

Constructing Confidential Channels from Authenticated Channels—Public-Key Encryption Revisited

Sandro Coretti, Ueli Maurer, and Björn Tackmann

Department of Computer Science, ETH Zürich, Switzerland
{corettis,maurer,bjoernt}@inf.ethz.ch

Abstract. The security of public-key encryption (PKE), a widely-used cryptographic primitive, has received much attention in the cryptology literature. Many security notions for PKE have been proposed, including several versions of CPA-security, CCA-security, and non-malleability. These security notions are usually defined via a game that no efficient adversary can win with non-negligible probability or advantage.

If a PKE scheme is used in a larger protocol, then the security of this protocol is proved by showing a reduction of breaking a certain security property of the PKE scheme to breaking the security of the protocol. A major problem is that each protocol requires in principle its own tailor-made security reduction. Moreover, which security notion of the PKE scheme should be used in a given context is a priori not evident; the employed games model the use of the scheme abstractly through oracle access to its algorithms, and the sufficiency for specific applications is neither explicitly stated nor proven.

In this paper we propose a new approach to investigating the application of PKE, based on the constructive cryptography framework [24, 25]. The basic use of PKE is to enable confidential communication from a sender A to a receiver B , assuming A is in possession of B 's public key. One can distinguish two relevant cases: The (non-confidential) communication channel from A to B can be authenticated (e.g., because messages are signed) or non-authenticated. The application of PKE is shown to provide the construction of a secure channel from A to B from two (assumed) authenticated channels, one in each direction, or, alternatively, if the channel from A to B is completely insecure, the construction of a confidential channel without authenticity. Composition then means that the assumed channels can either be physically realized or can themselves be constructed cryptographically, and also that the resulting channels can directly be used in any applications that require such a channel. The composition theorem of constructive cryptography guarantees the soundness of this approach, which eliminates the need for separate reduction proofs.

We also revisit several popular game-based security notions (and variants thereof) and give them a constructive semantics by demonstrating which type of construction is achieved by a PKE scheme satisfying which notion. In particular, the necessary and sufficient security notions for the above two constructions to work are CPA-security and a variant of CCA-security, respectively.

1 Introduction

Public-key encryption (PKE) is a cryptographic primitive devised to achieve confidential communication in a context where only authenticated (but not confidential) communication channels are available [11, 34]. The cryptographic security of PKE is traditionally defined in terms of a certain distinguishing game in which no efficient adversary is supposed to achieve a non-negligible advantage. There exists quite a wide spectrum of security notions and variants thereof. These notions are motivated by clearly captured attacks (e.g., a chosen-ciphertext attack) that should be prevented, but in some cases they seem to have been proposed mainly because they are stronger than previous notions or can be shown to be incomparable.

This raises the question of which security notion for PKE is suitable or necessary for a certain higher-level protocol (using PKE) to be secure. The traditional answer to this question is that for each protocol one (actually, a cryptography expert) needs to identify the right security notion and provide a reduction proof to show that a PKE satisfying this notion yields a secure protocol.¹

An alternative approach is to capture the semantics of a security notion by characterizing directly what it achieves, making explicit in which applications it can be used securely. The constructive cryptography paradigm [24, 25] was proposed with this general goal in mind. Resources such as different types of communication channels are modeled explicitly, and the goal of a cryptographic protocol or scheme π is to *construct* a stronger or more useful resource S from an assumed resource R , denoted as $R \stackrel{\pi}{\Longrightarrow} S$. Two such construction steps can then be composed, i.e., if we additionally consider a protocol ψ that assumes the resource S and constructs a resource T , the composition theorem states that

$$R \stackrel{\pi}{\Longrightarrow} S \quad \wedge \quad S \stackrel{\psi}{\Longrightarrow} T \quad \Longrightarrow \quad R \stackrel{\psi \circ \pi}{\Longrightarrow} T,$$

where $\psi \circ \pi$ denotes the composed protocol.

Following the constructive paradigm, a protocol is built in a modular fashion from isolated construction steps. A security proof guarantees the soundness of one such step, and each proof is independent of the remaining steps. The composition theorem then guarantees that several such steps can be composed. While the general approach to protocol design based on reduction proofs is in principle sound, it is substantially more complex, more error-prone, and not suitable for re-use. This is part of the reason why it is generally not applied to the design of real-world protocols (e.g., TLS), which in turn is the main reason for the large number of protocol flaws discovered in the past. A major goal in cryptography must be to break the cycle of flaw discovery and fixes by providing solid proofs. Modularity appears to be the key in achieving this goal.

¹ Note that this work is orthogonal to the foundational problem of designing practical PKE schemes provably satisfying certain security notions, based on realistic hardness assumptions. The seminal CCA-secure PKE scheme based on the DDH-assumption by Cramer and Shoup [9, 10] falls into this category, as do, e.g., [13, 32, 19, 21, 35].

In this spirit, we treat the use of PKE as such a construction step. The contributions of this paper are two-fold. First, we show how one can construct, using PKE, confidential channels from authenticated and insecure channels (cf. Section 1.1 and Section 3). Second, we revisit several known game-based security notions (and variants thereof) and give them a constructive semantics, providing an explicit understanding of the application contexts for which a given notion is suitable (cf. Section 1.2 and Section 4). In Section 1.3 we describe how our results, although stated in a simpler setting, capture settings with multiple senders and the notion of corruption that exists in other frameworks, and in Section 1.4 we contrast the constructive paradigm with the approach of idealizing the properties of cryptographic schemes. Related work is discussed in Section 1.5.

1.1 Constructing Confidential Channels using PKE

From the perspective of constructive cryptography [24, 25], the purpose of a public-key encryption scheme is to construct a confidential channel from non-confidential channels. Here, a channel is a resource (or functionality) that involves a sender, a receiver, and—to model channels with different levels of security—an attacker. A channel generally allows the sender to transmit a message to the receiver; the security properties of a particular channel are captured by the capabilities available to the attacker, which might, e.g., include reading or modifying the messages in transmission.

The parties access the channel through interfaces that the channel provides and that are specific for each party. For example, the sender’s interface allows to input messages, and the receiver’s interface allows to receive them. We refer to the interfaces by labels A , B , and E , where A and B are the sender’s and the receiver’s interfaces, respectively, and E is the adversary’s interface. In this work, we consider the following four types of channels (from A to B ; channels in the opposite direction are defined analogously), using the notation from [27]:²

- An *insecure channel*, denoted $- \rightarrow$, allows the adversary to read, deliver, and to delete all messages input at A , as well as to inject its own messages.
- An *authenticated channel*, denoted $\bullet \diamond \rightarrow$, still allows to read all messages, but the adversary is limited to forwarding or deleting messages input at A .
- A *confidential channel*, denoted $\diamond \rightarrow \bullet$, only leaks the length of the messages but does not necessarily prevent injections.
- A *secure channel*, denoted $\bullet \diamond \rightarrow \bullet$, also only leaks the message length, and only allows the adversary to forward or delete messages input at A .

To use public-key encryption, the receiver initially generates a key pair and transmits the public key to the sender. The sender needs to obtain the correct public key, which corresponds to assuming that the channel from B to A is

² The “ \bullet ” in the notation signifies that the capabilities at the marked interface, i.e., sending or receiving, are exclusive to the respective party. If the “ \bullet ” is missing, the adversary also has these capabilities. The \diamond -symbol is explained in Section 2.4.

authenticated ($\leftarrow\bullet$ ³). To transmit a message confidentially, the sender then encrypts the message under the received public key and sends the ciphertext to the receiver over a channel that could be authenticated or completely insecure.

The exact type of channel that is constructed depends on the type of assumed channel used to transmit the ciphertext to the receiver: We show that if the assumed channel is authenticated ($\bullet\rightarrow$) and the PKE scheme is *ind-cpa*-secure, the constructed channel is a secure channel ($\bullet\rightarrow\bullet$). If the assumed channel is insecure (\rightarrow) and the PKE scheme is *ind-cca*-secure, the constructed channel is only confidential ($\rightarrow\bullet$). Using the above notation, for protocols π and π' based on *ind-cpa* and *ind-cca* encryption schemes, respectively, these constructions can be written as

$$[\leftarrow\bullet, \bullet\rightarrow] \stackrel{\pi}{\Longrightarrow} \bullet\rightarrow\bullet \quad \text{and} \quad [\leftarrow\bullet, \rightarrow] \stackrel{\pi'}{\Longrightarrow} \rightarrow\bullet,$$

where the bracket notation means that both resources in the brackets are available.

The notion of constructing the confidential (or secure) channel from the two assumed non-confidential ones is made precise in a simulation-based sense [25, 24], where the simulator can be interpreted as translating all attacks on the protocol into attacks on the constructed (ideal) channel. As the constructed channel is secure by definition, there are no attacks on the protocol.

The composability of the construction notion then means that the constructed channel can again be used as an assumed resource (possibly along with additional assumed or constructed resources) in other protocols. For instance, if a higher-level protocol uses the confidential channel to transmit a message together with a shared secret value in order to achieve an additionally authenticated (and hence fully secure) transmission of the message, then the proof of this protocol is based on the “idealized” confidential channel and does not (need to) include a reduction to the security of the encryption scheme. In the same spirit, the authenticated channel from B to A could be a physically authenticated channel, but it could also be constructed by using, for instance, a digital signature scheme to authenticate the transmission of the public key (which is done by certificates in practice).

1.2 Constructive Semantics of Game-Based Security Notions

Security properties for PKE are often formalized via a game between a hypothetical challenger and an attacker. We assign constructive semantics to several existing game-based definitions by first characterizing the appropriate assumed and constructed resources and then showing that the “standard use” of a PKE scheme over those channels (as illustrated in Section 1.1) achieves the construction if (and sometimes only if) it has the considered property.⁴

³ The simple arrow indicates that $\leftarrow\bullet$ is a single-use channel, i.e., only one message can be transmitted.

⁴ We point out that our negative results do *not* rule out the existence of other protocols that are derived from the scheme in some possibly more complicated way; those could still achieve the respective construction.

In particular, we show that `ind-cpa`-security is not only sufficient but also necessary for constructing a secure channel from two authenticated channels. For the construction of a confidential channel from an authenticated and an insecure channel, it turns out that `ind-cca`-security, while sufficient, is unnecessarily strong. The transformation only requires the weaker notion of `ind-rcca`-security, which was introduced by Canetti et al. [8] to avoid the artificial strictness of `ind-cca`. We continue the analysis of `ind-cca`-security and follow up on work by Bellare et al. [4], where several non-equivalent definitional variants are considered. We show that only the stricter notions they consider are sufficient for the channel construction, leaving the exact semantics of the weaker notions unclear.

We also consider non-adaptive CCA-security (`ind-cca1`) and non-malleability (`nm-cpa`). We show that both notions correspond to transformations between somewhat artificial channels, but might still be useful for specific applications.

1.3 Capturing Settings with Potentially Corrupted Senders

Although our security definitions for public-key encryption are phrased in a setting where there is only one legitimate sender (at the A -interface), our treatment can be “lifted” to a setting with multiple senders generically, cf. [29]. In a scenario with multiple senders, it is important to formulate the guarantees that are maintained if one or more of the senders deviate from the protocol because their machines are controlled by some attacker (or virus). This is captured in most security frameworks by considering an external adversary that has the capability of corrupting some of the parties. In the context of PKE and secure communication, the goal is to still provide confidentiality guarantees to non-corrupted senders. (If the receiver is corrupted, then no security can be guaranteed.)

The ability of an attacker to act on behalf of corrupted senders means that it can directly send (bogus) ciphertexts to the receiver, even if the communication to the receiver is authenticated. This capability corresponds exactly to the case of assuming only an unauthenticated channel, where the messages are injected via the E -interface. Hence, our treatment extends to the case of (static) sender corruption by considering the lifting that relates the interfaces of the parties in the multi-party scenario to the A -interface in the three-party setting, and provides all capabilities of the statically corrupted parties also at the E -interface.

In summary, the security of public-key encryption in the presence of potentially (statically) corrupted senders corresponds exactly to the construction of a confidential channel $\dashv\rightarrow\bullet$ from one insecure channel \dashrightarrow and one authenticated channel $\leftarrow\bullet$ in the opposite direction, as discussed in Section 1.1. This implies that in the presence of (static) corruption, `ind-rcca` security is required and sufficient both in the case where the channel from the sender to the receiver is authenticated, and also where it is not authenticated.

1.4 Idealizing Properties vs. Constructing Resources

The security guarantees that one requires from a cryptographic scheme can be modeled in fundamentally different ways, even within a single formal security

framework. One approach, which underlies the PKE functionality \mathcal{F}_{PKE} in [8], is to idealize the properties of the algorithms that comprise the scheme. Such a functionality corresponds to a cryptographic scheme, and its interfaces closely resemble the interfaces of the algorithms (although, e.g., the private key is never output by \mathcal{F}_{PKE}). In such a treatment, elements that are essential for using the scheme, such as the ciphertext or the public key, will still appear in the functionality, but they are idealized in that, e.g., the ciphertext is independent of the corresponding plaintext; the idealized scheme is unbreakable by definition.

Another—fundamentally different—approach is to explicitly model *resources* that are available to one or more parties. The communication channels we describe in Section 1.1 can be considered *network resources*; there are also functionalities in the UC framework, such as $\mathcal{F}_{\text{AUTH}}$ or \mathcal{F}_{SC} in [7], that can be interpreted in this way. More generally, one can also think of randomness, memory, or even computation as resources of this type. Following the constructive paradigm, the guarantees of a cryptographic scheme are *not* a resource, but modeled as the guarantee that the scheme transforms one (assumed) resource into another (constructed) resource.⁵ Compared to ideal functionalities of the above type, the description of resources tends to be simpler and easier to understand. For example, in the case of public-key encryption, the confidential channel does not need to specify implementation artifacts such as ciphertexts or public keys.

While both approaches allow to divide the security proof of a composite protocol into several steps that can be proven independently, only the second approach enables a fully modular protocol design. Each sub-protocol achieves a well-defined construction step transforming a resource R into a resource S , which abstracts from how S is achieved. A higher-level protocol can thus use such a resource S independently of how it is obtained, and the construction of S can be replaced with a different one without affecting the design or proof of the higher-level protocol. Concretely, a protocol using the resource $\dashv\diamond\blacktriangleright$ does not depend on whether or not the channel is constructed by a PKE scheme, whereas a protocol using the functionality \mathcal{F}_{PKE} will always be specific to this step.

1.5 Related Work

We provide here an abridged comparison with related work. A more comprehensive comparison can be found in the full version of this work.

Game-based security. The study of PKE security was initiated by Goldwasser and Micali [17], who introduced the notions of indistinguishability and semantic security. Yao’s [36] definition, based on computational entropy, was shown

⁵ By contrast, a typical UC security statement is that a cryptographic scheme implements some functionality. While statements about *hybrid* protocols in UC appear similar to constructive statements, they are less expressive since, e.g., the UC framework technically does not allow to make statements about assuming only *bounded* resources, as protocols that use hybrid functionalities can always instantiate arbitrarily many functionalities of a given type.

equivalent to variants of [17] by Micali et al. [30]. Goldreich [14, 15] made important modifications and also dealt with uniform adversaries. Today’s widely-used variant, *indistinguishability* under chosen-plaintext attack or *ind-cpa*, has been strengthened by considering more powerful attackers that can additionally obtain decryptions of arbitrary ciphertexts. This led to the notions of *ind-cca1* and *ind-cca2* (e.g., [31, 37]). Different variants of *ind-cca2*-security were compared by Bellare et al. [4]. Canetti et al. [8] introduced the weaker notion *ind-rcca* that suffices for many applications. A second important security property is *non-malleability*, introduced by Dolev et al. [12]. Informally, it requires that an adversary cannot change a ciphertext into one that decrypts to a related message. Variations of this notion have been considered in subsequent work [3, 5].

Real-world/ideal-world security. The idea of defining protocol security with respect to an ideal execution was first proposed by Goldreich et al. [16]; the concept of a simulator can be traced back to the seminal work by Goldwasser et al. [18] on zero-knowledge proofs. General security frameworks that allow the formalization of arbitrary functionalities to be realized by cryptographic protocols have been introduced by Canetti [6] as universal composability (UC) as well as by Backes et al. [33, 1] as reactive simulatability (RSIM). Treatments of PKE exist in both frameworks. The treatment in UC is with respect to an “ideal PKE” functionality; realizing this functionality is equivalent to *ind-cca2*-security [8]. Canetti and Krawczyk [7] formulate UC functionalities that model different types of communication channels and can be interpreted as network resources; they do not treat public-key encryption from this perspective. The formalization of the functionalities in [33] is closer to our approach, but less modular and hence formally more complex. In particular, the treatment is restricted to the case where the authenticated transmission of the ciphertexts is achieved by digital signatures instead of using a generic composition statement. More generally, both frameworks [6] and [33] are designed from a bottom-up perspective (starting from a selected machine model), whereas we follow the top-down approach of [25], which leads to simpler, more abstract definitions and statements.

Maurer et al. [26] described symmetric encryption following the constructive cryptography paradigm as the construction of confidential channels from non-confidential channels and shared keys, and compared the security definitions they obtained with game-based definitions. The goal of this work is to provide a comparable treatment for the case of PKE. In the same spirit, specific anonymity-related properties of PKE have been discussed by Kohlweiss et al. [22].

2 Preliminaries

2.1 Systems: Resources, Converters, Distinguishers, and Reductions

At the highest level of abstraction (following the hierarchy in [25]), systems are objects with interfaces by which they connect to (interfaces of) other systems; each interface is labeled with an element of a label set and connects to only a

single other interface. This concept of *abstract systems* captures the topological structures that result when multiple systems are connected in this manner.

The abstract systems concept, however, does not model the behavior of systems, i.e., *how* the systems interact via their interfaces. Consequently, statements about cryptographic protocols are statements at the next (lower) abstraction level. In this work, we describe all systems in terms of (probabilistic) discrete systems, which we explain in Section 2.2.

Resources and converters. *Resources* in this work are systems with three interfaces labeled by A , B , and E . A protocol is modeled as a pair of two so-called *converters* (one for each honest party), which are directed in that they have an *inside* and an *outside* interface, denoted by **in** and **out**, respectively. As a notational convention, we generally use upper-case, bold-face letters (e.g., \mathbf{R} , \mathbf{S}) or channel symbols (e.g., $\bullet \dashrightarrow$) to denote resources and lower-case Greek letters (e.g., α , β) or sans-serif fonts (e.g., **enc**, **dec**) for converters. We denote by Φ the set of all resources and by Σ the set of all converters.

The topology of a composite system is described using a term algebra, where each expression starts from one (or more) resources on the right-hand side and is subsequently extended with further terms on the left-hand side. An expression is interpreted in the way that all interfaces of the system it describes can be connected to interfaces of systems which are appended on the left. For instance, for a single resource $\mathbf{R} \in \Phi$, all its interfaces A , B , and E are accessible.

For $I \in \{A, B, E\}$, a resource $\mathbf{R} \in \Phi$, and a converter $\alpha \in \Sigma$, the expression $\alpha^I \mathbf{R}$ denotes the composite system obtained by connecting the inside interface of α to interface I of \mathbf{R} ; the outside interface of α becomes the I -interface of the composite system. The system $\alpha^I \mathbf{R}$ is again a resource (cf. Figure 1 on page 14).

For two resources \mathbf{R} and \mathbf{S} , $[\mathbf{R}, \mathbf{S}]$ denotes the parallel composition of \mathbf{R} and \mathbf{S} . For each $I \in \{A, B, E\}$, the I -interfaces of \mathbf{R} and \mathbf{S} are merged and become the *sub-interfaces* of the I -interface of $[\mathbf{R}, \mathbf{S}]$, which we denote by $I.1$ and $I.2$. A converter α that connects to the I -interface of $[\mathbf{R}, \mathbf{S}]$ has two inside sub-interfaces, denoted by **in.1** and **in.2**, where the first one connects to $I.1$ of \mathbf{R} and the second one connects to $I.2$ of \mathbf{S} .

Any two converters α and β can be composed sequentially by connecting the inside interface of β to the outside interface of α , written $\beta \circ \alpha$, with the effect that $(\beta \circ \alpha)^I \mathbf{R} = \beta^I \alpha^I \mathbf{R}$. Moreover, converters can also be taken in parallel, denoted by $[\alpha, \beta]$, with the effect that $[\alpha, \beta]^I [\mathbf{R}, \mathbf{S}] = [\alpha^I \mathbf{R}, \beta^I \mathbf{S}]$.

We assume the existence of an identity converter $\text{id} \in \Sigma$ with $\text{id}^I \mathbf{R} = \mathbf{R}$ for all resources $\mathbf{R} \in \Phi$ and interfaces $I \in \{A, B, E\}$ and of a special converter $\perp \in \Sigma$ with an inactive outside interface.

Distinguishers. A *distinguisher* is a special type of system \mathbf{D} that connects to all interfaces of a resource \mathbf{U} and outputs a single bit at the end of its interaction with \mathbf{U} . In the term algebra, this appears as the expression $\mathbf{D}\mathbf{U}$, which defines a binary random variable. The *distinguishing advantage of a distinguisher \mathbf{D} on two systems \mathbf{U} and \mathbf{V}* is defined as

$$\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) := |\mathbb{P}[\mathbf{D}\mathbf{U} = 1] - \mathbb{P}[\mathbf{D}\mathbf{V} = 1]|$$

and as $\Delta^{\mathcal{D}}(\mathbf{U}, \mathbf{V}) := \sup_{\mathbf{D} \in \mathcal{D}} \Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V})$ for a distinguisher class \mathcal{D} .

The distinguishing advantage measures how much the output distribution of \mathbf{D} differs when it is connected to either \mathbf{U} or \mathbf{V} . There is an equivalence notion on systems (which is defined on the discrete systems level), denoted by $\mathbf{U} \equiv \mathbf{V}$, which implies that $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) = 0$ for all distinguishers \mathbf{D} . The distinguishing advantage satisfies the triangle inequality, i.e., $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{W}) \leq \Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) + \Delta^{\mathbf{D}}(\mathbf{V}, \mathbf{W})$ for all resources \mathbf{U} , \mathbf{V} , and \mathbf{W} and distinguishers \mathbf{D} .

Games. We capture games defining security properties as distinguishing problems in which an adversary \mathbf{A} tries to distinguish between two *game systems* \mathbf{G}_0 and \mathbf{G}_1 . Game systems (or simply *games*) are single-interface systems, which appear, similarly to resources, on the right-hand side of the expressions in the term algebra. The adversary is a distinguisher that connects to a game (instead of a resource). We denote by \mathcal{A} the class of *all* adversaries for games.

Reductions. When relating two distinguishing problems, it is convenient to use a special type of system \mathbf{C} that translates one setting into the other. Formally, \mathbf{C} is a converter that has an *inside* and an *outside* interface. When it is connected to a system \mathbf{S} , which is denoted by \mathbf{CS} , the inside interface of \mathbf{C} connects to the (merged) interface(s) of \mathbf{S} and the outside interface of \mathbf{C} is the interface of the composed system. \mathbf{C} is called a *reduction system* (or simply *reduction*).

To reduce distinguishing two systems \mathbf{S}, \mathbf{T} to distinguishing two systems \mathbf{U}, \mathbf{V} , one exhibits a reduction \mathbf{C} such that $\mathbf{CS} \equiv \mathbf{U}$ and $\mathbf{CT} \equiv \mathbf{V}$.⁶ Then, for all distinguishers \mathbf{D} , we have $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) = \Delta^{\mathbf{D}}(\mathbf{CS}, \mathbf{CT}) = \Delta^{\mathbf{D}^{\mathbf{C}}}(\mathbf{S}, \mathbf{T})$. The last equality follows from the fact that \mathbf{C} can also be thought of as being part of the distinguisher.

2.2 Discrete Systems

Protocols that communicate by passing messages and the respective resources are described as (probabilistic) discrete systems. Their behavior can be formalized by random systems as in [23], i.e., as families of conditional probability distributions of the outputs (as random variables) given all previous inputs and outputs of the system. For systems with multiple interfaces, the interface to which an input or output is associated is specified as part of the input or output. For the restricted (but here sufficient) class of systems that for each input provide (at most) one output, an execution of a collection of systems is defined as the consecutive evaluation of the respective random systems (similarly to the models in [6, 20]).

2.3 The Notion of Construction

Recall that we consider resources with interfaces A , B , and E , where A and B are interfaces of honest parties and E is the interface of the adversary. We

⁶ For instance, we consider reductions from distinguishing game systems to distinguishing resources. Then, \mathbf{C} connects to a game on the inside and provides interfaces A , B , and E on the outside.

formalize the security of protocols via the following notion of *construction*, which was introduced in [24] (and is a special case of the abstraction notion from [25]):

Definition 1. Let Φ and Σ be as in Section 2.1. A protocol $\pi = (\pi_1, \pi_2) \in \Sigma^2$ constructs resource $\mathbf{S} \in \Phi$ from resource $\mathbf{R} \in \Phi$ within ε and with respect to distinguisher class \mathcal{D} , denoted

$$\mathbf{R} \stackrel{(\pi, \varepsilon)}{\Longrightarrow} \mathbf{S},$$

if

$$\left\{ \begin{array}{ll} \Delta^{\mathcal{D}}(\pi_1^A \pi_2^B \perp^E \mathbf{R}, \perp^E \mathbf{S}) \leq \varepsilon & (\text{availability}) \\ \exists \sigma \in \Sigma : \Delta^{\mathcal{D}}(\pi_1^A \pi_2^B \mathbf{R}, \sigma^E \mathbf{S}) \leq \varepsilon & (\text{security}). \end{array} \right.$$

The availability condition captures that a protocol must correctly implement the functionality of the constructed resource in the absence of the adversary. The security condition models the requirement that everything the adversary can achieve in the *real-world system* (i.e., the assumed resource with the protocol) he can also accomplish in the *ideal-world system* (i.e., the constructed resource with the simulator).

An important property of Definition 1 is its composability. Intuitively, if a resource \mathbf{S} is used in the construction of a larger system, then the composability implies that \mathbf{S} can be replaced by a construction $\pi_1^A \pi_2^B \mathbf{R}$ without affecting the security of the composed system. Security and availability are preserved under composition. More formally, if for some resources \mathbf{R} , \mathbf{S} , and \mathbf{T} and protocols π and ϕ , $\mathbf{R} \stackrel{(\pi, \varepsilon)}{\Longrightarrow} \mathbf{S}$ and $\mathbf{S} \stackrel{(\phi, \varepsilon')}{\Longrightarrow} \mathbf{T}$, then

$$\mathbf{R} \stackrel{(\phi \circ \pi, \varepsilon + \varepsilon')}{\Longrightarrow} \mathbf{T},$$

as well as

$$[\mathbf{R}, \mathbf{U}] \stackrel{([\pi, \text{id}], \varepsilon)}{\Longrightarrow} [\mathbf{S}, \mathbf{U}] \quad \text{and} \quad [\mathbf{U}, \mathbf{R}] \stackrel{([\text{id}, \pi], \varepsilon)}{\Longrightarrow} [\mathbf{U}, \mathbf{S}]$$

for any resource \mathbf{U} . More details can be found in [24].

2.4 Channels

We consider the types of channels shown on the right. Each channel initially expects a special cheating bit $b \in \{0, 1\}$ at interface E , indicating whether the adversary is present and intends to interfere

Channel Name	Symbol	$\ell(m)$	inj
Insecure Channel	$- \rightarrow$	m	\checkmark
Confidential Channel	$-\diamond \rightarrow \bullet$	$ m $	\checkmark
Authenticated Channel	$\bullet \diamond \rightarrow$	m	\times
Secure Channel	$\bullet \diamond \rightarrow \bullet$	$ m $	\times

with the transmission of the messages. The special converter \perp (cf. Section 2.1) always sets $b = 0$. For simplicity, we will assume that whenever \perp is not present, all cheating bits are set to 1.

A channel from A to B with leakage ℓ and message space $\mathcal{M} \subseteq \{0, 1\}^*$ is a resource with interfaces A , B , and E and behaves as follows:⁷ When the i^{th} message $m \in \mathcal{M}$ is input at interface A , it is recorded as (i, m) and $(i, \ell(m))$ is output at interface E . When (dlv, i') is input at interface E , if (i', m') has been recorded, m' is delivered at interface B . If injections are permissible, when (inj, m') is input at interface E , m' is output at interface B .⁸

The security statements in this work are parameterized by the number of messages that are transmitted over the channels. More precisely, for each of the above channels and each $n \in \mathbb{N}$, we define the n -bounded channel as the one that processes (only) the first n queries at the A -interface and the first n queries at the E -interface (as described above) and ignores all further queries at these interfaces. We then require from a protocol that it constructs, for all $n \in \mathbb{N}$, the n -bounded “ideal” channel from the n -bounded assumed channel. Wherever the number n is significant, such as in the theorem statements, we denote the n -bounded versions of channels by writing the n on top of the channel symbol (e.g., $\overset{n}{-\diamond-\twoheadrightarrow\bullet}$); we omit it in places that are of less formal nature.

Finally, a simple-arrow symbol (e.g., $\bullet \rightarrow$) denotes a *single-use* channel. That is, only one message may be transmitted.

2.5 Public-Key Encryption Schemes

A public-key encryption (PKE) scheme with message space $\mathcal{M} \subseteq \{0, 1\}^*$ and ciphertext space \mathcal{C} is defined as three algorithms $\Pi = (K, E, D)$, where the key-generation algorithm K outputs a key pair (pk, sk) , the (probabilistic) encryption algorithm E takes a message $m \in \mathcal{M}$ and a public key pk and outputs a ciphertext $c \leftarrow E_{\text{pk}}(m)$, and the decryption algorithm takes a ciphertext $c \in \mathcal{C}$ and a secret key sk and outputs a plaintext $m \leftarrow D_{\text{sk}}(c)$. The output of the decryption algorithm can be the special symbol \diamond , indicating an invalid ciphertext.

A PKE scheme is correct if $m = D_{\text{sk}}(E_{\text{pk}}(m))$ (with probability 1 over the randomness in the encryption algorithm) for all messages m and all key pairs (pk, sk) generated by K .

It will be more convenient to phrase bit-guessing games used in definitions of PKE security properties as a distinguishing problem between two game systems (cf. Section 2.1). We consider the following games, which correspond to the (standard) notions of ind-cpa (cpa for short), ind-cca2 (cca), ind-cca1 (cca1), ind-rcca (rcca), and nm-cpa (nm).⁹ Informally, a scheme is secure in the sense of a notion if efficient adversaries have negligible advantage in distinguishing the two corresponding game systems.

⁷ If the cheating bit is set to $b = 0$, all messages input at the sender interface A are immediately delivered to B .

⁸ Note that none of the channels prevents the adversary from reordering or replaying messages sent over the channel. The \diamond -symbol suggests the “internal buffer” in which the channel stores messages input at A .

⁹ We consider the so-called real-or-random versions of these games, which are equivalent to the more popular left-or-right formulations (as shown in [2] for symmetric encryption). For non-malleability, we use an indistinguishability-based version by [5].

CPA game. Consider the systems $\mathbf{G}_0^{\text{cpa}}$ and $\mathbf{G}_1^{\text{cpa}}$ defined as follows: For a PKE scheme Π , both initially run the key-generation algorithm to obtain (pk, sk) and output pk . Upon (the first) query (chall, m) , $\mathbf{G}_0^{\text{cpa}}$ outputs an encryption $c \leftarrow E_{\text{pk}}(m)$ of m and $\mathbf{G}_1^{\text{cpa}}$ an encryption $c \leftarrow E_{\text{pk}}(\bar{m})$, called the *challenge*, of a randomly chosen message \bar{m} of length $|m|$.

CCA games. For $b \in \{0, 1\}$, system $\mathbf{G}_b^{\text{cca1}}$ proceeds as $\mathbf{G}_b^{\text{cpa}}$ but additionally answers decryption queries (dec, c') before the challenge is output by returning $m' \leftarrow D_{\text{sk}}(c')$. $\mathbf{G}_b^{\text{cca}}$ answers decryption queries at any time unless c' equals the challenge c (if defined), in which case the answer is **test**.

RCCA game. Consider the systems $\mathbf{G}_0^{\text{rcca}}$ and $\mathbf{G}_1^{\text{rcca}}$ defined as follows: Initially, both run the key-generation algorithm to obtain (pk, sk) and output pk . Upon (the first) query (chall, m) , *both* choose a random message \bar{m} of length $|m|$. $\mathbf{G}_0^{\text{rcca}}$ outputs $c \leftarrow E_{\text{pk}}(m)$ and $\mathbf{G}_1^{\text{rcca}}$ outputs $c \leftarrow E_{\text{pk}}(\bar{m})$. Both systems answer decryption queries (dec, c') , but if $D_{\text{sk}}(c') \in \{m, \bar{m}\}$ (if m and \bar{m} are defined), the answer is **test**.

For more details about RCCA-security, see Section 4.2 or consult [8], where the notion was introduced.

NM game. Consider the systems \mathbf{G}_0^{nm} and \mathbf{G}_1^{nm} defined as follows: Both initially run the key-generation algorithm to obtain (pk, sk) and output pk . Upon (the first) query (chall, m) , \mathbf{G}_0^{nm} outputs an encryption $c \leftarrow E_{\text{pk}}(m)$ of m and \mathbf{G}_1^{nm} an encryption $c \leftarrow E_{\text{pk}}(\bar{m})$ of a randomly chosen message \bar{m} of length $|m|$. When a query $(\text{dec}, c_1, \dots, c_\ell)$ is input, both systems decrypt c_1, \dots, c_ℓ , return the resulting plaintexts (if any of the ciphertexts equal c , the corresponding plaintexts are replaced by **test**), and terminate the interaction.

2.6 Asymptotics

To allow for asymptotic security definitions, cryptographic protocols are often equipped with a so-called *security parameter*. We formulate all statements in this paper in a non-asymptotic fashion, but asymptotic statements can be obtained by treating systems \mathbf{S} as asymptotic families $\{\mathbf{S}_\kappa\}_{\kappa \in \mathbb{N}}$ and letting the distinguishing advantage be a real-valued function of κ . Then, for a given notion of efficiency, one can consider security w.r.t. classes of efficient distinguishers and a suitable negligibility notion. All reductions in this work are efficient with respect to the standard polynomial-time notions.

3 Constructing Confidential Channels with PKE

The main purpose of public-key encryption (PKE) is to achieve confidential communication. As a constructive statement, this means that we view a PKE scheme Π as a protocol, a pair of converters (enc, dec) , whose goal is to construct a confidential channel from non-confidential channels. Differentiating between the two

cases where the communication from the sender to the receiver is authenticated and unauthenticated, this is written as

$$[\leftarrow\bullet, \bullet\rightarrow] \stackrel{(\text{enc}, \text{dec})}{\Longrightarrow} \bullet\rightarrow\bullet \quad (1) \quad \text{and} \quad [\leftarrow\bullet, -\rightarrow] \stackrel{(\text{enc}, \text{dec})}{\Longrightarrow} -\rightarrow\bullet, \quad (2)$$

respectively.

In both cases, the *single-use* channel $\leftarrow\bullet$ captures the ability of the sender to obtain the receiver's public key in an authenticated fashion. In construction (1), the communication from the sender A to the receiver B is authenticated, which is modeled by the channel $\bullet\rightarrow$. The goal is to achieve a secure channel $\bullet\rightarrow\bullet$, which only leaks the length of the messages sent at interface A . In construction (2), the communication from A to B is completely insecure, which is captured by the insecure channel $-\rightarrow$. Here, the goal is to achieve a confidential channel $-\rightarrow\bullet$, which still hides messages input at the A -interface but also allows to inject arbitrary messages at E .

In the following, we first show how a PKE scheme Π can be seen as a converter pair (enc, dec) . We then prove that (enc, dec) achieves construction (1) if the underlying PKE scheme is cpa -secure, and construction (2) if the underlying PKE scheme is cca -secure. We also briefly discuss the usefulness of the constructed channels.

3.1 PKE Schemes as Protocols

Let $\Pi = (K, E, D)$ be a PKE scheme. Based on Π , we define a pair of protocol converters (enc, dec) for constructions (1) and (2). Both converters have two sub-interfaces in.1 and in.2 on the inside, as we connect them to a resource that is a parallel composition of two other resources (cf. Section 2.1).

Converter enc works as follows: It initially expects a public key pk at in.1 . When a message m is input at the outside interface out , enc outputs $c \leftarrow E_{\text{pk}}(m)$ at in.2 . Converter dec initially generates a key pair (pk, sk) using key-generation algorithm K and outputs pk at in.1 . When dec receives c' at in.2 , it computes $m' \leftarrow D_{\text{sk}}(c')$ and, if $m' \neq \diamond$, outputs m' at the outside interface out .

3.2 Constructing a Secure from Two Authenticated Channels

Towards proving that the protocol (enc, dec) indeed achieves construction (1), note first that the correctness of Π implies that the *availability* condition of Definition 1 is satisfied. To prove *security*, we need to exhibit a simulator σ such that the assumed resource $[\leftarrow\bullet, \bullet\rightarrow]$ with the protocol converters is indistinguishable from the constructed resource $\bullet\rightarrow\bullet$ with the simulator (cf. Figure 1).

Theorem 1 implies that (enc, dec) realizes (1) if the underlying PKE scheme is cpa -secure.

Theorem 1. *There exists a simulator σ and for any $n \in \mathbb{N}$ there exists a (efficient) reduction \mathbf{C} such that for every \mathbf{D} ,*

$$\Delta^{\mathbf{D}}(\text{enc}^A \text{dec}^B [\leftarrow\bullet, \bullet\rightarrow], \sigma^E \bullet\rightarrow) \leq n \cdot \Delta^{\mathbf{DC}}(\mathbf{G}_0^{\text{cpa}}, \mathbf{G}_1^{\text{cpa}}).$$




Fig. 1. Left: The assumed resource (two authenticated channels) with protocol converters enc and dec attached to interfaces A and B , denoted $\text{enc}^A \text{dec}^B [\leftarrow \bullet, \bullet \rightarrow \bullet]$. Right: The constructed resource (a secure channel) with simulator σ attached to the E -interface, denoted $\sigma^E \bullet \rightarrow \bullet$. In particular, σ must simulate the E -interfaces of the two authenticated channels. The protocol is secure if the two systems are indistinguishable.

Proof. First, consider the following simulator σ for interface E of $\bullet \rightarrow \bullet$, which has two sub-interfaces, denoted by out.1 and out.2 , on the outside (since the real-world system has two sub-interfaces at E): Initially, σ generates a key pair (pk, sk) and outputs $(1, \text{pk})$ at out.1 . When it receives (i, l) at the inside interface in , σ generates an encryption $c \leftarrow E_{\text{pk}}(\bar{m})$ of a randomly chosen message \bar{m} of length l and outputs (i, c) at out.2 . When (dlv, i') is input at out.2 , σ simply outputs (dlv, i') at in .

Consider the two systems $\mathbf{U} := \text{enc}^A \text{dec}^B [\leftarrow \bullet, \bullet \rightarrow \bullet]$ and $\mathbf{V} := \sigma^E \bullet \rightarrow \bullet$. Distinguishing $\mathbf{G}_0^{\text{cpa}}$ from $\mathbf{G}_1^{\text{cpa}}$ can be reduced to distinguishing these two systems via the following reduction system \mathbf{C}' , which connects to a game on the inside and provides interfaces A , B , and E on the outside (cf. Section 2.1 for details on reduction systems): Initially, \mathbf{C}' takes a value pk from the game (on the inside) and outputs $(1, \text{pk})$ at the (outside) $E.1$ -interface. When a message m is input at the A -interface of \mathbf{C}' , it is passed as (chall, m) to the game. The resulting challenge c is output as $(1, c)$ at the $E.2$ -interface. When $(\text{dlv}, 1)$ is input at the $E.2$ -interface, \mathbf{C}' outputs m at interface B .

We have $\mathbf{C}' \mathbf{G}_0^{\text{cpa}} \equiv \mathbf{U}$ and $\mathbf{C}' \mathbf{G}_1^{\text{cpa}} \equiv \mathbf{V}$, and thus

$$\begin{aligned} \Delta^{\mathbf{D}}(\text{enc}^A \text{dec}^B [\leftarrow \bullet, \bullet \rightarrow \bullet], \sigma^E \bullet \rightarrow \bullet) &\leq n \cdot \Delta^{\mathbf{DC}''}(\mathbf{U}, \mathbf{V}) \\ &= n \cdot \Delta^{\mathbf{DC}''}(\mathbf{C}' \mathbf{G}_0^{\text{cpa}}, \mathbf{C}' \mathbf{G}_1^{\text{cpa}}) \\ &= n \cdot \Delta^{\mathbf{DC}}(\mathbf{G}_0^{\text{cpa}}, \mathbf{G}_1^{\text{cpa}}), \end{aligned}$$

where $\mathbf{C} := \mathbf{C}'' \mathbf{C}'$ and the first inequality follows from a standard hybrid argument for a reduction system \mathbf{C}'' (deferred to the full version). \square

3.3 Confidential Channels from Authenticated and Insecure Ones

To prove that the protocol (enc, dec) achieves construction (2), we need to again exhibit a simulator σ such that the assumed resource $[\leftarrow \bullet, - \rightarrow \bullet]$ with the protocol converters is indistinguishable from the constructed resource $\rightarrow \bullet$ with

the simulator, as done in Theorem 2, which implies that (enc, dec) realizes (2) if the underlying PKE scheme is cca -secure. We defer the proof to the full version.

Theorem 2. *There exists a simulator σ and for any $n \in \mathbb{N}$ there exists a (efficient) reduction \mathbf{C} such that for every \mathbf{D} ,*

$$\Delta^{\mathbf{D}}(\text{enc}^A \text{dec}^B[\leftarrow \bullet, - \xrightarrow{n}], \sigma^E \dashrightarrow \bullet) \leq n \cdot \Delta^{\mathbf{DC}}(\mathbf{G}_0^{\text{cca}}, \mathbf{G}_1^{\text{cca}}).$$

The confidential channel $\dashrightarrow \bullet$ is the best channel one can construct from the two assumed channels. As the E -interface has the same capabilities as the A -interface at both the authenticated (from B to A) and the insecure channels, it will necessarily also be possible to inject messages to the receiver via the E -interface by simply applying the sender's protocol converter.

3.4 Applicability of the Constructed Channels

The plain use of PKE yields constructions (1) and (2), i.e., one obtains the resources $\bullet \dashrightarrow \bullet$ and $\dashrightarrow \bullet$. Both channels allow the adversary to reorder or replace the messages sent by A . In practice, where PKE is often used to encapsulate symmetric keys, it is important, however, that keys used in various protocols by different users are independent. Thus, it is more useful to obtain independent single-use channels $[\bullet \rightarrow \bullet, \dots, \bullet \rightarrow \bullet]$ and $[\rightarrow \bullet, \dots, \rightarrow \bullet]$ instead of $\bullet \dashrightarrow \bullet$ and $\dashrightarrow \bullet$, respectively.

In the authenticated setting, given independent authenticated channels, protocol (enc, dec) (with only formal modifications) achieves the construction

$$[\leftarrow \bullet, \bullet \rightarrow, \dots, \bullet \rightarrow] \stackrel{(\text{enc}, \text{dec})}{\Longleftrightarrow} [\bullet \rightarrow \bullet, \dots, \bullet \rightarrow \bullet].$$

In the unauthenticated setting, however, the analogous construction

$$[\leftarrow \bullet, \rightarrow, \dots, \rightarrow] \stackrel{(\text{enc}, \text{dec})}{\Longleftrightarrow} [\rightarrow \bullet, \dots, \rightarrow \bullet]$$

is not achieved by (enc, dec) since, due to the absence of authenticity, the adversary can freely take a ciphertext it observes on one of the insecure channels \rightarrow and insert it into another one. Thus, the ideal resource cannot consist of independent channels. This issue can be taken care of by (explicitly) introducing session identifiers (SIDs). A systematic treatment of handling multiple sessions and senders can be found in [29].

4 Constructive Semantics of Game-Based Notions

We analyze several game-based security notions from a constructive viewpoint. We complete the analysis of cpa -security from Section 3.2 by showing that it is also necessary to achieve construction (1). Moreover, we explain why the notion of cca is unnecessarily strict for construction (2) and that the construction in fact only requires the weaker notion of rcca introduced in [8].

Then, we follow up on work by Bellare et al. [4], who compared several variants of defining *cca*-security, and showed that only the stricter notions they consider are sufficient for construction (2). We also provide constructive semantics for non-adaptive chosen-ciphertext security and non-malleability.

4.1 CPA Security is Necessary for Construction (1)

We prove in Section 3.2 that indistinguishability under chosen-plaintext attacks, *cpa*-security, suffices to construct a secure channel from two authenticated channels. Here, we show that it is also necessary. That is, if protocol (enc, dec) , based on a PKE scheme Π as shown in Section 3.1, achieves the construction, then Π must be *cpa*-secure.

In the following, let

$$\mathbf{U} := \text{enc}^A \text{dec}^B [\leftarrow \bullet, \bullet \diamond \rightarrow] \quad \text{and} \quad \mathbf{V} := \sigma^E \bullet \diamond \rightarrow \bullet,$$

where σ is an *arbitrary* simulator.

Theorem 3. *There exist (efficient) reductions \mathbf{C}_0 and \mathbf{C}_1 such that for all adversaries \mathbf{A} ,*

$$\Delta^{\mathbf{A}}(\mathbf{G}_0^{\text{cpa}}, \mathbf{G}_1^{\text{cpa}}) \leq \Delta^{\mathbf{AC}_0}(\mathbf{U}, \mathbf{V}) + \Delta^{\mathbf{AC}_1}(\mathbf{U}, \mathbf{V}).$$

Proof. Consider the following reduction systems \mathbf{C}_0 and \mathbf{C}_1 , both connecting to an $\{A, B, E\}$ -resource on the inside and providing a single interface on the outside (for the adversary): Initially, both obtain $(1, \text{pk})$ at the inside $E.1$ -interface and output pk at the outside interface. When (chall, m) is received on the outside, \mathbf{C}_0 outputs m at the inside A -interface and \mathbf{C}_1 a randomly chosen message \bar{m} of length $|m|$. Subsequently, $(1, c)$ is received at the inside $E.2$ -interface, and c is output (as the challenge) on the outside by both systems. We have

$$\mathbf{C}_0 \mathbf{U} \equiv \mathbf{G}_0^{\text{cpa}} \quad \text{and} \quad \mathbf{C}_1 \mathbf{U} \equiv \mathbf{G}_1^{\text{cpa}} \quad \text{and} \quad \mathbf{C}_0 \mathbf{V} \equiv \mathbf{C}_1 \mathbf{V},$$

where the last equivalence follows from the fact that, in \mathbf{V} , the input from $\bullet \diamond \rightarrow \bullet$ to σ is the same in both systems (the length of the message input at the A -interface of $\bullet \diamond \rightarrow \bullet$), and therefore they behave identically. Hence,

$$\begin{aligned} \Delta^{\mathbf{A}}(\mathbf{G}_0^{\text{cpa}}, \mathbf{G}_1^{\text{cpa}}) &= \Delta^{\mathbf{A}}(\mathbf{C}_0 \mathbf{U}, \mathbf{C}_1 \mathbf{U}) \\ &\leq \Delta^{\mathbf{A}}(\mathbf{C}_0 \mathbf{U}, \mathbf{C}_0 \mathbf{V}) + \Delta^{\mathbf{A}}(\mathbf{C}_0 \mathbf{V}, \mathbf{C}_1 \mathbf{V}) + \Delta^{\mathbf{A}}(\mathbf{C}_1 \mathbf{V}, \mathbf{C}_1 \mathbf{U}) \\ &= \Delta^{\mathbf{AC}_0}(\mathbf{U}, \mathbf{V}) + \Delta^{\mathbf{AC}_1}(\mathbf{U}, \mathbf{V}). \end{aligned}$$

□

4.2 RCCA Security is Necessary for Construction (2)

Indistinguishability under chosen-ciphertext attacks, *cca*-security, suffices to construct a confidential channel from an authenticated and an insecure one (cf. Section 3.3). It is, however, unnecessarily strict, as can be seen from the following

example, adapted from [8]: Let Π be a PKE scheme and assume it is cca -secure. Consider a modified scheme Π' that works exactly as Π , except that a 0-bit is appended to every encryption, which is ignored during decryption. It is easily seen that Π' is not cca -secure, since the adversary can obtain a decryption of the challenge ciphertext by flipping its last bit and submitting the result to the decryption oracle. PKE scheme Π' can, however, still be used to achieve construction (2) using a simulator that issues the dlv -instruction to $\leftarrow \diamond \rightarrow \bullet$ whenever a recorded ciphertext is received at the outside interface or one where flipping the last bit results in a recorded ciphertext (cf. full version for more details).

Canetti et al. [8] introduced the notion of *replayable chosen ciphertext* security, rcca , which is more permissive in that it allows the adversary to transform a ciphertext into one that decrypts to the same message. In the full version of this paper, we show that if protocol (enc, dec) , based on a PKE scheme Π (cf. Section 3.1), achieves (2), then Π must be rcca -secure, and that rcca is also sufficient for the construction if the message space of Π is sufficiently large.

4.3 Variants of Chosen-Ciphertext Security

Bellare et al. [4] analyze several ways of enforcing the condition that the adversary must not query the challenge ciphertext c to the decryption oracle. They consider modifications along two axes: First, the condition can be enforced during the entire game (b for *both* phases) or only in the second phase (s for *second* phase), i.e., after the c has been given to the adversary. Second, one can either exclude adversaries with a non-zero probability of violating the condition from consideration (e for *exclusion*) or penalize an adversary (by declaring the game lost) whenever he asks the challenge c (p for *penalty*). The combination of these choices yields four *non-equivalent* notions ind-cca-sp , ind-cca-se , ind-cca-bp , ind-cca-be . The s-notions are equivalent to each other and to our formulation of cca -security (cf. Section 2.5). The e-notions are strictly weaker and do in fact not even imply cca1 -security [4]. Since cca1 -security is weaker than rcca -security and rcca is needed for construction (2), they are not sufficient for (2).

4.4 Non-Malleability

Informally, a non-malleable PKE scheme is such that the adversary cannot transform a ciphertext into one that decrypts to a related message. We consider the notion of non-malleability under chosen-plaintext attacks, nm-cpa , and show that from a PKE scheme with this property we can build a protocol $(\text{enc}'', \text{dec}'')$ that achieves the construction

$$[\leftarrow \bullet, - \rightarrow] \stackrel{(\text{enc}'', \text{dec}'')}{\iff} \leftarrow \diamond \rightarrow \bullet, \tag{3}$$

where $- \rightarrow$ works like $- \rightarrow$ but halts when halt is input at B and where the channel $\leftarrow \diamond \rightarrow \bullet$ is defined as follows: It internally keeps an initially empty list \mathcal{L} of messages. When the i^{th} message m is input at interface A , it is recorded as (i, m) and $(i, |m|)$ is output at interface E . When (dlv, i') is input at interface E

and if (i', m') has been recorded, m' is appended to \mathcal{L} . When (inj, m') is input at interface E , m' is appended to \mathcal{L} . When $dlv\text{-all}$ is input at B , all messages in \mathcal{L} are output at B , and the channel halts.

The protocol converters (enc'', dec'') are built as (enc, dec) in Section 3.1, except that dec'' only outputs the messages it received once $dlv\text{-all}$ is input at the outside interface, at which time it also outputs $halt$ at its inside interface and halts. In the full version of this paper, we prove that (enc'', dec'') achieves construction (3) if Π is $nm\text{-cpa}$ -secure.

The assumed channel $\dashv\!\!\dashv$ could itself be constructed in a setting where A and B have synchronized clocks and B buffers all messages until an agreed point in time, when A also stops sending. By the composition theorem, the channel that is constructed in this manner can then serve as the assumed channel in construction (3) to construct the channel $\circ\text{-}\dashv\!\!\dashv\!\!\bullet$ using PKE. This channel may then for instance be useful for running a protocol implementing a blind auction.

4.5 Non-Adaptive Chosen-Ciphertext Security

$ind\text{-cca1}$ -security, is defined via a game \mathbf{G}^{cca1} , which works as \mathbf{G}^{cca} except that no decryption queries are answered once the adversary has been given the challenge ciphertext. The most natural way to translate this into a constructive statement is to consider the construction of a (type of) confidential channel $\circ\text{-}\dashv\!\!\dashv\!\!\bullet$ where the adversary can inject messages at interface E only as long as no message has been input at A from an insecure channel $\circ\text{-}\dashv\!\!\dashv$ with the same property.

In the full version of this paper, we show that protocol (enc, dec) built from a $cca1$ -secure PKE scheme Π as in Section 3.1 achieves

$$[\leftarrow\bullet, \circ\text{-}\dashv\!\!\dashv] \stackrel{(enc'', dec'')}{\iff} \circ\text{-}\dashv\!\!\dashv\!\!\bullet. \quad (4)$$

Although this construction seems artificial, as with construction (3), it can be used in any setting where the assumed channel is an appropriate modeling of an available physical channel (or can itself be constructed from such a channel).

5 Conclusions

The purpose of this paper is to present the basic ways of applying PKE (within a larger protocol) as constructive steps, to be used for the modular design of complex protocols, thus taming the complexity of security-protocol design. To be ultimately applicable to full-fledged real-world protocols, other relevant cryptographic primitives also need to be modeled in the same way. While for symmetric encryption and MACs this was explained in [28, 26], and for commitments in [25], treating digital signatures and other cryptographic schemes and security mechanisms (sequence numbers, session identifiers, etc.) in constructive cryptography is left for future work (cf. [29]).

Acknowledgments. The work was supported by the Swiss National Science Foundation (SNF), project no. 200020-132794.

References

1. Backes, M., Pfitzmann, B., Waidner, M.: The Reactive Simulatability (RSIM) Framework for Asynchronous Systems. *Information and Computation* 205(12), 1685–1720 (December 2007)
2. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A Concrete Security Treatment of Symmetric Encryption. In: 38th FOCS. pp. 394–403 (1997)
3. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations Among Notions of Security for Public-Key Encryption Schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998)
4. Bellare, M., Hofheinz, D., Kiltz, E.: Subtleties in the Definition of IND-CCA: When and How Should Challenge-Decryption be Disallowed? *Cryptology ePrint Archive* 2009/418 (2009)
5. Bellare, M., Sahai, A.: Non-malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 519–536. Springer, Heidelberg (1999)
6. Canetti, R.: Universally Composable Security: A New Paradigm for Cryptographic Protocols. *Cryptology ePrint Archive*, Report 2000/067 (2000)
7. Canetti, R., Krawczyk, H.: Universally Composable Notions of Key Exchange and Secure Channels. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 337–351. Springer, Heidelberg (2002)
8. Canetti, R., Krawczyk, H., Nielsen, J.B.: Relaxing Chosen-Ciphertext Security. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 565–582. Springer, Heidelberg (2003)
9. Cramer, R., Shoup, V.: A Practical Public Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
10. Cramer, R., Shoup, V.: Design and Analysis of Practical Public-Key Encryption Schemes Secure against Adaptive Chosen Ciphertext Attack. *SIAM Journal on Computing* 33, 167–226 (2001)
11. Diffie, W., Hellman, M.E.: New Directions in Cryptography. *IEEE Transactions on Information Theory* 22(6), 644–654 (1976)
12. Dolev, D., Dwork, C., Naor, M.: Non-Malleable Cryptography (Extended Abstract). In: 23rd ACM STOC. pp. 542–552 (1991)
13. Fujisaki, E., Okamoto, T., Pointcheval, D., Stern, J.: RSA-OAEP Is Secure under the RSA Assumption. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 260–274. Springer, Heidelberg (2001)
14. Goldreich, O.: *Foundations of Cryptography*. Class Notes (Spring 1989), Technion University
15. Goldreich, O.: A Uniform-Complexity Treatment of Encryption and Zero-Knowledge. *Journal of Cryptology* 6(1), 21–53 (1993)
16. Goldreich, O., Micali, S., Wigderson, A.: How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In: 19th ACM STOC. pp. 218–229 (1987)
17. Goldwasser, S., Micali, S.: Probabilistic Encryption. *Journal of Computer and System Sciences* 28(2), 270–299 (1984)
18. Goldwasser, S., Micali, S., Rackoff, C.: The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In: 17th ACM STOC. pp. 291–304 (1985)
19. Hofheinz, D., Kiltz, E.: Practical Chosen Ciphertext Secure Encryption from Factoring. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 313–332. Springer, Heidelberg (2009)

20. Hofheinz, D., Shoup, V.: GNUC: A New Universal Composability Framework. Cryptology ePrint Archive, Report 2011/303 (2011)
21. Kiltz, E., Pietrzak, K., Stam, M., Yung, M.: A New Randomness Extraction Paradigm for Hybrid Encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 590–609. Springer, Heidelberg (2009)
22. Kohlweiss, M., Maurer, U., Onete, C., Tackmann, B., Venturi, D.: Anonymity-Preserving Public-Key Encryption: A Constructive Approach. In: Cristofaro, E.D., Wright, M. (eds.) PETS 2013. LNCS, vol. 7981, pp. 19–39. Springer, Heidelberg (2013)
23. Maurer, U.: Indistinguishability of Random Systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002)
24. Maurer, U.: Constructive Cryptography—A New Paradigm for Security Definitions and Proofs. In: Moedersheim, S., Palamidessi, C. (eds.) TOSCA 2011. LNCS, vol. 6993, pp. 33–56. Springer, Heidelberg (Apr 2011)
25. Maurer, U., Renner, R.: Abstract Cryptography. In: Chazelle, B. (ed.) The Second Symposium in Innovations in Computer Science, ICS 2011. pp. 1–21. Tsinghua University Press (Jan 2011)
26. Maurer, U., Rüdlinger, A., Tackmann, B.: Confidentiality and Integrity: A Constructive Perspective. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 209–229. Springer, Heidelberg (2012)
27. Maurer, U., Schmid, P.E.: A Calculus for Security Bootstrapping in Distributed Systems. *Journal of Computer Security* 4(1), 55–80 (1996)
28. Maurer, U., Tackmann, B.: On the Soundness of Authenticate-then-Encrypt: Formalizing the Malleability of Symmetric Encryption. In: ACM CCS. pp. 505–515 (2010)
29. Maurer, U., Tackmann, B., Coretti, S.: Key Exchange with Unilateral Authentication: Composable Security Definition and Modular Protocol Design. Cryptology ePrint Archive, Report 2013/555 (2013)
30. Micali, S., Rackoff, C., Sloan, B.: The Notion of Security for Probabilistic Cryptosystems. *SIAM Journal on Computing* 17(2), 412–426 (1988)
31. Naor, M., Yung, M.: Public-key Cryptosystems Provably Secure Against Chosen Ciphertext Attacks. In: 22nd ACM STOC. pp. 427–437 (1990)
32. Peikert, C., Waters, B.: Lossy Trapdoor Functions and Their Applications. In: 40th ACM STOC. pp. 187–196 (2008)
33. Pfitzmann, B., Waidner, M.: A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission. In: IEEE Symposium on Security and Privacy. pp. 184–200 (2001)
34. Rivest, R.L., Shamir, A., Adleman, L.M.: A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Communications of the ACM* 21(2), 120–126 (1978)
35. Rosen, A., Segev, G.: Chosen-Ciphertext Security via Correlated Products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009)
36. Yao, A.C.C.: Theory and Applications of Trapdoor Functions (Extended Abstract). In: 23rd FOCS. pp. 80–91 (1982)
37. Zheng, Y., Seberry, J.: Practical Approaches to Attaining Security Against Adaptively Chosen Ciphertext Attacks (Extended Abstract). In: Brickell, E.F. (ed.) CRYPTO '92. LNCS, vol. 740, pp. 292–304. Springer, Heidelberg (1992)