

Robust Authenticated Encryption and the Limits of Symmetric Cryptography

Christian Badertscher¹, Christian Matt¹, Ueli Maurer¹,
Phillip Rogaway², and Björn Tackmann³

¹ Department of Computer Science, ETH Zurich, Switzerland
{badi, mattc, maurer}@inf.ethz.ch

² Department of Computer Science, University of California, Davis, USA
rogaway@cs.ucdavis.edu

³ Department of Computer Science & Engineering,
University of California, San Diego, USA
btackmann@eng.ucsd.edu

Abstract. Robust authenticated encryption (RAE) is a primitive for symmetric encryption that allows to flexibly specify the ciphertext expansion, i.e., how much longer the ciphertext is compared to the plaintext. For every ciphertext expansion, RAE aims at providing the best-possible authenticity and confidentiality. To investigate whether this is actually achieved, we characterize exactly the guarantees symmetric cryptography can provide for any given ciphertext expansion. Our characterization reveals not only that RAE reaches the claimed goal, but also, contrary to prior belief, that one cannot achieve full confidentiality without ciphertext expansion. This provides new insights into the limits of symmetric cryptography.

Moreover, we provide a rigorous treatment of two previously only informally stated additional features of RAE; namely, we show how redundancy in the message space can be exploited to improve the security and we analyze the exact security loss if multiple messages are encrypted with the same nonce.

1 Introduction

Authenticity and confidentiality are arguably among the most important cryptographic objectives. Authenticated encryption is a symmetric primitive that aims to achieve both at the same time, allowing efficiency gains and reducing the risk of misuse compared to combined schemes. Several notions of authenticated encryption have emerged over a series of works [2,3,5,7,8,13], including authenticated encryption with associated data [4,12] and misuse-resistant authenticated encryption [14]. In this development, *robust authenticated encryption (RAE)*, introduced by Hoang, Krovetz, and Rogaway [6], is the latest and most ambitious notion. Robust authenticated encryption allows to specify the ciphertext expansion λ that determines how much longer ciphertexts are compared to the corresponding plaintexts. Its self-declared goal in [6] is to provide the best-possible authenticity

and confidentiality for every choice of λ . This raises the question of what best-possible authenticity and confidentiality is, and whether RAE actually achieves it. We provide a formal model that allows us to investigate this question and answer it in the affirmative. We further show how to use verifiable redundancy to improve security, and we show what security guarantees remain if values intended as nonces are reused. Both questions were addressed in [6] but not proven formally.

1.1 Robust Authenticated Encryption

An RAE scheme consists of a key distribution \mathcal{K} , a deterministic encryption algorithm \mathcal{E} , and a deterministic decryption algorithm \mathcal{D} . The encryption algorithm takes as input a key K , a nonce N , associated data A , the ciphertext expansion λ , and a message M . It outputs a ciphertext C . The decryption algorithm takes as input K , N , A , λ , and C , and returns the corresponding message M (or \perp if C is an invalid ciphertext). In [6], the security of an RAE scheme is defined via a game in which an adversary has access to two oracles and has to distinguish between two possible settings. In the first setting, the oracles correspond to the encryption and decryption algorithm of the RAE scheme, where the key is fixed in the beginning and chosen according to \mathcal{K} . In the second setting, the first oracle chooses for each N , A , λ , and message length ℓ an injective function that maps strings of length ℓ to strings of length $\ell + \lambda$. On input (N, A, λ, M) , the oracle answers by evaluating the corresponding function. The second oracle corresponds to the partially defined inverse of that function that answers \perp if the given value has no preimage. An RAE scheme is secure if these two settings are indistinguishable for efficient adversaries. While this seems to be a strong guarantee, it is not clear which security such a scheme actually provides in a specific application and whether it is best-possible.

1.2 Security Definitions and Constructive Cryptography

Since game-based security definitions only capture what an adversary can do in a specific attack-scenario, they inherently fall short of providing guarantees that hold in any possible application of the scheme. To capture what RAE schemes achieve, we formulate our results using the constructive cryptography framework by Maurer and Renner [9,10]. The central idea of this framework is that the resources available to the parties, such as communication channels or shared randomness like cryptographic keys, are made explicit. The goal of a cryptographic protocol is then to construct, from certain existing resources, another resource that can again be used by higher-level protocols or applications. For example, the goal of an authentication scheme can be formalized as constructing an authenticated channel from a shared secret key and an insecure channel. The insecure channel allows a sender, say Alice, to send messages to a receiver, say Bob, but entirely leaks the transmitted messages to the adversary and additionally allows the adversary to delete messages and inject arbitrary messages; an authenticated channel still leaks the messages but only allows the adversary to delete messages and to deliver the messages originally sent. A conventional encryption scheme

is supposed to construct a secure channel from a shared secret key and an authenticated channel, where the secure channel restricts the leakage to the length of the transmitted messages. The composition theorem of constructive cryptography guarantees that if two protocols achieve these constructions, the composed protocol constructs a secure channel from two shared secret keys and an insecure channel, i.e, the security of the overall construction follows from the security of the individual construction steps. On the other hand, authenticated encryption directly achieves the overall construction.

1.3 Our Contributions

In the vein of [1] and accounting for the associated data RAE schemes support, we formalize the goal of RAE as constructing an *augmented secure channel (ASC)* from a shared secret key and an insecure channel. An ASC takes as input from the sender a tuple (A, M) , leaks A and the length of M to the adversary, and allows the adversary to either deliver the pair (A, M) or to terminate the channel. This channel provides authenticity for both A and M , but confidentiality is only guaranteed for the message M . The value A can for example be used to authenticate non-private header information; see [1] for an application of ASC in the context of TLS 1.3.

Uniform random injection resource. Instead of directly constructing channels from a shared secret key and an insecure channel, we introduce an intermediate system *URI (uniform random injection)* that provides the sender and receiver access to the same uniform random injections and their inverses chosen as follows: For each combination of N , A , λ , and message length ℓ , an injective function that maps strings of length ℓ to strings of length $\ell + \lambda$ is chosen uniformly at random.

As we shall see, this resource can be constructed from a shared secret key using an RAE scheme in a straightforward manner. We then construct several channels from URI and an insecure channel. The advantage of this approach is that all further constructions in this paper are information-theoretic, i.e., we do not have to relate the security of each construction step to the RAE security game. Instead, we can rely on the composition theorem to guarantee the security of the overall construction.

Random injection channel. We show that one can construct a channel we call *RIC (random injection channel)* from URI and an insecure channel by fixing λ and using a counter as the nonce. RIC can be seen as a further intermediate step towards constructing ASC, that in addition allows us to analyze best-possible security.

The channel RIC takes as input a pair (A, M) from the sender and leaks A and the length of M to the adversary. The adversary can deliver the pair (A, M) , and further at any point in time try to inject a new message of length ℓ and some value A . The probability with which such an injection is successful depends

on λ and ℓ . In case of a success, an almost uniform message of length ℓ from the message space together with A is delivered to the receiver. If an injection was successful and the tuple (A, M) was received, and if the sender subsequently sends exactly the pair (A, M) , then the adversary is notified about this repetition.

Best possible authenticity and confidentiality. If ASC is considered as the ultimate goal of RAE and authenticated encryption in general, the only shortcomings of RIC are that it is possible to inject messages with positive probability and that, if an attempted message injection was successful, the channel leaks a certain repetition to the adversary. While the first shortcoming is a lack of authenticity, the second one is a lack of confidentiality. While the type of leakage violating confidentiality might seem artificial, we describe an application in which such leakage might be problematic. Briefly, the leakage can reveal hidden information flow from the receiver to the sender.

We then analyze whether RAE really achieves the best-possible authenticity and confidentiality by bounding the probabilities of successful message injections and of leaking this particular repetition pattern for arbitrary schemes for achieving authenticity and confidentiality. While it is straightforward to see that authenticity requires redundancy and therefore a large ciphertext expansion, one might hope that the repetition leakage can be avoided. We prove that this is not the case, i.e., we show that the probability of an adversary being able to observe such a repetition is at least as high as in RIC, no matter what scheme is used or which setup assumptions are made.

To illustrate this lack of confidentiality for a concrete scheme, consider the following scenario in which the one-time pad is used over an insecure channel: Assume the attacker injects a ciphertext to Bob who decrypts it using the shared secret key and outputs the resulting message. Further assume Alice afterwards sends a message to Bob which results in the attacker seeing the same ciphertext. In that case, the attacker learns the fact that the message sent by Alice equals the message output by Bob. This contradicts the understanding of confidentiality as revealing nothing except the length of the transmitted message.⁴ Our results generalize this observation to arbitrary schemes. We thereby refine the understanding of what symmetric cryptography can and cannot achieve by showing that confidentiality, quite surprisingly, also requires redundancy in the ciphertexts when only insecure channels and an arbitrary setup are assumed, even if the protocol can keep state.

Augmented secure channels and message redundancy. Since the probability of successful message injections decreases exponentially with λ , RIC is indistinguishable from ASC for large λ . We further provide a construction that

⁴ This also contradicts a prior result in [11] that claims that the one-time pad constructs a certain (fully) confidential channel, a so-called XOR-malleable channel, from an insecure channel and a shared key. The proof in that paper is flawed in that the simulation fails if more ciphertexts are injected than messages sent.

incorporates an idea from [6] to exploit the redundancy in messages to achieve a better bound. Our construction reveals the exact trade-off between ciphertext expansion and redundancy to achieve a required security level.

Nonce-reuse resistance. It was claimed in [6] that reusing nonces only results in leaking the repetition pattern of messages, but does not compromise security beyond that. However, the claim was neither formalized nor proven. We fill this gap by introducing the channel resource *RASC (Repetition ASC)* that, aside of the length of each message, leaks the repetition pattern of the transmitted messages to the adversary. Furthermore, the adversary can deliver messages out-of-order and arbitrarily replay messages. We show that RASC can be constructed from URI and an insecure channel if the used nonce is always the same. This confirms the informal claim from [6] and makes explicit that some authenticity is lost by allowing the adversary to reorder messages.

2 Preliminaries

2.1 Notation for Systems and Algorithms

We describe our systems with pseudocode using the following conventions: We write $x \leftarrow y$ for assigning the value y to the variable x . For a distribution \mathcal{X} over some set, $x \leftarrow \mathcal{X}$ denotes sampling x according to \mathcal{X} . For a finite set X , $x \leftarrow X$ denotes assigning to x a uniformly random value in X .

We denote the empty list by $[]$ and for a list L , $L \parallel x$ denotes the list L with x appended. Furthermore, $|L|$ denotes the number of elements in L and the i th element in L is denoted by $L[i]$ for $i \in \{1, \dots, |L|\}$. For a FIFO queue \mathcal{Q} , we write $\mathcal{Q}.\text{enqueue}(x)$ to insert x into the queue and $\mathcal{Q}.\text{dequeue}()$ to retrieve (and remove) the element of the queue that was inserted first among all remaining elements.

For $n, m \in \mathbb{N}$, $\text{Inj}(\Sigma^n, \Sigma^m)$ denotes the set of injective functions $\Sigma^n \rightarrow \Sigma^m$. For an injective function $f: X \rightarrow Y$, we denote by f^{-1} the function $Y \rightarrow X \cup \{\perp\}$ that maps y to the preimage of y under f if existing, and to the distinct element \perp otherwise.

Typically queries to systems consist of a suggestive keyword and a list of arguments (e.g., (send, M) to send the message M). We ignore keywords in writing the domains of arguments, e.g., $(\text{send}, M) \in \mathcal{M}$ indicates that $M \in \mathcal{M}$.

2.2 Constructive Cryptography

Constructive cryptography makes statements about *constructions* of *resources* from other resources. A resource is a system with interfaces via which the resource interacts with its environment and which can be thought of as being assigned to parties. All resources in this paper have an interface **A** for the sender (Alice), an interface **B** for the receiver (Bob), and an interface **E** for the adversary (Eve). In

our security statements, we are interested in the advantage of a distinguisher \mathbf{D} in distinguishing two resources, say \mathbf{R} and \mathbf{S} which is defined as

$$\Delta^{\mathbf{D}}(\mathbf{R}, \mathbf{S}) = \Pr[\mathbf{DR} = 1] - \Pr[\mathbf{DS} = 1],$$

where $\Pr[\mathbf{DR} = 1]$ denotes the probability that \mathbf{D} outputs 1 when connected to resource \mathbf{R} . More concretely, \mathbf{DR} is a random experiment, where the distinguisher repeatedly provides an input to one of the interfaces \mathbf{A} , \mathbf{B} , or \mathbf{E} and observes the output generated in reaction to that input before it decides on its output bit.

Converters are systems that can be attached to an interface of a resource to change the inputs and outputs at that interface, which yields another resource. A converter is a system with two interfaces: the *inner interface* \mathbf{in} is connected to an interface of a resource and the *outer interface* \mathbf{out} , becomes the new connection point of that resource towards the environment. The protocols of the honest parties and simulators correspond to converters.

We directly state the central definition of a construction of [9] and briefly explain the relevant conditions.

Definition 1. *Let \mathbf{R} and \mathbf{S} be resources and let $\mathbf{noAtck}_{\mathbf{R}}$ and $\mathbf{noAtck}_{\mathbf{S}}$ be converters that describe the default behavior at interface \mathbf{E} when no attacker is present. Let ε be a function that maps distinguishers to a value in $[-1, 1]$ and let \mathbf{sim} be a converter (the simulator). A protocol, i.e., a pair $(\mathbf{conv}_1, \mathbf{conv}_2)$ of converters, constructs resource \mathbf{S} from resource \mathbf{R} within ε and with respect to the pair $(\mathbf{noAtck}_{\mathbf{R}}, \mathbf{noAtck}_{\mathbf{S}})$ and the simulator \mathbf{sim} , if for all distinguishers \mathbf{D} ,*

$$\Delta^{\mathbf{D}}(\mathbf{conv}_1^{\mathbf{A}} \mathbf{conv}_2^{\mathbf{B}} \mathbf{noAtck}_{\mathbf{R}}^{\mathbf{E}} \mathbf{R}, \mathbf{noAtck}_{\mathbf{S}}^{\mathbf{E}} \mathbf{S}) \leq \varepsilon(\mathbf{D}) \quad (\text{Availability})$$

$$\Delta^{\mathbf{D}}(\mathbf{conv}_1^{\mathbf{A}} \mathbf{conv}_2^{\mathbf{B}} \mathbf{R}, \mathbf{sim}^{\mathbf{E}} \mathbf{S}) \leq \varepsilon(\mathbf{D}). \quad (\text{Security})$$

The first condition ensures that the protocol implements the required functionality if there is no attacker. For example, for communication channels, all sent messages have to be delivered when no attacker interferes with the protocol.

The second condition ensures that whatever Eve can do with the assumed resource, she could do as well with the constructed resource by using the simulator \mathbf{sim} . Turned around, if the constructed resource is secure by definition, there is no successful attack on the protocol.

The notion of construction is composable, which intuitively means that the constructed resource can be replaced in any context by the assumed resource with the protocol attached without affecting the security. This is proven in [9].

2.3 Robust Authenticated Encryption

Let Σ be an alphabet (a finite nonempty set). Typically an element of Σ is a bit ($\Sigma = \{0, 1\}$) or a byte ($\Sigma = \{0, 1\}^8$). For a string $x \in \Sigma^*$, $|x|$ denotes its length. We define the syntax of a robust authenticated encryption scheme following [6].

Definition 2. A robust authenticated encryption (RAE) scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ consists of a key distribution \mathcal{K} , a deterministic encryption algorithm \mathcal{E} that maps a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \mathcal{A}$, ciphertext expansion $\lambda \in \mathbb{N}$, and a message $M \in \mathcal{M}$ to a ciphertext $C \in \mathcal{C}$, and a deterministic decryption algorithm \mathcal{D} that maps (K, N, A, λ, C) to an element in $\mathcal{M} \cup \{\perp\}$. We assume the domains \mathcal{N} , \mathcal{A} , \mathcal{M} , and \mathcal{C} are equal to Σ^* . We write $\mathcal{E}_K^{N,A,\lambda}$ and $\mathcal{D}_K^{N,A,\lambda}$ for the functions $\mathcal{E}(K, N, A, \lambda, \cdot)$ and $\mathcal{D}(K, N, A, \lambda, \cdot)$, respectively. We require that $\mathcal{D}_K^{N,A,\lambda}(\mathcal{E}_K^{N,A,\lambda}(M)) = M$ for all K, N, A, λ, M .

3 Shared Uniform Random Injections and RAE Security

In this section, we describe the resource **URI** that grants access to shared uniform random injections and their inverses at interfaces **A** and **B**, and no access at interface **E**. We then use **URI** to define the security of RAE schemes and show that any RAE scheme that satisfies this definition can be used to construct **URI** from a shared secret key. Though syntactically different, it is easy to see that our definition is equivalent to the security definition from [6]. We recall that definition and prove the equivalence in the full version of this paper.

We first give a definition for the uniform random injection system **URI**.

Definition 3. The resource **URI** has interfaces **A**, **B**, and **E** and takes inputs of the form $(\text{fun}, N, A, \lambda, x)$ and $(\text{inv}, N, A, \lambda, y)$ at interfaces **A** and **B** for $N \in \mathcal{N}$, $A \in \mathcal{A}$, $\lambda \in \mathbb{N}$, $x \in \mathcal{M}$, and $y \in \mathcal{C}$. Any input at interface **E** is ignored. We assume the domains \mathcal{N} , \mathcal{A} , \mathcal{M} , and \mathcal{C} are equal to Σ^* . On input $(\text{fun}, N, A, \lambda, x)$ at interface **A** or **B**, it returns $f_{N,A,\lambda,|x|}(x)$ at the same interface. On input $(\text{inv}, N, A, \lambda, y)$, it returns $f_{N,A,\lambda,|y|-\lambda}^{-1}(y)$ if $|y| > \lambda$, and \perp otherwise. The function $f_{N,A,\lambda,\ell}$ is chosen uniformly at random from the set $\text{Inj}(\Sigma^\ell, \Sigma^{\ell+\lambda})$ when needed for the first time and reused for later inputs.

3.1 Definition of RAE Security and Construction of URI

We define a shared key resource $\mathbf{SK}_{\mathcal{K}}$ for some key distribution \mathcal{K} . The resource initially chooses a key according to \mathcal{K} and outputs this key to interfaces **A** and **B** while interface **E** remains inactive, see Fig. 1. Slightly abusing notation, we will also refer to the key space by \mathcal{K} whenever no confusion can arise. We further define the converter rae_{Π} that is based on an RAE scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. First, rae_{Π} requests the key from $\mathbf{SK}_{\mathcal{K}}$. For any input at the outer interface, it evaluates \mathcal{E} or \mathcal{D} using that key (and the arguments provided in the input) and returns the result. The code is given in Fig. 1.

We consider an RAE scheme secure if all efficient distinguishers have poor advantage with respect to the following definition.

Definition 4. The advantage of a distinguisher \mathbf{D} for an RAE scheme Π is quantified as

$$\text{Adv}_{\Pi}^{\text{rae}}(\mathbf{D}) := \Delta^{\mathbf{D}}(\text{rae}_{\Pi}^{\mathbf{A}} \text{rae}_{\Pi}^{\mathbf{B}} \mathbf{SK}_{\mathcal{K}}, \mathbf{URI}).$$

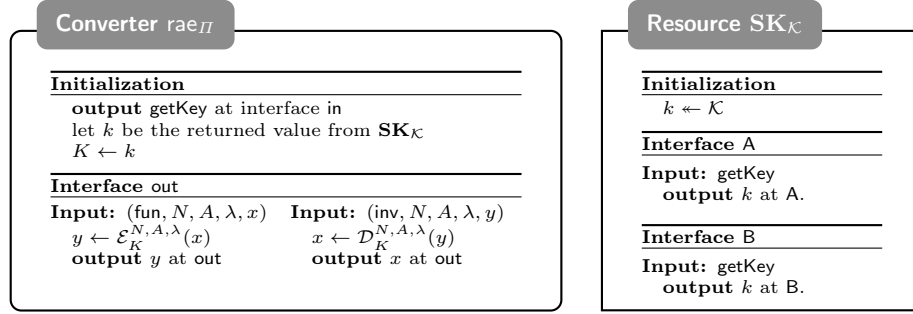


Fig. 1. Protocol that constructs **URI** from a shared secret key (left) and the shared secret key resource (right). For the shared key resource, interface **E** remains inactive.

It is straightforward to see that the definition implies the following construction statement, where the converters **sim** and **noAtck** are defined as the converter that blocks any interaction at the interface it is connected to.

Lemma 1. *The protocol $(\text{rae}_\Pi, \text{rae}_\Pi)$ constructs **URI** from $\text{SK}_\mathcal{K}$ within $\text{Adv}_\Pi^{\text{rae}}$ with respect to $(\text{noAtck}, \text{noAtck})$ and simulator **sim** defined above.*

Proof. Since interface **E** of $\text{SK}_\mathcal{K}$ and **URI** are inactive, the converters **sim** and **noAtck** have no effect when connected to that interface, i.e., $\text{noAtck}^{\text{E}} \text{SK}_\mathcal{K} = \text{SK}_\mathcal{K}$ and $\text{noAtck}^{\text{E}} \text{URI} = \text{sim}^{\text{E}} \text{URI} = \text{URI}$. Thus, both the availability and the security condition of the construction are equivalent to

$$\Delta^{\text{D}}(\text{rae}_\Pi^{\text{A}} \text{rae}_\Pi^{\text{B}} \text{SK}_\mathcal{K}, \text{URI}) \leq \text{Adv}_\Pi^{\text{rae}}(\text{D})$$

for all distinguishers **D**, which trivially holds by definition of $\text{Adv}_\Pi^{\text{rae}}$. \square

4 Random Injection Channels: Security for any Expansion

The goal of the current section is to examine the exact security achieved by RAE schemes when used to protect communication. We present constructions of specific secure channels from insecure channels and resource **URI** where each type of secure channel precisely captures the amount of leakage to an eavesdropper and the possible influence of an adversary interfering with the protocol execution. As an additional result, we are able to answer what best-possible communication security is and observe that RAE schemes achieve this level of security.

The insecure channel **IC** allows messages $m \in \mathcal{M}$ to be input repeatedly at interface **A**. Each message is subsequently leaked at the **E**-interface. At interface **E**, arbitrary messages (including those that were previously input at interface **A**) can be injected such that they are delivered to **B**. This channel does not give any security guarantees to Alice and Bob. A formal description is provided in Fig. 2. For the rest of this paper, the message space of the insecure channel is Σ^* .




Fig. 2. The insecure channel resource.

4.1 Constructing Random Injection Channels

The constructed channel. The channel we construct in this section is defined in Fig. 3 and can be roughly described as follows: It allows to repeatedly send pairs (A_i, M_i) in an ordered fashion from a sender to a receiver. Each pair consists of the associated data A_i and the message M_i . The attacker is limited to seeing the associated data A_i and the length of the message $|M_i|$ of each transmitted pair. Additionally, the attacker learns whether the i th injected pair equals the one that is currently sent.

The attacker can either deliver the next legitimate pair (A_i, M_i) or try to inject a pair (A, M) that is different from (A_i, M_i) . Such an injection is only successful with a certain probability. The associated data A and the length ℓ of the message are chosen by the attacker and M is a uniformly random message of length ℓ if $A \neq A_i$. Otherwise, M is a uniformly random message $M \neq M_i$ of length ℓ . If an injection attempt is not successful, the resource does not deliver messages at interface B any more and signals an error by outputting \perp . If the adversary injects the i th message, the legitimate i th message cannot be delivered anymore.⁵

The success probability of an injection attempt depends on the expansion λ and the specified message length ℓ and whether the sender's queue \mathcal{S} is empty or not. The exact probabilities are quantified by the two sampling functions SAMPLE and SAMPLEEXCL. The function SAMPLE first samples a bit according to the probability that a fixed element from $\Sigma^{\ell+\lambda}$ has a preimage under a uniform random injection $\Sigma^\ell \rightarrow \Sigma^{\ell+\lambda}$. If the bit is 1, a uniform random preimage is returned. The function SAMPLEEXCL essentially does the same, but the domain and codomain are both reduced by one element.⁶

Protocol. We construct resource \mathbf{RIC}_λ from $[\mathbf{URI}, \mathbf{IC}]$ which denotes the resource that provides at each interface access to the corresponding interface of

⁵ This relates to the security of RAE schemes which ensures that the message cannot be decrypted using a wrong nonce. In our construction, the nonce is implemented as the sequence number.

⁶ This ensures that the injected message is different from the one that the sender provided.




Fig. 3. Description of RIC_λ . In the description, $\text{Bernoulli}(p)$ denotes the distribution over $\{0, 1\}$, where 1 has probability p and 0 has probability $1 - p$.




Fig. 4. The converters for the sender (left) and the receiver (right) to construct RIC_λ .

both resources. Our protocol specifies a particular but very natural usage of **URI** where the nonce is implemented as a counter value.⁷ We present the protocol as pseudocode in Fig. 4. The converter for the sender, snd_λ , accepts inputs of the form (send, A, M) at its outer interface. It outputs $(\text{fun}, i, A, \lambda, M)$ at the inner interface to **URI**. The nonce is implemented as a counter and λ is the parameter of the protocol. Once a ciphertext is received as a return value from **URI**, it is output together with its associated data at the inner interface for the insecure channel **IC**. The receiver converter rcv_λ receives ciphertexts together with the associated data at its inner interface from **IC** and decrypts C using parameters A, i and λ . On success, the corresponding plaintext is output at the outer interface. If decryption fails, the converter stops and signals an error by outputting \perp .

Construction statement. In order to show that the protocol $(\text{snd}_\lambda, \text{rcv}_\lambda)$ constructs \mathbf{RIC}_λ from $[\mathbf{URI}, \mathbf{IC}]$, we prove both conditions of Definition 1. For all channels, the converter noAtck corresponds to the converter dlv that on any input at its inner interface, outputs deliver to the channel connected to its inner interface and blocks any interaction at its outer interface.

Theorem 1. *Let $\lambda \in \mathbb{N}$. The protocol $(\text{snd}_\lambda, \text{rcv}_\lambda)$ constructs resource \mathbf{RIC}_λ from $[\mathbf{URI}, \mathbf{IC}]$ with respect to (dlv, dlv) and simulator $\text{sim}_{\mathbf{RIC}}$ as defined in Fig. 5. More specifically, for all distinguishers \mathbf{D}*

$$\Delta^{\mathbf{D}} \left(\text{snd}_\lambda^{\mathbf{A}} \text{rcv}_\lambda^{\mathbf{B}} \text{dlv}^{\mathbf{E}}[\mathbf{URI}, \mathbf{IC}], \text{dlv}^{\mathbf{E}} \mathbf{RIC}_\lambda \right) = 0 \quad (1)$$

$$\text{and} \quad \Delta^{\mathbf{D}} \left(\text{snd}_\lambda^{\mathbf{A}} \text{rcv}_\lambda^{\mathbf{B}}[\mathbf{URI}, \mathbf{IC}], \text{sim}_{\mathbf{RIC}}^{\mathbf{E}} \mathbf{RIC}_\lambda \right) = 0. \quad (2)$$

We prove Theorem 1 in the full version of this paper.

4.2 What is Best-Possible Security?

We observe that \mathbf{RIC}_λ has two undesirable properties: messages can be injected and the output at interface **E** leaks more than only the length of the payload in that it reveals whether Alice sends the pair (A, M) that has been output by Bob upon an adversarial injection. In contrast, a channel that only leaks the message length is considered fully confidential.

We first illustrate an application in which this lack of full confidentiality is problematic. The main purpose of our example is to show that one cannot exclude the existence of an application where exactly this (intuitively small) difference to full confidentiality yields a security problem.

Second, we show that a successful injection followed by the undesired information leakage about the repetition is possible for any scheme, even if it is stateful and uses an arbitrary setup before starting communication, and that the probability of this is minimized if \mathbf{RIC}_λ is used.

⁷ Implementing the nonce as a counter allows to maintain the order of messages.




Fig. 5. Simulator for the security condition of the construction of RIC_λ .

Sample scenario: On the difference to full confidentiality. Assume a setting in which party A is allowed to send information to party B via a fully confidential channel but not vice versa. Suppose now that B finds a possibility to send information to A via a covert channel and the two parties use the confidential channel for messages from A to B and the covert channel for messages from B to A. Suppose now that the confidential channel is in fact a channel that leaks the above repetition event instead of only the message length. This gives an investigator E a means to test for the existence of a covert channel from B to A as follows: At some point, E injects a random message M to B. Assuming information flow from B to A, party B might start a discussion about M with party A. As part of this conversation, A might send M to B, which would signal a repetition-event to E. For large message spaces, it is very unlikely that A comes up with the exact same message that was randomly injected to B before, unless there is a (hidden) flow of information. The occurrence of the event is therefore a witness for the existence of a channel from B to A. In contrast, a fully confidential channel would not reveal the existence of the covert channel.

RIC provides best-possible security. In RIC_λ , an injection attempt is successful with probability at most $\frac{|\mathcal{M}|}{|\mathcal{C}|} = |\Sigma|^{-\lambda}$ and given a successful injection and that Alice subsequently sends the corresponding output of Bob, the above described leakage occurs with probability 1. Overall, the total probability that an undesired property can be observed is bounded by $|\Sigma|^{-\lambda}$.

We show that this probability is optimal and that no protocol can achieve a better bound. Hence, \mathbf{RIC}_λ maximizes authenticity and confidentiality. We first prove the following general lemma.

Lemma 2. *Let \mathcal{M} and \mathcal{C} be finite nonempty sets and let E and D be random variables over functions $\mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$ and $\mathcal{A} \times \mathcal{C} \rightarrow \mathcal{M} \cup \{\perp\}$, respectively, such that*

$$\forall m \in \mathcal{M}, a \in \mathcal{A}: \Pr[D(a, E(a, m)) = m] \geq p$$

for some $p \in [0, 1]$. We then have for all $a \in \mathcal{A}$ and any random variable C that is distributed uniformly over \mathcal{C} and independent from E and D ,

$$\Pr[D(a, C) \neq \perp \wedge E(a, D(a, C)) = C] \geq p \cdot \frac{|\mathcal{M}|}{|\mathcal{C}|}.$$

Proof. We have for all $a \in \mathcal{A}$

$$\begin{aligned} & \Pr[D(a, C) \neq \perp \wedge E(a, D(a, C)) = C] \\ &= \sum_{m \in \mathcal{M}} \sum_{c \in \mathcal{C}} \Pr[D(a, c) = m \wedge E(a, m) = c \wedge C = c] \\ &= \frac{1}{|\mathcal{C}|} \sum_{m \in \mathcal{M}} \underbrace{\sum_{c \in \mathcal{C}} \Pr[D(a, c) = m \wedge E(a, m) = c]}_{= \Pr[D(a, E(a, m)) = m] \geq p} \\ &\geq p \cdot \frac{|\mathcal{M}|}{|\mathcal{C}|}, \end{aligned}$$

where we used in the second step that C is distributed uniformly over \mathcal{C} and independent from E and D . \square

Lemma 2 can be applied to our usual setting $\mathbf{enc}_\lambda^A \mathbf{dec}_\lambda^B[\mathbf{SK}_K, \mathbf{IC}]$ for a generic protocol $(\mathbf{enc}_\lambda, \mathbf{dec}_\lambda)$ in a straightforward manner: we only have to observe that for the i th input $(\mathbf{send}, A_i, M_i)$, for all $i \in \mathbb{N}$, converter \mathbf{enc}_λ is characterized by a probabilistic map $\mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$, that may depend on previous inputs and outputs and on the key k . Similarly, the converter \mathbf{dec}_λ is characterized by a probabilistic map $\mathcal{A} \times \mathcal{C} \rightarrow \mathcal{M} \cup \{\perp\}$ for any $i \in \mathbb{N}$.

Correctness of the protocol implies that if $(\mathbf{send}, A_i, M_i)$ is input to \mathbf{enc}_λ as the i th query and yields ciphertext C_i , then the probability that on the i th input (A_i, C_i) to \mathbf{dec}_λ , and if \mathbf{dec}_λ has not halted yet, the converter decrypts the ciphertext to M_i with probability p ; note that $p = 1$ for RAE schemes. Hence, Lemma 2 implies that the probability that any of the two undesirable properties can be observed during protocol execution is at least $|\Sigma|^{-\lambda}$.

5 Augmented Secure Channels and Verifiable Redundancy

Looking at the specification of \mathbf{RIC}_λ , we observe that for large values λ , the probability of successful injections becomes exponentially small, and so are the

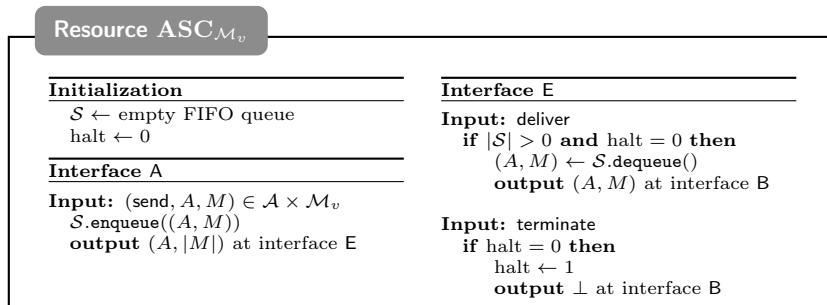


Fig. 6. Description of **ASC**, an augmented secure channel.

repetition events at interface E. We are particularly interested in the resource that specifies this abstraction: a channel that allows to securely transmit messages consisting of an associated data and a payload part such that an attacker is limited to seeing the associated data and the length of the payload, to deliver the next message, or to abort the whole communication. This channel abstraction corresponds to an augmented secure channel. Such channels were introduced in [1] and shown to be achievable by the AEAD notion of [12]. Not surprisingly, this confirms that RAE and AEAD achieve the same security goals for large ciphertext expansion.

Additionally, we formally show how redundancy in messages can be exploited to improve authenticity, where redundancy restricts the set of valid messages to a subset of $\mathcal{M} = \Sigma^*$.

The following theorem provides the exact security bound in terms of redundancy in the message space and ciphertext expansion λ . We thereby confirm a conjecture of [6]. Let $v : \mathcal{M} \mapsto \{\text{true}, \text{false}\}$ be a predicate on the message space. We define the subset $\mathcal{M}_v := \{M \mid M \in \mathcal{M} \wedge v(M)\}$ which we call the set of valid messages. Following [6], the density of \mathcal{M}_v is defined as

$$d_v := \max_{\ell \in \mathbb{N}} \frac{|\Sigma^\ell \cap \mathcal{M}_v|}{|\Sigma^\ell|}.$$

A similar statement, namely for a specific encryption mode that allows the attacker to either forward messages or to inject purely random ones into the channel, and for a specific redundant encoding, has been shown by Ruedlinger [15].

The constructed channel. The augmented secure channel $ASC_{\mathcal{M}_v}$ is described in Fig. 6. The channel is derived from RIC_λ by requiring that $M \in \mathcal{M}_v$ and by removing undesired capabilities that vanish due to the exponentially small success probability for large λ .

Protocol. The protocol for the sender, $sndChk_v$, accepts inputs of the form $(send, A, M)$ at its outer interface and forwards the pair to the channel RIC_λ if and only if $v(M)$ (and otherwise ignores the request). The receiver converter

rcvChk_v , on receiving the pair (A, M) from \mathbf{RIC}_λ , outputs (A, M) at its outer interface if and only if $v(M)$. If rcvChk_v receives \perp from \mathbf{RIC}_λ or if $\neg v(M)$, it outputs \perp at its outer interface and halts.

Theorem 2. *Let $\lambda \in \mathbb{N}$. The protocol $(\text{sndChk}_v, \text{rcvChk}_v)$ constructs $\mathbf{ASC}_{\mathcal{M}_v}$ from \mathbf{RIC}_λ with respect to (dlv, dlv) and simulator $\text{sim}_{\mathbf{ASC}}$ defined in Fig. 7. More specifically, for all distinguishers \mathbf{D}*

$$\Delta^{\mathbf{D}} \left(\text{sndChk}_v^{\mathbf{A}} \text{rcvChk}_v^{\mathbf{B}} \text{dlv}^{\mathbf{E}} \mathbf{RIC}_\lambda, \text{dlv}^{\mathbf{E}} \mathbf{ASC}_{\mathcal{M}_v} \right) = 0 \quad (3)$$

$$\text{and } \Delta^{\mathbf{D}} \left(\text{sndChk}_v^{\mathbf{A}} \text{rcvChk}_v^{\mathbf{B}} \mathbf{RIC}_\lambda, \text{sim}_{\mathbf{ASC}}^{\mathbf{E}} \mathbf{ASC}_{\mathcal{M}_v} \right) \leq d_v \cdot |\Sigma|^{-\lambda}. \quad (4)$$




Fig. 7. Simulator for the security condition of the construction of $\mathbf{ASC}_{\mathcal{M}_v}$.

Proof. The availability condition (3) is straightforward to verify. We only prove the security condition (4). It is easy to see that the two systems behave identically as long as no injection attempt is successful. This is because successful injections are necessary for observing repeat: as long as no injection is successful, for any send-query to \mathbf{RIC}_λ , the condition $i \leq |\mathcal{R}|$ is not satisfied after incrementing i . We thus only have to bound this probability. We hence consider the event that in an interaction of a distinguisher with the real system $\text{sndChk}_v^{\mathbf{A}} \text{rcvChk}_v^{\mathbf{B}} \mathbf{RIC}_\lambda$ the first attempt to inject a random message is successful (since in case of an unsuccessful attempt, both channels stop delivering messages). In any interaction of \mathbf{D} with the resource, the probability of the event is determined by \mathbf{RIC}_λ as one out of two possibilities, see Fig. 3. For any $i \in \mathbb{N}$, if the i th query at interface \mathbf{E} is the first attempt to inject a message, then the probability depends on whether the specified associated data and the length coincides with the length of the message and the associated data of the i th input at interface \mathbf{A} . Both probabilities are upper bounded by

$$\begin{aligned} \max \left\{ \frac{|\Sigma|^{C_i - \lambda} \cap \mathcal{M}_v| - 1}{|\Sigma|^{C_i}| - 1}, \frac{|\Sigma|^{C_i - \lambda} \cap \mathcal{M}_v|}{|\Sigma|^{C_i}|} \right\} &\leq \frac{|\Sigma|^{C_i - \lambda} \cap \mathcal{M}_v|}{|\Sigma|^{C_i}|} \\ &= \frac{|\Sigma|^{C_i - \lambda} \cap \mathcal{M}_v|}{|\Sigma|^{C_i - \lambda}|} \cdot |\Sigma|^{-\lambda} \leq d_v \cdot |\Sigma|^{-\lambda}, \end{aligned}$$

where we used $\frac{x-1}{y-1} \leq \frac{x}{y}$ for $x \leq y$ in the first step, and the definition of d_v in the last step. \square

6 Guarantees for Nonce-Reuse

One goal of robust authenticated encryption is to provide resilience to the misuse when nonces are repeated. While the expected security loss was only informally stated in [6], we rigorously derive the exact guarantees that can still be expected in such a scenario. To this end, we consider the extreme case where the nonce is a constant value.

Repetition ASC. The channel that is achieved if the nonce is repeating is denoted $\mathbf{RASC}_{\mathcal{M}_v}$ and its description is given in Fig. 8. There are two differences to $\mathbf{ASC}_{\mathcal{M}_v}$: First, not only the length of the message is leaked at interface \mathbf{E} but also the number i of the first transmitted message that equals the current message. This leaks the repetition pattern of transmitted values. Second, the adversary can replay messages and induce arbitrary out-of-order delivery.

Protocol. The protocol, which we denote by $(\mathbf{rsnd}, \mathbf{rcv})$, invokes \mathbf{URI} using the constant nonce 0. Furthermore, the protocol verifies that all messages are from the set \mathcal{M}_v . The protocol is specified in Fig. 9.

Theorem 3. *Let $\lambda \in \mathbb{N}$. The protocol $(\mathbf{rsnd}_\lambda, \mathbf{rcv}_\lambda)$ constructs $\mathbf{RASC}_{\mathcal{M}_v}$ from $[\mathbf{URI}, \mathbf{IC}]$ with respect to $(\mathbf{dlv}, \mathbf{dlv})$ and simulator $\mathbf{sim}_{\mathbf{RASC}}$ defined in Fig. 10. More specifically, we have for all distinguishers \mathbf{D}*

$$\Delta^{\mathbf{D}} \left(\mathbf{rsnd}_\lambda^{\mathbf{A}} \mathbf{rcv}_\lambda^{\mathbf{B}} \mathbf{dlv}^{\mathbf{E}}[\mathbf{URI}, \mathbf{IC}], \mathbf{dlv}^{\mathbf{E}} \mathbf{RASC}_{\mathcal{M}_v} \right) \leq \frac{2d_v + q \cdot (q-1)}{2} \cdot |\Sigma|^{-\lambda} \quad (5)$$

and

$$\Delta^{\mathbf{D}} \left(\mathbf{rsnd}_\lambda^{\mathbf{A}} \mathbf{rcv}_\lambda^{\mathbf{B}}[\mathbf{URI}, \mathbf{IC}], \mathbf{sim}_{\mathbf{RASC}}^{\mathbf{E}} \mathbf{RASC}_{\mathcal{M}_v} \right) \leq \frac{2d_v + q \cdot (q-1)}{2} \cdot |\Sigma|^{-\lambda}, \quad (6)$$

where q is the total number of inputs made by \mathbf{D} .

We prove Theorem 3 in the full version of this paper.

Acknowledgments. Ueli Maurer was supported by the Swiss National Science Foundation (SNF), project no. 200020-132794. Björn Tackmann was supported by the Swiss National Science Foundation (SNF) via Fellowship no. P2EZP2_155566 and the NSF grants CNS-1228890 and CNS-1116800. Much of the work on this paper was done while Phil Rogaway was visiting Ueli Maurer’s group at ETH Zurich. Many thanks to Ueli for hosting that sabbatical. Rogaway was also supported by NSF grants CNS-1228828 and CNS-1314885.

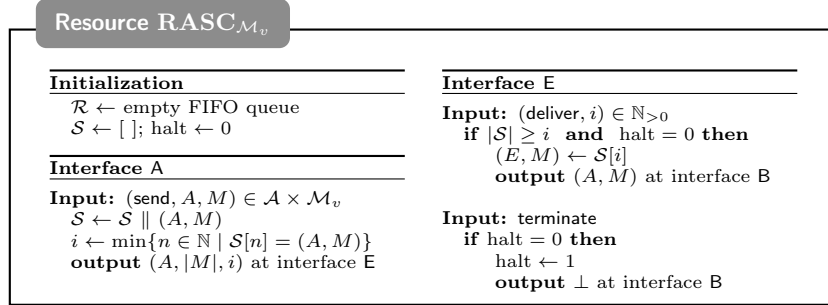


Fig. 8. Description of RASC.

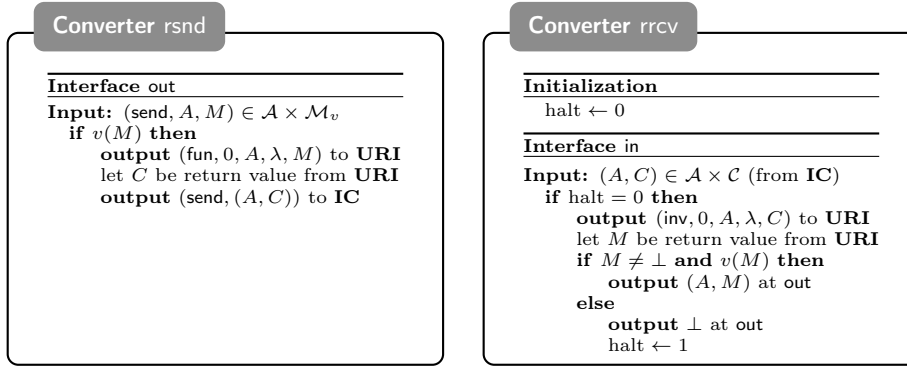
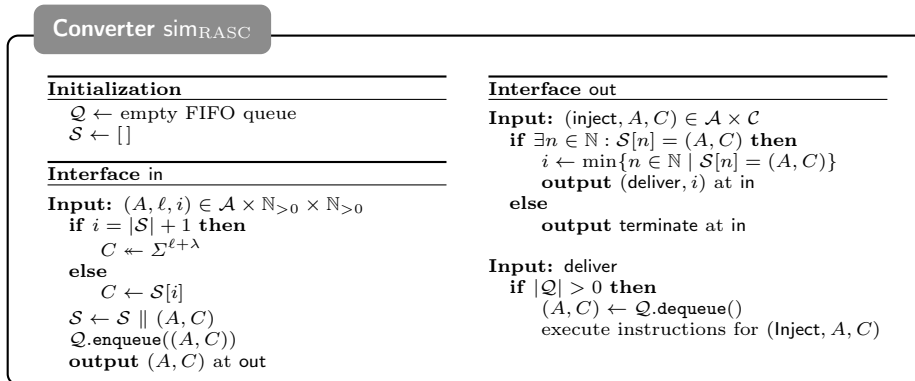


Fig. 9. The converters for the sender (left) and the receiver (right).


 Fig. 10. Simulator for the security condition of the construction of $\text{RASC}_{\mathcal{M}_v}$.

References

1. Badertscher, C., Matt, C., Maurer, U., Rogaway, P., Tackmann, B.: Augmented Secure Channels as the Goal of the TLS 1.3 Record Layer. *Cryptology ePrint Archive*, Report 2015/394, 2015
2. Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In: *Advances in Cryptology – ASIACRYPT 2000*, pp. 531–545. Springer Berlin Heidelberg (2000)
3. Bellare, M., Rogaway, P.: Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography. In: *Advances in Cryptology – ASIACRYPT 2000*, pp. 317–330. Springer (2000)
4. Bellare, M., Rogaway, P., Wagner, D.: The EAX mode of operation. In: *Fast Software Encryption*, pp. 389–407. Springer (2004)
5. Gligor, V., Donescu, P.: Fast Encryption and Authentication: XCBC Encryption and XECB Authentication Modes. In: *Fast Software Encryption*, pp. 92–108. Springer (2002)
6. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption: AEZ and the problem that it solves. In: *Advances in Cryptology – EUROCRYPT 2015*, pp. 15–44. Springer (2015)
7. Jutla, C.: Encryption modes with almost free message integrity. In: *Advances in Cryptology – EUROCRYPT 2001*, pp. 529–544. Springer (2001)
8. Katz, J., Yung, M.: Unforgeable encryption and chosen ciphertext secure modes of operation. In: *Fast Software Encryption*, pp. 284–299. Springer (2001)
9. Maurer, U.: Constructive cryptography – a new paradigm for security definitions and proofs. In: Mödersheim, S., Palamidessi, C. (eds.) *Theory of Security and Applications*, Lecture Notes in Computer Science, vol. 6993, pp. 33–56. Springer (2012)
10. Maurer, U., Renner, R.: Abstract cryptography. In: Chazelle, B. (ed.) *The Second Symposium on Innovations in Computer Science*, ICS 2011, pp. 1–21. Tsinghua University Press (2011)
11. Maurer, U., Rüdinger, A., Tackmann, B.: Confidentiality and integrity: a constructive perspective. In: *TCC 2012*, pp. 209–229. Springer (2012)
12. Rogaway, P.: Authenticated-encryption with associated-data. In: *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pp. 98–107. ACM (2002)
13. Rogaway, P., Bellare, M., Black, J.: OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Transactions on Information and System Security (TISSEC)* 6(3), 365–403 (2003)
14. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: *Advances in Cryptology – EUROCRYPT 2006*, pp. 373–390. Springer (2006)
15. Rüdinger, A.: Restricted types of malleability in encryption schemes. Masters thesis, ETH Zürich (2011)