\mathsf{M}_i}, \mathfrak{K})$. If this computation fails, abort with an error (return invalid).
2. Split K in the form

$$K = K_{\mathsf{MAC}} \parallel K_{\mathsf{STREAM}}$$

such that $|K_{\mathsf{MAC}}| = \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen}$.
3. Compute

$$\widetilde{\mathfrak{M}} = \mathsf{MAC}_{\mathsf{M}_i}(K_{\mathsf{MAC}}, \mathfrak{C})$$

and test whether

$$\widetilde{\mathfrak{M}} = \mathfrak{M}.$$

If this is not the case, abort with an error (return invalid).
4. Compute the string

$$stream = \mathsf{STREAM}_{\mathsf{M}_i}(K_{\mathsf{STREAM}})$$

of length $\mathsf{STREAM}_{\mathsf{M}_i}.\mathsf{OutLen}$.
5. Compute

$$stream \oplus (\mathfrak{C} \parallel 0^{\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen}+\mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}+\mathsf{PlainLen}_{\mathsf{M}_i}})$$

and split the resulting string in the form

$$plain_i \parallel P$$

where $plain_i$ is of length $\mathsf{PlainLen}_{\mathsf{M}_i}$ (and thus, by (3), P is of length ℓ).
6. Return the pair $(plain_i, P)$.

If this algorithm finishes without an error, $plain_i$ is a control message directed to the current mix, and P is the message to be forwarded to another mix or to the final recipient (as requested by the control message).

It is straightforward to verify that for ciphertexts computed as

$$chain\_encrypt_{\mathsf{M}_1,\dots,\mathsf{M}_n}(plain_1,\dots,plain_n; payload),$$

iterative decryption by mixes $\mathsf{M}_1,\dots,\mathsf{M}_n$ will indeed work without an error and recover the respective strings $plain_i$ and finally also the message $payload$ concatenated with some (useless) extra data. Appendix A.2 illustrates decryption.

3   Provable Security

The mix concept is intended to provide security when at least one mix can be trusted and behaves correctly. (However note that denial-of-service attacks by incorrectly operating mixes cannot be ruled out; the only option is to avoid mixes that appear to malfunction.) Thus we assume that some single mix M_i works as expected while all other mixes are controlled by the adversary and may not follow the protocol (also the cryptographic schemes $\mathsf{KEM}_{\mathsf{M}_j}$, $\mathsf{MAC}_{\mathsf{M}_j}$, and $\mathsf{STREAM}_{\mathsf{M}_j}$ associated with mixes M_j, $j \neq i$, might be not secure, and $\mathsf{KEM}_{\mathsf{M}_j}$ might even not be a valid key encapsulation mechanism according to the description given in section 2).

This leaves only M_i to be attacked: outer layers of encryption for mixes that appear before M_i in a mix chain are easily removed by the adversary and thus are not relevant for security; and inner layers of encryption for mixes that appear after M_i in a mix chain are involved in the encryption process

$$chain\_encrypt_{\mathsf{M}_i,\dots,\mathsf{M}_n}(plain_i,\dots,plain_n; payload; C_i)$$

described in section 2.1, but cannot provide protection against the adversary.

An important security property for mixes is *unlinkability:* an adversary who observes a batch of encrypted messages as these are sent to a mix and who sees the resulting decrypted messages when these are forwarded should not be able to tell better than by chance which decrypted messages correspond to which encrypted messages.

Beyond this basic security notion, we also want to achieve security against active attacks. As noted in the introduction, it is crucial to detect message replay and ignore replayed message. However, this does not rule out active attacks based on message modifications. To show that our construction for length-preserving mixes is secure against this kind of attacks, we will model an active adversary who is able to launch an *adaptive chosen ciphertext attack (CCA).*

We point out that in the model of operation that we assume, unlinkability cannot be fully guaranteed in the presence of active attacks: an active adversary who suppresses all but a single one of the legitimate encrypted messages in a batch and substitutes new messages for them can easily trace the remaining legitimate message. This can be avoided only heuristically (each sender of mix messages should occasionally route a message to himself to see if it gets through). To avoid related flooding attacks, mixes should not start a new batch whenever

8

a certain fixed number of messages has arrived; instead, each batch should be kept open for additional messages during some time frame.

We show how the security of the mix construction can be captured formally by describing unlinkability (section 3.1) and CCA security (section 3.2). We then provide security definitions for the underlying key encapsulation method (section 3.3), one-time message authentication code (section 3.4), and pseudorandom bit string generator (section 3.5). Finally, section 3.6 present security results that relates the security of the mix encryption scheme to the security of these cryptographic schemes: we will see that if the mix encryption scheme is not secure, then this is because one of the underlying cryptographic schemes is not secure.

3.1 Unlinkability of the Mix Encryption Scheme

An intuitive approach to defining unlinkability is as follows: the adversary sees a batch of m ciphertexts and a random permutation of the resulting m decryption results; for each decryption result, the adversary makes a guess on the index of the corresponding ciphertext. With random guessing, the expected number of correct guesses is one. An adversary who can do better than this breaks unlinkability.

More specifically, we let the adversary select the plaintexts as far as possible (remember that as we have length-expanding decryption, it is not possible to choose all of the decryption result arbitrarily when generating a ciphertext). In the formal definition, we confine to a setting with two plaintexts where only one ciphertext is shown to the adversary (an adversary in the unrestricted setting can be used to build an adversary in this setting). This is captured in the following attack game:

1. The adversary queries a *key generation oracle*, which uses $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{KeyGen}$ to compute a key pair

$$(\mathsf{PK}, \mathsf{SK})$$

   and responds with PK (and secretly stores SK).
2. The adversary uses an *encryption oracle* as follows (cf. the encryption algorithm from section 2.1 in the case of a length-1 mix chain, i.e. with no recursion and $|C_i| = 0$):

   – The adversary submits a pair of string pairs

$$\big((plain_{i,0}, m_0), (plain_{i,1}, m_1)\big)$$

   where

$$|plain_{i,b}| = \mathsf{PlainLen}_{\mathsf{M}_i}$$

   and

$$|m_b| = \ell - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} - \mathsf{PlainLen}_{\mathsf{M}_i}$$

   for $b = 0, 1$.

– The encryption oracle chooses a bit $b \in \{0, 1\}$ uniformly at random. Then it uses $\mathsf{KEM_{M_i}.Encrypt(PK)}$ to generate a pair $(K_{\mathsf{oracle}}, \mathfrak{K}_{\mathsf{oracle}})$ and splits K_{oracle} in the form

$$K_{\mathsf{oracle}} = K_{\mathsf{MAC}} \parallel K_{\mathsf{STREAM}}$$

such that $|K_{\mathsf{MAC}}| = \mathsf{MAC_{M_i}.KeyLen}$; it computes $\mathsf{STREAM_{M_i}}(K_{\mathsf{STREAM}})$ and splits the resulting string $stream$ in the form

$$stream = stream_{\mathsf{L}} \parallel stream_{\mathsf{R}}$$

such that $|stream_{\mathsf{L}}| = \ell - \mathsf{KEM_{M_i}.CipherLen} - \mathsf{MAC_{M_i}.OutLen}$; and it computes

$$\mathfrak{C} = stream_{\mathsf{L}} \oplus (plain_{i,b} \parallel m_b)$$

and

$$\mathfrak{M} = \mathsf{MAC_{M_i}}(K_{\mathsf{MAC}}, \mathfrak{C}).$$

Then it outputs the pair

$$(\mathfrak{K}_{\mathsf{oracle}} \parallel \mathfrak{M} \parallel \mathfrak{C}, stream_{\mathsf{R}}).$$

3. The adversary outputs a bit $\widetilde{b} \in \{0, 1\}$.

The bit $\widetilde{b}$ output by the adversary is its guess for the value of b. Note that the decryption result corresponding to the ciphertext $\mathfrak{K}_{\mathsf{oracle}} \parallel \mathfrak{M} \parallel \mathfrak{C}$ is

$$(plain_{i,b}, m_b \parallel stream_{\mathsf{R}})$$

(see section 2.2), and all of this except for the choice of the bit b is known to the adversary.

Let A be any adversary (interactive probabilistic algorithm with bounded runtime) in this attack game. Its *advantage against unlinkability* for M_i's instantiation of the mix encryption scheme is

$$\mathsf{AdvLink_{M_i, A}} = \left| \Pr\left[\widetilde{b} = 1 \mid b = 1\right] - \Pr\left[\widetilde{b} = 1 \mid b = 0\right] \right|.$$

3.2 CCA Security of the Mix Encryption Scheme

Due to the recursive nature of our encryption algorithm for mix chains, we cannot directly apply the usual definitions of security under adaptive chosen ciphertext attack (CCA) for ordinary public-key encryption. We adapt the attack game described in [8, section 3.2] (which goes back to [10] and [18]) as follows to take into account the special properties of our construction:

1. The adversary queries a *key generation oracle*, which uses $\mathsf{KEM_{M_i}.KeyGen}$ to compute a key pair

$$(\mathsf{PK}, \mathsf{SK})$$

and responds with PK (and secretly stores SK).

2. The adversary makes a sequence of queries to a *decryption oracle.* Each query is an arbitrary string s of length ℓ, and the oracle responds with

$$mix\_decrypt_{\mathsf{M}_i}(s),$$

using the secret key SK from step 1. That is, each oracle response is either a pair $(plain, P)$ where $|plain| = \mathsf{PlainLen}_{\mathsf{M}_i}$ and $|P| = \ell$, or the special value invalid.

3. The adversary uses an *interactive encryption oracle* as follows (compare with the encryption algorithm in section 2.1):

   – First the adversary submits some string C_i subject only to the condition that

   $$0 \le |C_i| < \ell - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}.$$

   – The interactive encryption oracle uses $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK})$ to generate a pair $(K_{\mathsf{oracle}}, \mathfrak{K}_{\mathsf{oracle}})$ and splits K_{oracle} in the form

   $$K_{\mathsf{oracle}} = K_{\mathsf{MAC}} \parallel K_{\mathsf{STREAM}}$$

   such that $|K_{\mathsf{MAC}}| = \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{KeyLen}$; it computes $\mathsf{STREAM}_{\mathsf{M}_i}(K_{\mathsf{STREAM}})$ and splits the resulting string $stream$ in the form

   $$stream = stream_{\mathsf{L}} \parallel stream_{\mathsf{R}}$$

   such that $|stream_{\mathsf{L}}| = \ell - |C_i| - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen}$; and it computes

   $$C_{i+1} = stream_{\mathsf{R}} \oplus \left(C_i \parallel 0^{\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} + \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} + \mathsf{PlainLen}_{\mathsf{M}_i}}\right)$$

   and sends C_{i+1} to the adversary.

   – The adversary submits a pair of string pairs

   $$\left((plain_{i,0}, m_0), (plain_{i,1}, m_1)\right)$$

   satisfying

   $$|plain_{i,b}| = \mathsf{PlainLen}_{M_i}$$

   and

   $$|m_b| = \ell - |C_i| - \mathsf{KEM}_{\mathsf{M}_i}.\mathsf{CipherLen} - \mathsf{MAC}_{\mathsf{M}_i}.\mathsf{OutLen} - \mathsf{PlainLen}_{\mathsf{M}_i}$$

   for $b = 0, 1$.

   – The interactive encryption oracle chooses a uniformly random bit $b \in \{0, 1\}$, determines

   $$\mathfrak{C} = \left(stream_{\mathsf{L}} \oplus (plain_{i,b} \parallel m_b)\right) \parallel C_i$$

   and

   $$\mathfrak{M} = \mathsf{MAC}_{\mathsf{M}_i}(K_{\mathsf{MAC}}, \mathfrak{C}),$$

   and responds with the *challenge ciphertext*

   $$\mathfrak{K}_{\mathsf{oracle}} \parallel \mathfrak{M} \parallel \mathfrak{C}.$$

4. The adversary again makes a sequence of queries to a *decryption oracle* as in step 2, except that this time the decryption oracle refuses being asked for the challenge ciphertext $\mathfrak{K}_{\mathsf{oracle}} \parallel \mathfrak{M} \parallel \mathfrak{C}$ from step 3 (it returns invalid for this case).

5. The adversary outputs a bit $\widetilde{b} \in \{0, 1\}$.

The bit $\widetilde{b}$ output by the adversary is its guess for the value of b.

The essential difference to the adaptive chosen ciphertext attack game for ordinary public-key encryption is that we have made the encryption oracle interactive to reflect the recursiveness of the encryption algorithm for mix chains, where encryption to mixes later in the chain than M_i can have potentially arbitrary effects on a large part of the plaintext. Appendix A.3 illustrates the ciphertext resulting from the invocation of the interactive encryption oracle.

Let A be any adversary (interactive probabilistic algorithm with bounded runtime) in the above attack game. Its *CCA advantage* against M_i's instantiation of the mix encryption scheme is

$$\mathsf{AdvCCA}_{M_i, A} = \left| \Pr\left[\widetilde{b} = 1 \mid b = 1\right] - \Pr\left[\widetilde{b} = 1 \mid b = 0\right] \right|.$$

3.3 Security of the Key Encapsulation Mechanism

CCA security for the key encryption mechanism KEM_{M_i} is defined through the following attack game (cf. [8, section 7.1.2]):

1. The adversary queries a *key generation oracle*, which uses $\mathsf{KEM}_{M_i}.\mathsf{KeyGen}$ to compute a key pair $(\mathsf{PK}, \mathsf{SK})$ and responds with PK (and secretly stores SK).
2. The adversary makes a sequence of queries to a *decryption oracle.* Each query is an arbitrary string s of length $\mathsf{KEM}_{M_i}.\mathsf{CipherLen}$; the oracle responds with

$$\mathsf{KEM}_{M_i}.\mathsf{Decrypt}(\mathsf{SK}, s).$$

   Thus each oracle response is either a string of length $\mathsf{KEM}_{M_i}.\mathsf{OutLen}$ or the special value invalid.
3. The adversary queries an *encryption oracle,* which works as follows: it uses $\mathsf{KEM}_{M_i}.\mathsf{Encrypt}(\mathsf{PK})$ to obtain a pair $(K_0, \mathfrak{K}_{\mathsf{oracle}})$, it generates a uniformly randomly string K_1 such that $|K_0| = |K_1|$, chooses a uniformly random bit $b_{\mathsf{KEM}} \in \{0, 1\}$, and responds with $(K_{b_{\mathsf{KEM}}}, \mathfrak{K}_{\mathsf{oracle}})$.
4. The adversary again makes a sequence of queries to a *decryption oracle* as in step 2, but this time the oracle refuses the specific query $\mathfrak{K}_{\mathsf{oracle}}$ and responds invalid in this case.
5. The adversary outputs a bit $\widetilde{b}_{\mathsf{KEM}} \in \{0, 1\}$.

The *CCA advantage* of an adversary A (an interactive probabilistic algorithm with bounded runtime) against KEM_{M_i} in this attack game is

$$\mathsf{AdvCCA}_{\mathsf{KEM}_{M_i}, A} = \left| \Pr\left[\widetilde{b}_{\mathsf{KEM}} = 1 \mid b_{\mathsf{KEM}} = 1\right] - \Pr\left[\widetilde{b}_{\mathsf{KEM}} = 1 \mid b_{\mathsf{KEM}} = 0\right] \right|.$$

3.4 Security of the One-Time Message Authentication Code

To define the security of MAC_{M_i}, we use the following attack game (cf. [8, section 7.2.2]):

1. The adversary submits a string s (which for our purposes may be assumed to be of fixed length $\ell - \mathsf{KEM_{M_i}.CipherLen} - \mathsf{MAC_{M_i}.OutLen}$) to an oracle. The oracle generates a uniformly random string K of length $\mathsf{MAC_{M_i}.KeyLen}$ and responds with $\mathsf{MAC_{M_i}}(K, s)$.
2. The adversary outputs a list $(s_1, t_1), (s_2, t_2), \ldots, (s_m, t_m)$ of pairs of strings.

An adversary A again is an interactive probabilistic algorithm with bounded runtime; the runtime bound implies a bound on the list length m. We say that adversary A has *produced a forgery* if $s_k \neq s$ and $\mathsf{MAC_{M_i}}(K, s_k) = t_k$ for some k $(1 \leq k \leq m)$. Its advantage against $\mathsf{MAC_{M_i}}$, denoted $\mathsf{AdvForge_{MAC_{M_i}, A}}$, is the probability that it produces a forgery in the above game.

3.5 Security of the Pseudo-Random Bit String Generator

To define the security of $\mathsf{STREAM_{M_i}}$, we use the following attack game:

1. The adversary queries an oracle. The oracle generates a uniformly random string K of length $\mathsf{STREAM_{M_i}.KeyLen}$, computes $stream_0 = \mathsf{STREAM_{M_i}}(K)$, generates a uniformly random string $stream_1$ with $|stream_0| = |stream_1|$, chooses a uniformly random bit $b_{\mathsf{STREAM}} \in \{0, 1\}$, and responds with $stream_b$.
2. The adversary outputs a bit $\widetilde{b}_{\mathsf{STREAM}} \in \{0, 1\}$.

The advantage of an adversary A (again an interactive probabilistic algorithm with bounded runtime) against $\mathsf{STREAM_{M_i}}$ is

$$\mathsf{Adv_{STREAM_{M_i}, A}}$$
$$= \left| \Pr\left[\widetilde{b}_{\mathsf{STREAM}} = 1 \mid b_{\mathsf{STREAM}} = 1\right] - \Pr\left[\widetilde{b}_{\mathsf{STREAM}} = 1 \mid b_{\mathsf{STREAM}} = 0\right] \right|.$$

3.6 Security Results

First let A be an adversary against the unlinkability of the mix encryption scheme, attacking a mix $\mathsf{M_i}$ as described in section 3.1. It can be shown that there are adversaries A_1 and A_2 against $\mathsf{KEM_{M_i}}$ and $\mathsf{STREAM_{M_i}}$, respectively, with essentially the same runtime as A such that

$$\mathsf{AdvLink_{M_i, A}} \leq 2 \cdot \left(\mathsf{AdvCCA_{KEM_{M_i}, A_1}} + \mathsf{Adv_{STREAM_{M_i}, A_2}} \right).$$

Details of the proof can be found in appendix B.1.

Now let A be an adversary against the CCA security of the mix encryption scheme, attacking a mix $\mathsf{M_i}$ as described in section 3.2. It can be shown that there are adversaries A_1, A_2, A_3 against $\mathsf{KEM_{M_i}}$, $\mathsf{MAC_{M_i}}$, $\mathsf{STREAM_{M_i}}$ all having essentially the same runtime as A such that

$$\mathsf{AdvCCA_{M_i, A}} \leq 2 \cdot \left(\mathsf{AdvCCA_{KEM_{M_i}, A_1}} + \mathsf{AdvForge_{MAC_{M_i}, A_2}} + \mathsf{Adv_{STREAM_{M_i}, A_3}} \right).$$

Details of the proof are given in appendix B.2.

Acknowledgement

References

1. ABDALLA, M., BELLARE, M., AND ROGAWAY, P. DHAES: An encryption scheme based on the Diffie-Hellman problem. Submission to IEEE P1363a. http://grouper.ieee.org/groups/1363/P1363a/Encryption.html, 1998.

2. ABDALLA, M., BELLARE, M., AND ROGAWAY, P. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In *Progress in Cryptology - CT-RSA 2001* (2001), D. Naccache, Ed., vol. 2020 of *Lecture Notes in Computer Science*, pp. 143–158.

3. BELLARE, M., CANETTI, R., AND KRAWCZYK, H. Keying hash functions for message authentication. In *Advances in Cryptology – CRYPTO '96* (1996), N. Koblitz, Ed., vol. 1109 of *Lecture Notes in Computer Science*, pp. 1–15.

4. BELLARE, M., DESAI, A., JOKIPII, E., AND ROGAWAY, P. A concrete security treatment of symmetric encryption. In *38th Annual Symposium on Foundations of Computer Science (FOCS '97)* (1997), IEEE Computer Society, pp. 394–403.

5. BLACK, J., HALEVI, S., KRAWCZYK, H., KROVETZ, T., AND ROGAWAY, P. UMAC: Fast and secure message authentication. In *Advances in Cryptology – CRYPTO '99* (1999), M. Wiener, Ed., vol. 1666 of *Lecture Notes in Computer Science*, pp. 216–233.

6. CHAUM, D. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM 24* (1981), 84–88.

7. COTTRELL, L. Mixmaster & remailer attacks. http://www.obscura.com/%7Eloki/remailer/remailer-essay.html, 1997.

8. CRAMER, R., AND SHOUP, V. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. Manuscript, http://shoup.net/papers/, 2001.

9. DIFFIE, W., AND HELLMAN, M. E. New directions in cryptography. *IEEE Transactions on Information Theory 22*, 6 (1976), 644–654.

10. GOLDWASSER, S., AND MICALI, S. Probabilistic encryption. *Journal of Computer and System Sciences 28* (1984), 270–299.

11. JAKOBSSON, M., AND JUELS, A. An optimally robust hybrid mix network. In *20th Annual ACM Symposium on Principles of Distributed Computing (PODC 2001)* (2001), ACM Press, pp. 284–292.

12. KROVETZ, T., BLACK, J., HALEVI, S., HEVIA, A., KRAWCZYK, H., AND ROGAWAY, P. UMAC: Message authentication code using universal hashing. Internet-Draft draft-krovetz-umac-01.txt, http://www.cs.ucdavis.edu/~rogaway/umac/, 2000.

13. LIPMAA, H., ROGAWAY, P., AND WAGNER, D. Comments to NIST concerning AES modes of operation: CTR-mode encryption. http://csrc.nist.gov/encryption/modes/workshop1/papers/lipmaa-ctr.pdf, 2000.

14. MILLER, V. S. Use of elliptic curves in cryptography. In *Advances in Cryptology – CRYPTO '85* (1986), H. C. Williams, Ed., vol. 218 of *Lecture Notes in Computer Science*, pp. 417–428.

15. Mixmaster anonymous remailer software. `http://sourceforge.net/projects/mixmaster/`.

16. MÖLLER, U., AND COTTRELL, L. Mixmaster protocol version 2. `http://www.eskimo.com/~rowdenw/crypt/Mix/draft-moeller-v2-01.txt`, 2000.

17. OHKUBO, M., AND ABE, M. A length-invariant hybrid mix. In *Advances in Cryptology – ASIACRYPT 2000* (2000), T. Okamoto, Ed., vol. 1976 of *Lecture Notes in Computer Science*, pp. 178–191.

18. RACKOFF, C. W., AND SIMON, D. R. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Advances in Cryptology – CRYPTO '91* (1992), J. Feigenbaum, Ed., vol. 576 of *Lecture Notes in Computer Science*, pp. 433–444.

19. SHOUP, V. A proposal for an ISO standard for public key encryption. Version 2.1, December 20, 2001. `http://shoup.net/papers/`.

20. WEGMAN, M. N., AND CARTER, J. L. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences 22* (1981), 265–279.

A   Illustrations

A.1   Encryption

We depict a ciphertext generated by the encryption algorithm

$$chain\_encrypt_{\mathsf{M_1,M_2,M_3}}(plain_1, plain_2, plain_3; payload)$$

as described in section 2.1. In the illustration, we have concatenation horizontally and XOR vertically (i.e. boxes in the same row represent bit strings that are concatenated, and the ciphertext is the XOR of the multiple rows shown).

$\mathfrak{K}_1$	$\mathfrak{M}_1$							$plain_3$	$payload$

(illustration table representing the encryption layout:)

|		$plain_1$	$\mathfrak{K}_2$	$\mathfrak{M}_2$	$plain_2$	$\mathfrak{K}_3$	$\mathfrak{M}_3$	$plain_3$	$payload$
									$stream_{3,\mathsf{L}}$
							$stream_{2,\mathsf{L}}$		
$\mathfrak{K}_1$	$\mathfrak{M}_1$	$stream_{1,\mathsf{L}}$							

A.2   Decryption

The result obtained by mix M_1 when applying the decryption algorithm from section 2.2 to the ciphertext from appendix A.1 is composed as follows:

$plain_1$	$\mathfrak{K}_2$	$\mathfrak{M}_2$	$plain_2$	$\mathfrak{K}_3$	$\mathfrak{M}_3$	$plain_3$	$payload$
							$stream_{3,\mathsf{L}}$
		$stream_{2,\mathsf{L}}$					
						$stream_{1,\mathsf{R}}$	

The string $plain_1$ is directed to M_1. The remainder of the decryption result is a ciphertext that should be forwarded to the next mix in the chain, M_2, which will then obtain the following result:

15

			$plain_3$	$payload$	
$plain_2$	$\mathfrak{K}_3$	$\mathfrak{M}_3$	$stream_{3,\mathsf{L}}$		
					$stream_{2,\mathsf{R}}$
				$stream_{1,\mathsf{R}}$	

Similarly, mix M_3 will obtain the following final decryption result:

| $plain_3$	$payload$		
		$stream_{3,\mathsf{R}}$	
		$stream_{2,\mathsf{R}}$	
	$stream_{1,\mathsf{R}}$		

A.3 Interactive Encryption Oracle

In the adaptive chosen ciphertext attack game from section 3.2, recursive encryption is replaced by an interactive encryption oracle to generate the challenge ciphertext, which has the following structure:

	$plain_{i,b}$	m_b	C_i
$\mathfrak{K}_{\mathsf{oracle}}$ $\mathfrak{M}$		$stream_{\mathsf{L}}$	

Applying algorithm $mix\_decrypt_{\mathsf{M}_i}$ from section 2.2 to this ciphertext would yield the following result:

$plain_{i,b}$	m_b	C_i	
			$stream_{\mathsf{R}}$

This can also be written as follows:

$plain_{i,b}$	m_b	C_{i+1}

B Security Proofs

B.1 Unlinkability

We prove the security result from section 3.6 for the unlinkability of the mix encryption scheme. Let $\mathbf{G}_0$ denote the attack game from section 3.1. Let $\mathbf{G}_1$ be like $\mathbf{G}_0$ except that a uniformly random string is used for K_{oracle}, whereas $\mathfrak{K}_{\mathsf{oracle}}$ is still generated by $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK})$. Let $\mathbf{G}_2$ be like $\mathbf{G}_1$ except that the encryption oracle uses a uniformly random string $stream$ instead of computing $\mathsf{STREAM}_{\mathsf{M}_i}(K_{\mathsf{STREAM}})$.

Now we consider an adversary A as in section 3.1, exposed to these different attack games $\mathbf{G}_x$, $0 \leq x \leq 2$, and look at the respective success probabilities $\Pr_{\mathbf{G}_x}\big[\tilde{b} = b\big]$. Based on A, adversaries A_1 and A_2 against $\mathsf{KEM}_{\mathsf{M}_i}$ and $\mathsf{STREAM}_{\mathsf{M}_i}$, respectively, will be built, each with essentially the same runtime as A.

A_1 attacks $\mathsf{KEM}_{\mathsf{M}_i}$ as follows (cf. section 3.3; note that this A_1 never actually uses its decryption oracle). At first, it generates $b \in \{0,1\}$ uniformly at random. Then, it runs the adversary A; when A queries its encryption oracle, A_1 queries its own encryption oracle to obtain a pair $(K_{\mathsf{oracle}}, \mathfrak{K}_{\mathsf{oracle}})$ and performs step 2 from section 3.1 using this pair and the pregenerated bit b. Finally, when A outputs its bit $\widetilde{b}$, A_1 outputs 1 if $\widetilde{b} = b$ and 0 otherwise. Observe that

$$\left| \Pr{}_{\mathbf{G}_1}\left[\widetilde{b} = b\right] - \Pr{}_{\mathbf{G}_0}\left[\widetilde{b} = b\right] \right| = \mathsf{AdvCCA}_{\mathsf{KEM}_{\mathsf{M}_i}, A_1}$$

($\mathbf{G}_0$ corresponds to $b_{\mathsf{KEM}} = 0$, $\mathbf{G}_1$ corresponds to $b_{\mathsf{KEM}} = 1$ in section 3.3).

A_2 attacks $\mathsf{STREAM}_{\mathsf{M}_i}$ as follows (cf. section 3.5). First, it generates $b \in \{0,1\}$ and a string K_{oracle} of length $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{OutLen}$ uniformly at random. Then, it runs the adversary A, playing the role of the key generation oracle and the role of the encryption oracle, wherein it substitutes the pregenerated string K_{oracle} and the pregenerated bit b in step 2 of section 3.1. When A outputs its bit $\widetilde{b}$, A_2 outputs 1 if $\widetilde{b} = b$ and 0 otherwise. Observe that

$$\left| \Pr{}_{\mathbf{G}_2}\left[\widetilde{b} = b\right] - \Pr{}_{\mathbf{G}_1}\left[\widetilde{b} = b\right] \right| = \mathsf{Adv}_{\mathsf{STREAM}_{\mathsf{M}_i}, A_2}$$

($\mathbf{G}_1$ corresponds to $b_{\mathsf{STREAM}} = 0$, $\mathbf{G}_2$ corresponds to $b_{\mathsf{STREAM}} = 1$ in section 3.5).

Finally observe that $\Pr_{\mathbf{G}_2}\left[\widetilde{b} = b\right] = \frac{1}{2}$.

From this we obtain the inequality

$$\begin{aligned}
\mathsf{AdvLink}_{\mathsf{M}_i, A} &= 2 \cdot \left| \frac{1}{2} - \Pr{}_{\mathbf{G}_0}\left[\widetilde{b} = b\right] \right| \\
&= 2 \cdot \left| \sum_{1 \le x \le 2} \left(\Pr{}_{\mathbf{G}_x}\left[\widetilde{b} = b\right] - \Pr{}_{\mathbf{G}_{x-1}}\left[\widetilde{b} = b\right] \right) \right| \\
&\le 2 \cdot \left(\mathsf{AdvCCA}_{\mathsf{KEM}_{\mathsf{M}_i}, A_1} + \mathsf{Adv}_{\mathsf{STREAM}_{\mathsf{M}_i}, A_2} \right),
\end{aligned}$$

which concludes the proof.

B.2 CCA Security

We prove the CCA security result given in section 3.6. Let $\mathbf{G}_0$ denote the attack game from section 3.2. We will modify it in multiple steps, essentially disabling the underlying cryptographic schemes (key encryption method, one-time message authentication code, and pseudo-random bit string generator) one after another.

$\mathbf{G}_1$ is like $\mathbf{G}_0$ except that a uniformly random string is used for K_{oracle}, whereas $\mathfrak{K}_{\mathsf{oracle}}$ is still generated by $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK})$. In the decryption oracle, K is substituted whenever $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Decrypt}(\mathsf{SK}, \mathfrak{K}_{\mathsf{oracle}})$ would have to be computed.

$\mathbf{G}_2$ is like $\mathbf{G}_1$ except that the decryption oracle always responds with invalid when faced with any query prefixed with string $\mathfrak{K}_{\mathsf{oracle}}$. (This applies to both step 2 and step 4 in section 3.2. Thus, the invocation of $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{Encrypt}(\mathsf{PK})$ to generate $(K_{\mathsf{oracle}}, \mathfrak{K}_{\mathsf{oracle}})$ must be advanced from step 3 to an earlier step; in $\mathbf{G}_1$, this is only a descriptive change and does not affect the behaviour observed by the adversary.)

$\mathbf{G}_3$ is like $\mathbf{G}_2$ except that the interactive encryption oracle uses a uniformly random string *stream* instead of computing $\mathsf{STREAM}_{\mathsf{M}_i}(K_{\mathsf{STREAM}})$.

Now we consider an adversary A as in section 3.2, exposed to these different attack games $\mathbf{G}_x$, $0 \leq x \leq 3$, and look at the respective success probabilities $\Pr_{\mathbf{G}_x}[\widetilde{b} = b]$. Based on A, adversaries A_1, A_2, A_3 against $\mathsf{KEM}_{\mathsf{M}_i}$, $\mathsf{MAC}_{\mathsf{M}_i}$, $\mathsf{STREAM}_{\mathsf{M}_i}$ will be built all having essentially the same runtime as A.

A_1 attacks $\mathsf{KEM}_{\mathsf{M}_i}$ as follows (cf. section 3.3). At first, it generates $b \in \{0,1\}$ uniformly at random. Then, it runs the adversary A; when A queries its decryption oracle, A_1 uses its own decryption oracle to perform the decryption algorithm from section 2.2, and when A queries its interactive encryption oracle, A_1 queries its own encryption oracle to obtain a pair $(K_{\mathsf{oracle}}, \mathfrak{K}_{\mathsf{oracle}})$ and performs step 3 from section 3.2 using this pair and the pregenerated bit b. Finally, when A outputs its bit $\widetilde{b}$, A_1 outputs 1 if $\widetilde{b} = b$ and 0 otherwise. Observe that

$$\left| \Pr_{\mathbf{G}_1}[\widetilde{b} = b] - \Pr_{\mathbf{G}_0}[\widetilde{b} = b] \right| = \mathsf{AdvCCA}_{\mathsf{KEM}_{\mathsf{M}_i}, A_1}$$

($\mathbf{G}_0$ corresponds to $b_{\mathsf{KEM}} = 0$, $\mathbf{G}_1$ corresponds to $b_{\mathsf{KEM}} = 1$ in section 3.3).

A_2 attacks $\mathsf{MAC}_{\mathsf{M}_i}$ as follows (cf. section 3.4). At first, it generates $b \in \{0,1\}$ and a string K_{oracle} of length $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{OutLen}$ uniformly at random. Then, it runs the adversary A, playing the roles of the key generation oracle, decryption oracle, and interactive encryption oracle, substituting the pregenerated string K_{oracle} and the pregenerated bit b in step 3 of section 3.2. Whenever A submits a query $\mathfrak{K} \parallel \mathfrak{M} \parallel \mathfrak{C}$ to the decryption oracle, A_2 adds the pair $(\mathfrak{C}, \mathfrak{M})$ to its own output (A's final output bit $\widetilde{b}$ is ignored). Observe that

$$\left| \Pr_{\mathbf{G}_2}[\widetilde{b} = b] - \Pr_{\mathbf{G}_1}[\widetilde{b} = b] \right| \leq \mathsf{AdvForge}_{\mathsf{MAC}_{\mathsf{M}_i}, A_2}$$

($\mathbf{G}_2$ behaves differently from $\mathbf{G}_1$ only if a forgery has been produced).

A_3 attacks $\mathsf{STREAM}_{\mathsf{M}_i}$ as follows (cf. section 3.5). First, it generates $b \in \{0,1\}$ and a string K_{oracle} of length $\mathsf{KEM}_{\mathsf{M}_i}.\mathsf{OutLen}$ uniformly at random. Then, it runs the adversary A, playing the roles of the key generation oracle, decryption oracle, and interactive encryption oracle, substituting the pregenerated string K_{oracle} and the pregenerated bit b in step 3 of section 3.2. When A outputs its bit $\widetilde{b}$, A_3 outputs 1 if $\widetilde{b} = b$ and 0 otherwise. Observe that

$$\left| \Pr_{\mathbf{G}_3}[\widetilde{b} = b] - \Pr_{\mathbf{G}_2}[\widetilde{b} = b] \right| = \mathsf{Adv}_{\mathsf{STREAM}_{\mathsf{M}_i}, A_3}$$

($\mathbf{G}_2$ corresponds to $b_{\mathsf{STREAM}} = 0$, $\mathbf{G}_3$ corresponds to $b_{\mathsf{STREAM}} = 1$ in section 3.5).

Finally observe that $\Pr_{\mathbf{G}_3}[\widetilde{b} = b] = \frac{1}{2}$.

From this we obtain the inequality

$$\mathsf{AdvCCA}_{\mathsf{M}_i, A} = 2 \cdot \left| \frac{1}{2} - \Pr_{\mathbf{G}_0}[\widetilde{b} = b] \right|$$

$$= 2 \cdot \left| \sum_{1 \leq x \leq 3} \left(\Pr_{\mathbf{G}_x}[\widetilde{b} = b] - \Pr_{\mathbf{G}_{x-1}}[\widetilde{b} = b] \right) \right|$$

$$\leq 2 \cdot \left(\mathsf{AdvCCA}_{\mathsf{KEM}_{\mathsf{M}_i}, A_1} + \mathsf{AdvForge}_{\mathsf{MAC}_{\mathsf{M}_i}, A_2} + \mathsf{Adv}_{\mathsf{STREAM}_{\mathsf{M}_i}, A_3} \right),$$

which concludes the proof.