

State Separation for Code-Based Game-Playing Proofs

Chris Brzuska¹, Antoine Delignat-Lavaud², Cédric Fournet², Konrad Kohbrok¹, and Markulf Kohlweiss^{2,3}

¹Aalto University

²Microsoft Research

³University of Edinburgh

May 17, 2021

Abstract

The security analysis of real-world protocols involves reduction steps that are conceptually simple but still have to account for many protocol complications found in standards and implementations. Taking inspiration from universal composability, abstract cryptography, process algebras, and type-based verification frameworks, we propose a method to simplify large reductions, avoid mistakes in carrying them out, and obtain concise security statements.

Our method decomposes monolithic games into collections of stateful *packages* representing collections of oracles that call one another using well-defined interfaces. Every component scheme yields a pair of a real and an ideal package. In security proofs, we then successively replace each real package with its ideal counterpart, treating the other packages as the reduction. We build this reduction by applying a number of algebraic operations on packages justified by their state separation. Our method handles reductions that emulate the game perfectly, and leaves more complex arguments to existing game-based proof techniques such as the code-based analysis suggested by Bellare and Rogaway. It also facilitates computer-aided proofs, inasmuch as the perfect reductions steps can be automatically discharged by proof assistants.

We illustrate our method on two generic composition proofs: (1) a proof of self-composition using a hybrid argument; and (2) the composition of keying and keyed components. For concreteness, we apply them to the KEM-DEM proof of hybrid-encryption by Cramer and Shoup and to the composition of forward-secure game-based key exchange protocols with symmetric-key protocols.

1 Introduction

Code-based game-playing by Bellare and Rogaway [BR06] introduces pseudo-code as a precise tool for cryptographic reasoning. Following in their footsteps, we would like to

reason about games using code, rather than interactive Turing machines [vLW01]. Our code uses state variables and function calls, doing away with the details of operating on local tapes and shared tapes. Function calls enable straightforward code composition, defined for instance by inlining, and enjoy standard but useful properties, such as associativity. In the following, we refer to code units \mathcal{A} , \mathbf{R} and \mathbf{G} as *code packages*. If adversary \mathcal{A} calls reduction \mathbf{R} and \mathbf{R} calls game \mathbf{G} , we may see it either as code \mathbf{A} -calling- \mathbf{R} that calls code \mathbf{G} , or as code \mathcal{A} calling code \mathbf{R} -calling- \mathbf{G} . This form of associativity is used to define reductions, e.g., in abstract cryptography and in Rosulek’s book *The Joy of Cryptography* [Ros18].

As a first example, consider indistinguishability under chosen plaintext attacks, coded as a game IND-CPA^b with secret bit b , and let \mathcal{A} be an adversary that interacts with this game by calling its encryption oracle, which we write $\mathcal{A} \circ \text{IND-CPA}^b$. As a construction, consider a symmetric encryption scheme based on a pseudorandom function (PRF). We can decompose IND-CPA^b into some corresponding wrapper MOD-CPA that calls PRF^b , where b now controls idealization of the PRF. The equality $\text{IND-CPA}^b = \text{MOD-CPA} \circ \text{PRF}^b$ can be checked syntactically (and can be automatically discharged by proof assistants). IND-CPA security follows from PRF security using MOD-CPA as reduction:

$$\mathcal{A} \circ (\text{MOD-CPA}) \circ \text{PRF}^b \stackrel{\text{code}}{\equiv} (\mathcal{A} \circ \text{MOD-CPA}) \circ \text{PRF}^b.$$

Appendix A presents this example in more detail, including a discussion of our definitional choices. In particular, we encode all games as decisional games between a real game and an ideal game, following the tradition of [Can01], [Mau11] and [Bla08].

KEM-DEM.

Our second example, the composition of a key encapsulation mechanism (KEM) with a one-time deterministic encryption scheme (DEM), involves associativity and *interchange*, another form of code rearrangement (defined in Section 2). Cramer and Shoup [CS03] show that the composition of a KEM and a DEM that are both indistinguishable under chosen ciphertext attacks (IND-CCA) results in an IND-CCA public-key encryption scheme. We give a new formulation of their proof.

We here consider the standard IND-CCA notion, as did Cramer and Shoup. Note that in a previous version, we used $\$$ -IND-CCA (pseudorandom ciphertexts), since we originally (wrongly) believed that this might lead to a more tight reduction. We thank Théo Winterhalter for pointing out to us that the code equivalence step in the original version did not hold for $b = 1$.

We first reduce to the security of the KEM, replacing the encapsulated KEM key with a uniformly random key, then we reduce to the security of the DEM, which requires such a key, finally we ‘de-idealize’ the KEM and replace the key again with the real key. To facilitate these three reductions and analogously to the previous example, we decompose



Figure 1: Decomposed KEM and DEM games

the PKE-CCA game for public-key encryption into a wrapper MOD-CCA that calls the games for KEM and DEM security. That is, we use a *parallel* composition of the KEM and the DEM game. As the KEM and the DEM share the encapsulated KEM key, we need to enable state-sharing between both games. We achieve this by also decomposing the KEM and DEM security games into two packages such that they both contain a so-called KEY package that stores the shared key.

The KEM Game. Fig. 1a depicts the decomposed IND-CCA KEM game using a KEY package (also see page 15, Def. 16). The formal semantics of the graph-based notation of package composition is introduced in Section 2.2.

The IND-CCA KEM game allows the adversary to make a KEMGEN query to initialize the game as well as encapsulation queries ENCAP and decapsulation queries DECAP. Upon receiving an encapsulation query ENCAP, the KEM package makes a SET(k) query to KEY to store the real encapsulation key k , if the bit b is 0. In turn, if the bit b is 1, the KEM package makes a GEN query to the KEY package that samples a key uniformly at random.

In standard formulations of KEM security, the adversary not only receives an encapsulation, but also the encapsulated key (or a random key, if $b = 1$) as an answer to ENCAP. In our decomposed equivalent formulation, the adversary can access the encapsulated key (or a random key, if $b = 1$) via a GET query to the KEY package (also see page 18, Definition 21 for the IND-CCA KEM game).

The DEM Game. Fig. 1b depicts the decomposed IND-CCA DEM game that also contains a KEY package. Here, the adversary can ask a GEN query to the KEY package which induces the KEY package to sample a uniformly random key that the DEM package obtains via a GET query to the KEY package. Note that in the DEM game, the adversary only has access to the GEN oracle of the KEY package, but neither to SET nor to GET. Moreover, in the DEM game, the adversary can make encryption and decryption queries (see page 18, Definition 22 for the definition of IND-CCA security for DEMs).

KEM-DEM security. Recall that we prove that the KEM-DEM construction is a IND-CCA secure public-key encryption scheme. Using the packages `KEM`, `DEM` and `KEY`, we now write the IND-CCA security game for public-key encryption in a modular way, see Figure 2. In Appendix D, we prove via inlining, that the modular game in Figure 2a, is equivalent to the monolithic IND-CCA game for public-key encryption with secret bit 0 and that the modular game in Figure 2g, is equivalent to the monolithic IND-CCA game for public-key encryption with secret bit 1.

Thus, we first idealize the `KEM` package, then idealize the `DEM` package, and finally de-idealize the `KEM` package again. Technically, this works as follows. Starting from the composition in Fig. 2a, we lengthen the edges of the graph such that the `KEM`⁰ and `KEY` packages are on the right side of a vertical line (see Fig. 2b). Analogously to the first example, we use associativity (and additional rules, explained shortly) to reduce to the security of `KEM` by noticing that the packages on the left side of the vertical line call the packages on the right side of the vertical line, where the latter correspond to the `KEM` security game.

Reasoning on the graph corresponds to reasoning on compositions of packages, defined via the *sequential* operator \circ and the *parallel* composition operator, see Section 2. The lengthening of edges corresponds to inserting forwarding packages, denoted *identity* `ID`. The aforementioned *interchange* rule then allows to formally interpret the vertical line in the graph as a sequential composition of the packages on the left side of the line with the packages on the right side. For a graphical depiction of the identity rule and the interchange rule, see Section 2.2.

After applying the `KEM` assumption (which modifies `KEM`⁰ to `KEM`¹), we contract the graph which, again, corresponds to applying the interchange rule and then removing `ID`s, see Fig. 2c. Via the analogous mechanism, we stretch the graph edges such that the `DEM`⁰ and `KEY` appear on the right side of a vertical line, see Fig. 2d. We apply the `DEM` assumption and then contract the graph to obtain Fig. 2e. By applying the `KEM` assumption again in a similar manner we obtain Fig. 2g as desired.

Contents.

§2 Proof methodology. In this section, we set up the underlying code framework and define sequential and parallel composition. We specify rules to operate on package compositions such as the aforementioned associativity, interchange and identity rules. Those rules enable the graphical interpretation as a call graph which we explain in Section 2.2.

§3 KEY package composition. We introduce keying games (such as the `KEM` game) and keyed games (such as the `DEM` game) which both contain a `KEY` package, introduced in this section. In a single key lemma we prove indistinguishability properties of composed keyed and keying packages. A core argument in the proof of the lemma is that the idealization of the keying game leads to only calling the `GEN` oracle. As keyed games rely on uniformly

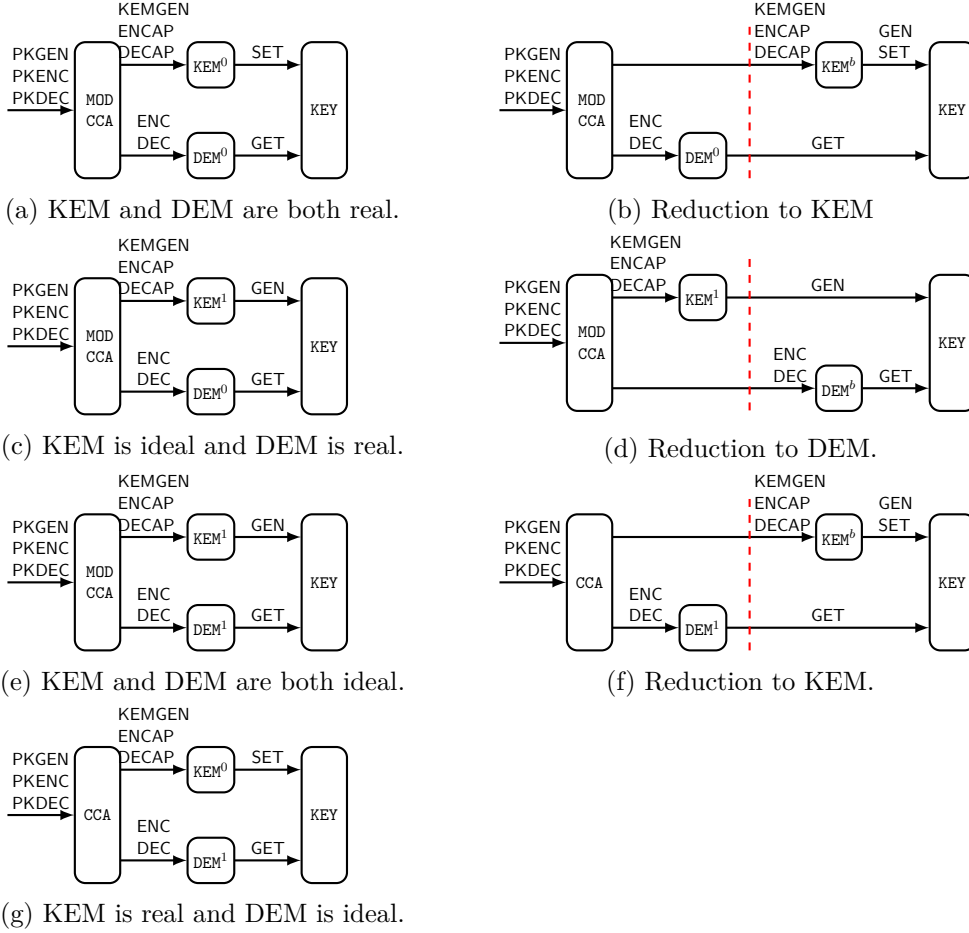


Figure 2: KEM-DEM Proof.

random keys, we model their security formally by inserting an identity package ID_{GEN} that only forwards the GEN oracle. Based on Section 2.2, we maintain a coherent mapping to the graphical notation in which accessible oracles are simply labels on edges.

§4 KEM-DEM. We provide the details of the KEM-DEM construction and proof discussed earlier. In particular, the security reduction is a straightforward application of the single key lemma.

§5 Multi-Instance Packages and Composition. In this section, we generalize to the multi-instance setting and carry out a multi-instance-to-single-instance composition proof. We then build on the multi-instance lemma to obtain multi-instance version of the single key lemma.

Avoiding multi-to-single instance reductions is one of the motivations of composition frameworks (see below). Hence, we see it as a sanity check that our proof methodology captures multi-to-single instance reductions. Note that also in the game-based setting, general multi-instance to single-instance reductions for classes of games have been provided before (see, e.g., Bellare, Boldyreva and Micali [BBM00]).

§6 Composition of forward-secure key exchange. To showcase our key-composition techniques in the multi-instance setting, we re-prove a composition theorem for forward-secure game-based key exchange and arbitrary symmetric-key based protocols such as secure channels. This result was proven in Brzuska, Fischlin, Warinschi, and Williams [BFWW11, Brz13] and becomes a straightforward application of the multiple keys lemma. Our results are closely related to composition results very recently shown in the framework of CryptoVerif [Bla18].

Limitations and Challenges.

Our method considers distinguishing games for *single-stage* adversaries [RSS11], that is, we do not consider games where the adversary is split into separate algorithms whose communications are restricted. Although suitable extensions might exist (e.g., by extending adversaries into packages that can call each other), we chose to restrict our current method to the simpler single-stage setting.

Another apparent restriction is that we encode all security properties via indistinguishability. Search problems such as strong unforgeability can also be encoded via indistinguishability. While the encoding might seem surprising when not used to it, at a second thought, an appropriate encoding of an unforgeability game also simplifies game-hopping: Imagine that we insert an abort condition whenever a message is accepted by verification that was not signed by the signer. This step corresponds to idealizing the verification of the signature scheme so that it only accepts messages that were actually signed before.¹

A challenge that all cryptographic works on real-world protocols face is to decompose a protocol that does not inherently have a modular structure into cryptographic building blocks. As demonstrated by [KPW13, KMO⁺15, BFK⁺14] this can be done even for archaic protocols such as TLS. Our method is influenced by the insights of the miTLS project to allow for the necessary flexibility.

Related Techniques.

Our approach is inspired by important conceptual works from cryptography and programming language. In particular, we would like to acknowledge the influences of Canetti’s universal composability framework (UC) [Can01], Renner’s and Maurer’s work on random systems and abstract cryptography [Mau02, MR11], process algebras, such as the π -calculus

¹CryptoVerif [Bla08] also encodes authentication properties as indistinguishability.

of Milner, Parrow, and Walker [MPW92], and type-based verification frameworks used, e.g., to verify the TLS protocol [BFK⁺13]. We now discuss these influences in detail.

Cryptographic Proof Frameworks. Composable proofs in the pen-and-paper world as pioneered by Backes, Pfitzmann, Waidner and by Canetti have a long history full of rich ideas [BPW04, Can01, KT13, MQU07, HS11, MT13, HS15, Wik16], such as considering an environment that cannot distinguish a real protocol from an ideal variant with strong security guarantees.

Likewise, Maurer’s and Renner’s work on random systems, abstract cryptography and constructive cryptography [Mau02, Mau10, MR11, Mau11] inspired and encouraged our view that a more abstract and algebraic approach to cryptographic proofs is possible and desirable. Several of our concepts have close constructive cryptography analogues: for instance, our use of associativity in this paper is similar to composition-order independence in Maurer’s frameworks [Mau11]. Sequential and parallel composition also appears in cryptographic algebras. An ambitious expression of the idea is found in [MR11, Section 6.2]. Abstract cryptography has an associativity law and neutral element for sequential composition and an interchange law for parallel composition. The same line of work [MR11, Mau11] introduces a distinguishing advantage between composed systems and makes use of transformations that move part of the system being considered into and out of the distinguisher.

Our focus is not on definitions but on writing game-based security proofs. As such we are also influenced by works on game-based composition, e.g., Brzuska, Fischlin, Warinschi, and Williams [BFWW11]. We aim to facilitate security proofs for full-fledged standardized protocols [JKSS12, KPW13, DFGS15, CCD⁺17]. Such proofs typically involve large reductions relating a complex monolithic game to diverse cryptographic assumptions through an intricate simulation of the protocol.

Language-Based Security and Cryptography. Algebraic reasoning is at the core of process calculi such as the π -calculus by Milner, Parrow and Walker [MPW92]. They focus on concurrency with non-determinism, which is also adequate for symbolic reasoning about security protocols. Subsequently, probabilistic process algebras have been used to reason computationally about protocols, e.g., in the work of Mitchell, Ramanathan, Scedrov, and Teague [MRST06] and the *computational indistinguishability logic* (CIL) of Barthe, Crespo, Lakhnech and Schmidt [BDKL10]. Packages can be seen as an improvement of CIL oracle systems, with oracle visibility and associativity corresponding to the context rules of CIL.

Monadic composition, a generalisation of function composition to effectful programs, is a central principle of functional languages such as Haskell, F^\sharp , and F^* [Jon03, SGC12, SHK⁺16]. Associativity is also used by Mike Rosulek in his rich undergraduate textbook draft *The Joy of Cryptography* to make the cryptographic reduction methodology accessible to undergraduate students with no background in complexity theory [Ros18]. Our concept

of packages is inspired by module systems in programming languages such as F^\sharp , OCaml, SML (see e.g. Tofte [Tof96]). Our oracles similarly define a public interface for calling functions that may share private state.

Existing techniques for overcoming the crisis of rigour in provable security as formalised by Bellare and Rogaway [BR06] and mechanised in EasyCrypt [BGHB11] have focused on the most intricate aspects of proofs. EasyCrypt supports a rich module system similar to the ones found in functional programming languages [BCLS15] (including parametric modules, i.e. functors), but it has not yet been used to simplify reasoning about large reductions in standardized protocols.

The closest to our idea of package-based reductions is the modular code structure of miTLS, an cryptographically verified implementation of TLS coded in F^* [FKS11, BFK⁺13, BFK⁺14, DFK⁺17]. Fournet, Kohlweiss and Strub [FKS11] show that code-based game rewriting can be conducted on actual implementation code, one module at a time, with the rest of the program becoming the reduction for distinguishing the *ideal* from the *real* version of the module. Packages are simpler than F^* modules, with interfaces consisting just of sets of oracle names, whereas F^* provides a rich type system for specifying module interfaces and verifying their implementations.

Our method draws from both formal language techniques and pen-and-paper approaches for cryptographic proofs. We see facilitating the flow of information between the two research communities as an important contribution of our work. In this paper, we use pseudo-code, treating the concrete syntax and semantics of our language as a parameter. This simplifies our presentation and makes it more accessible to the cryptographic community. Our method can be instantiated either purely as a pen-and-paper method or via using a full-fledged programming language, equipped with a formal syntax and operational semantics. The latter might also allow the development of tools for writing games and automating their proofs. Indeed, this is already the subject of recent follow-up work [DKO21, AHR⁺21].

2 Proof Methodology

As discussed in the introduction, we suggest to work with *pseudo-code* instead of Turing machines as a model of computation and thus, this section will start by providing a definition of code. We then continue to define functions and function calls (to probabilistic and stateful functions), also known as oracles and oracle calls in the cryptographic literature. We will then collect several such functions (oracles) into a package, and when the package itself does not make any function calls, we call a package *closed* or a *game*. We then define sequential composition of 2 packages, where the first package calls functions (oracles) defined by the second package. Moreover, we define parallel composition which allows to take the functions defined by two packages and to take their union.

Then, we move to more advanced packages and algebraic rules that allow to implement

the “moving to the right” operation that we hinted at in the introduction.

2.1 Composing Oracle Definitions

While we advocate to work with pseudo-code, we do not define a particular language, but rather *parametrize* our method by a language for writing algorithms, games, and adversaries. We specify below the properties of the syntax and semantics of any language capable of instantiating our approach. We first describe our pseudo-code and give a probabilistic semantics to whole programs, then we explain our use of functions for composing code.

Definition 2 (Pseudo-Code). *We assume given sets of values v, \dots , local variables x, y, \dots , expressions e , state variables a, T (uppercase denotes tables), \dots , and commands c .*

Values provide support for booleans, numbers, and bitstrings. Expressions provide support for operations on them. Expressions may use local variables, but not state variables.

*Commands include local-variable assignments $x \leftarrow a$ and $x \leftarrow e$, sampling from a distribution $x \leftarrow_s \mathcal{D}$, state updates $a \leftarrow e$ and $T[x] \leftarrow e$, sequential compositions $c; c'$, and **return** e for returning the value of e . We write $\text{fv}(c)$ for the state variables accessed in c . We assume given default initial values for all state variables, e.g. $T \leftarrow \emptyset$.*

We write $\Pr[v \leftarrow c]$ for the probability that command c returns v . (We only consider programs that always terminate.) We assume this probability is stable under injective renamings of local variables and state variables.

For brevity, we often write commands with expressions that depend on the current state, as a shorthand for using intermediate local variables for reading the state, e.g. we write $T[x] \leftarrow T[x] + 1$ as a shorthand for $t \leftarrow T[x]; T[x] \leftarrow t + 1$.

Definition 3 (Functions). *We assume given a set of names f, \dots for functions. We let \mathcal{O} range over function definitions of the form $f(x) \mapsto c$. and write $\Omega = \{f_i(x_i) \mapsto c_i\}_{i=1..n}$ for a set of n function definitions with distinct function names. We write $\text{dom}(\Omega)$ for the set of names $\{f_1, \dots, f_n\}$ defined in Ω and $\Sigma(\Omega)$ for the set of state variables accessed in their code.*

We extend commands with function calls, written $y \leftarrow f(e)$. We write $\text{fn}(c)$ for the set of function names called in c , and similarly define $\text{fn}(O)$ and $\text{fn}(\Omega)$. We say that a term is closed when this set is empty.

We interpret all function calls by inlining, as follows: given the definition $f(x) \mapsto c; \text{return } e'$, the call $y \leftarrow f(e)$ is replaced with $c; y \leftarrow e'$ after replacing x with e in the function body. We write $\text{inline}(c, \Omega)$ for the code obtained by inlining all calls to the functions f_1, \dots, f_n defined by Ω in the command c . Similarly, we write $\text{inline}(\Omega', \Omega)$ for the set of definitions obtained by inlining all calls to functions in Ω into the code of the definitions of Ω' .

We consider function definitions up to injective renamings of their local variables.

Packages. We now introduce the general definition of *packages* as collections of oracles that subsume adversaries, games and reductions. Packages are sets of oracles Ω s defined above. Intuitively, we will treat the state variables of their oracles as private to the package, i.e., the rest of the code only get oracle access. Looking ahead to the composition of packages we endow each package with an *output* interface consisting of the oracles names that it defines and an *input* interface consisting of the oracles names that it queries.

Definition 4 (Packages). *A package \mathbb{M} is a set of function definitions Ω (its oracles) up to injective renamings of its state variables $\Sigma(\Omega)$.*

We write $\text{in}(\mathbb{M}) = \text{fn}(\Omega)$ for its input interface and $\text{out}(\mathbb{M}) = \text{dom}(\Omega)$ for its output interface.

We disallow internal calls to prevent recursion. Technically, the disallowing of internal calls is captured (a) by the input interface of a package, since this input provides all oracles that are called by the oracles in Ω , and (b) by the Def. 5 of sequential composition that specifies that oracle calls are instantiated by the oracles of *another* package.

We often consider families of oracles \mathcal{O}^Π and packages \mathbb{M}^Π parametrized by Π , treating parameters as symbolic values in their code. We usually omit parameters and refer to oracles and packages by their name, unless context requires further clarification. In particular, we write $\text{in}(\mathbb{M}^\Pi)$ only if the input interface differs for different parameters; $\text{out}(\mathbb{M})$ never depends on the parameters.

Package composition. We say that \mathbb{M} *matches* the output interface of \mathbb{M}' iff $\text{in}(\mathbb{M}) \subseteq \text{out}(\mathbb{M}')$. When composing two matching packages $\mathbb{M} \circ \mathbb{M}'$, we *inline* the code of all oracles of \mathbb{M}' called by oracles in \mathbb{M} , as specified in Definition 3.

Definition 5 (Sequential Composition). *Given two packages \mathbb{M} with oracles Ω and \mathbb{M}' with oracles Ω' such that \mathbb{M} matches \mathbb{M}' and $\Sigma(\Omega) \cap \Sigma(\Omega') = \emptyset$, their sequential composition $\mathbb{M} \circ \mathbb{M}'$ has oracles $\text{inline}(\Omega, \Omega')$.*

Thus, we have $\text{out}(\mathbb{M} \circ \mathbb{M}') = \text{out}(\mathbb{M})$ and $\text{in}(\mathbb{M} \circ \mathbb{M}') = \text{in}(\mathbb{M}')$.

When describing a package composition, one cannot use the same package twice, e.g., it is not possible to have compositions such as $(\mathbb{M} \circ \mathbb{M}' \circ \mathbb{M})$. Note that this is a fundamental restriction, since it is unclear how to define the state of such a composition, since there would be copies of pointers to the same state (a.k.a. aliases).

Lemma 6 (Associativity). *Let $\mathbb{M}_0, \mathbb{M}_1, \mathbb{M}_2$ such that $\text{in}(\mathbb{M}_0) \subseteq \text{out}(\mathbb{M}_1)$ and $\text{in}(\mathbb{M}_1) \subseteq \text{out}(\mathbb{M}_2)$. We have $(\mathbb{M}_0 \circ \mathbb{M}_1) \circ \mathbb{M}_2 \stackrel{\text{code}}{\equiv} \mathbb{M}_0 \circ (\mathbb{M}_1 \circ \mathbb{M}_2)$.*

Proof outline. We rename the local variables and state variables of the three packages to prevent clashes, then unfold the definition of sequential compositions by inlining, and rely on the associativity of their substitutions of function code for function calls.

We now define parallel composition, which is essentially a disjoint union operator that takes two packages and builds a new package that implements both of them in parallel. It is important to note that only the output interfaces of M and M' need to be disjoint, while they can potentially share input oracles. This feature allows for parallel composition of several packages that use the same input interface.

Definition 7 (Parallel Composition). *Given two packages M with oracles Ω and M' with oracles Ω' such that $\text{out}(M) \cap \text{out}(M') = \emptyset$ and $\Sigma(\Omega) \cap \Sigma(\Omega') = \emptyset$, their parallel composition $\frac{M}{M'} \text{ (alternatively } (M|M'))$ has oracles $\Omega \uplus \Omega'$. Thus, $\text{out}(\frac{M}{M'}) = \text{out}(M) \uplus \text{out}(M')$ and $\text{in}(\frac{M}{M'}) = \text{in}(M) \cup \text{in}(M')$.*

(This composition may require preliminary renamings to prevent clashes between the state variables of M and M' .)

Lemma 8. *Parallel composition is commutative and associative.*

The proof of these properties directly follows from our definition of packages. Associativity enables us to write n -ary parallel compositions of packages. Next, we show that sequential composition distributes over parallel composition. (The conditions in the lemma guarantee that the statement is well defined.)

Lemma 9 (Interchange). *For all packages M_0, M_1, M'_0, M'_1 , if $\text{out}(M_0) \cap \text{out}(M_1) = \emptyset$, $\text{out}(M'_0) \cap \text{out}(M'_1) = \emptyset$, $\text{out}(M_0) \subseteq \text{in}(M'_0)$ and $\text{out}(M_1) \subseteq \text{in}(M'_1)$, then*

$$\frac{M_0}{M_1} \circ \frac{M'_0}{M'_1} \stackrel{\text{code}}{=} \frac{M_0 \circ M'_0}{M_1 \circ M'_1}.$$

Proof outline. The code equality relies on the property that function-call inlining applies pointwise to each of the oracle definitions in the 3 sequential compositions above.

Identity packages. Some proofs and definitions make one or more oracles of a package unavailable to the adversary, which is captured by sequential composition with a package that forwards a subset of their oracle calls:

Definition 10 (Identity Packages). *The identity package ID_X for the names X has oracles $\{f(x) \mapsto r \leftarrow f(x); \text{return } r\}_{f \in X}$.*

Hence, for $X \subseteq \text{out}(M)$, the package $\text{ID}_X \circ M$ behaves as M after deleting the definitions of oracles outside X . In particular, the next lemma gives some identity compositions that do not affect a package.

Lemma 11 (Identity Rules). *For all packages M , $M \stackrel{\text{code}}{=} \text{ID}_{\text{out}(M)} \circ M$ and $M \stackrel{\text{code}}{=} M \circ \text{ID}_{\text{in}(M)}$.*

Proof outline. By definition of sequential composition and basic properties of substitutions, we obtain the following from $ID_{out(M)} \circ M$:

We substitute ‘ $f(x) \mapsto c; \mathbf{return} r$ ’ in ‘ $f(x) \mapsto r \leftarrow f(x); \mathbf{return} r$ ’ and yield ‘ $f(x) \mapsto c; r \leftarrow r; \mathbf{return} r$ ’ which is equivalent to ‘ $f(x) \mapsto c; \mathbf{return} r$ ’. Analogously, for $M \circ ID_{in(M)}$:

We substitute ‘ $f(x) \mapsto r \leftarrow f(x); \mathbf{return} r$ ’ in ‘ $r' \leftarrow f(x)$ ’ and yield ‘ $r \leftarrow f(x); r' \leftarrow r$ ’ which is equivalent to ‘ $r' \leftarrow f(x)$ ’. \square


2.2 Graphical Representation of Package Composition

Writing fully-precise package compositions can be tedious. Recall the KEM-DEM proof of Fig. 2; the step from (a) to (b) corresponds to applying a mix of interchange and identity rules:

$$CCA \circ \left(\frac{KEM^0}{DEM^0} \circ KEY \right) \stackrel{\text{code}}{\equiv} CCA \circ \left(\frac{ID \circ KEM^0}{DEM^0 \circ ID} \circ KEY \right) \stackrel{\text{code}}{\equiv} CCA \circ \left(\left(\frac{ID}{DEM^0} \circ \frac{KEM^0}{ID} \right) \circ KEY \right)$$


Instead of writing such steps explicitly, we propose a graphical representation of package composition that allows us to reason about compositions “up to” applications of the interchange, identity and associativity rules.

From terms to graphs. Identity packages ID_S map to edges, one for each oracle in the set S . Other packages map to a node labelled with the package name. Each output oracle of the package maps to an incoming edge of the node, labelled with the oracle name. Similarly, input oracles map to outgoing edges.




Sequential composition $A \circ B$ simply consists of merging the outgoing edges of A with the incoming edges of B with the same label. Note that in this process, some of the incoming edges of B may be dropped, i.e. A may not use all of the oracles exported by B .


The parallel composition of A and B is simply the union of the graphs constructed from A and B . By definition of parallel composition, $out(A) \cap out(B) = \emptyset$, while *input* oracles may be used both by A and B . We merge shared input edges (i.e. unconnected outgoing edges) in the resulting graph to capture this sharing.



From graphs to terms. By inductive application of the above 3 rules, one can construct a graph representing any term. However, some information is lost in the process: most importantly, the order in which



sequential and parallel compositions are applied. For instance, consider the left-hand side and right-hand side of the interchange rule. Both terms map to the



same graph. This is by design, as we intend to represent terms modulo interchange. By drawing explicit boxes around parallel and sequential compositions, it is possible to ensure that a graph can be interpreted unambiguously as a term. For instance, the figure on the right shows how to depict the interchange rule on graphs with boxes.

2.3 Games and Adversaries

Games A *game* is a package with an empty input interface. We model security properties of a cryptographic scheme as indistinguishability between a *pair* of games, usually parameterized by a bit $b \in \{0, 1\}$ (which is equivalent to a single game that draws a bit and then runs one of the two games at random.).

Adversaries. An adversary \mathcal{A} is a package with output interface $\{\text{run}\}$ that returns a bit 0 or 1. We model the adversary as a package whose input interface is equal to the set of names of the oracles of the game that the adversary is meant to interact with.

Next, we define games and adversaries such that their composition $\mathcal{A} \circ \mathbf{G}$ be a closed package of the form $\mathbf{R} = \{\text{run}() \mapsto c; \text{return } g\}$.

Since Definition 2 defines our probabilistic semantics only on commands, we first extend it to such closed packages, defining $\Pr[1 \leftarrow \mathbf{R}]$ as $\Pr[1 \leftarrow c; \text{return } g]$. (The command $c; \text{return } g$ is the ‘top-level’ code $g \leftarrow \text{run}(); \text{return } g$ after inlining the definition of $\text{run}()$.)

Definition 12 (Games). A *game* is a package \mathbf{G} such that $\text{in}(\mathbf{G}) = \emptyset$. An *adversary against* \mathbf{G} is a package \mathcal{A} such that $\text{in}(\mathcal{A}) = \text{out}(\mathbf{G})$ and $\text{out}(\mathcal{A}) = \{\text{run}\}$. A *game pair* consists of two games \mathbf{G}^0 and \mathbf{G}^1 that define the same oracles: $\text{out}(\mathbf{G}^0) = \text{out}(\mathbf{G}^1)$. Naturally, a game \mathbf{G}^b with a binary parameter b defines a game pair. We thus use the two notions interchangeably.

We now define distinguishing advantages. Note that we operate in the concrete security setting as it is more adequate for practice-oriented cryptography and therefore only define

advantages rather than security in line with the critique of Rogaway [Rog06], and Bernstein and Lange [BL13]. Our ideas can be transferred analogously to the asymptotic setting.

Definition 13 (Distinguishing Advantage). *The advantage of an adversary \mathcal{A} against a game pair \mathbf{G} is*

$$\epsilon_{\mathbf{G}}(\mathcal{A}) = \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \mathbf{G}^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \mathbf{G}^1 \right] \right|.$$

In the rest of the paper, we may refer to the advantage function $\epsilon_{\mathbf{G}}$ in this definition by writing $\mathbf{G}^0 \stackrel{\epsilon_{\mathbf{G}}}{\approx} \mathbf{G}^1$; and we write $\mathbf{G}^0 \stackrel{\text{perf}}{\equiv} \mathbf{G}^1$ if $\epsilon_{\mathbf{G}} = 0$. For two packages (not only for games), \mathbf{M} and \mathbf{N} , we write $\mathbf{M} \stackrel{\text{code}}{\equiv} \mathbf{N}$ if they provide the same function definitions Ω up to injective renamings of state variables, after inlining (in case \mathbf{M} and/or \mathbf{N} was specified as a composition of packages). Note that if \mathbf{M} and \mathbf{N} are games, then $\mathbf{M} \stackrel{\text{code}}{\equiv} \mathbf{N}$ implies $\mathbf{M} \stackrel{\text{perf}}{\equiv} \mathbf{N}$. As an example of advantage, we restate below the usual triangular equality for three games with the same oracles.

Lemma 14 (Triangle Inequality). *Let \mathbf{F} , \mathbf{G} and \mathbf{H} be games such that $\text{out}(\mathbf{F}) = \text{out}(\mathbf{G}) = \text{out}(\mathbf{H})$. If $\mathbf{F} \stackrel{\epsilon_1}{\approx} \mathbf{G}$, $\mathbf{G} \stackrel{\epsilon_2}{\approx} \mathbf{H}$, and $\mathbf{F} \stackrel{\epsilon_3}{\approx} \mathbf{H}$, then $\epsilon_3 \leq \epsilon_1 + \epsilon_2$.*

The triangle inequality helps to sum up game-hops. Many game-hops will exploit simple associativity, as the following lemma illustrates.

Lemma 15 (Reduction). *Let \mathbf{G} be a game pair and let \mathbf{M} be a package such that $\text{in}(\mathbf{M}) \subseteq \text{out}(\mathbf{G})$. Let \mathcal{A} be an adversary that matches the output interface of \mathbf{M} , then for both $b \in \{0, 1\}$, the adversary $\mathcal{D} := \mathcal{A} \circ \mathbf{M}$ satisfies*

$$\Pr \left[1 \leftarrow \mathcal{A} \circ (\mathbf{M} \circ \mathbf{G}^b) \right] = \Pr \left[1 \leftarrow \mathcal{D} \circ \mathbf{G}^b \right].$$

As a corollary, we obtain $\mathcal{A} \circ \mathbf{M} \circ \mathbf{G}^0 \stackrel{\epsilon(\mathcal{A})}{\approx} \mathcal{A} \circ \mathbf{M} \circ \mathbf{G}^1$ for $\epsilon(\mathcal{A}) = \epsilon_{\mathbf{G}}(\mathcal{A} \circ \mathbf{M})$.

Proof. The proof follows by associativity of sequential composition, i.e., Lemma 6 yields $\mathcal{A} \circ (\mathbf{M} \circ \mathbf{G}^b) \stackrel{\text{code}}{\equiv} (\mathcal{A} \circ \mathbf{M}) \circ \mathbf{G}^b \stackrel{\text{code}}{\equiv} \mathcal{D} \circ \mathbf{G}^b$. \square

3 KEY Package Composition

Many cryptographic constructions emerge as compositions of two cryptographic building blocks: The first building block generates the (symmetric) key(s) and the second building block uses the (symmetric) key(s). In the introduction, we already discussed the popular composition of key encapsulation mechanisms (KEM) with a deterministic encryption mechanism (DEM). Likewise, complex protocols such as TLS first execute a key exchange protocol to generate symmetric keys for a secure channel. In composition proofs, the keying

building block and the keyed building block share the (symmetric) key(s). To capture this shared state, we introduce a key package KEY^λ that holds a single key k of length λ . (We handle multiple keys in Section 5.)

Definition 16 (Key Package). *For $\lambda \in \mathbb{N}$, KEY^λ is the package that defines the three oracles below, i.e., $\text{out}(\text{KEY}^\lambda) = \{\text{GEN}, \text{SET}, \text{GET}\}$.*

$\text{GEN}()$	$\text{SET}(k')$	$\text{GET}()$
$\text{assert } k = \perp$	$\text{assert } k = \perp$	$\text{assert } k \neq \perp$
$k \leftarrow_{\$} \{0, 1\}^\lambda$	$k \leftarrow k'$	$\text{return } k$

Hence, this package encapsulates the state variable k , initialized (once) by calling either GEN or SET , then accessed by calling GET . This usage restriction is captured using **asserts**, and all our definitions and theorems apply only to code that never violates assertions.

Definition 17 (Keying Games). *A keying game K is a game composed of a core keying package CK and the key package as follows:*

$$K^{b,\lambda} \stackrel{\text{code}}{\equiv} \frac{\text{CK}^{b,\lambda}}{\text{ID}_{\{\text{GET}\}}} \circ \text{KEY}^\lambda.$$

where $b \in \{0, 1\}$, $\text{in}(\text{CK}^{0,\lambda}) = \{\text{SET}\}$, and $\text{in}(\text{CK}^{1,\lambda}) = \{\text{GEN}\}$.

Definition 18 (Keyed Games). *A keyed game D is a game composed of a core keyed package CD and the key package as follows:*

$$D^{b,\lambda} \stackrel{\text{code}}{\equiv} \frac{\text{ID}_{\{\text{GEN}\}}}{\text{CD}^{b,\lambda}} \circ \text{KEY}^\lambda.$$

where $b \in \{0, 1\}$ and $\text{in}(\text{CD}^{b,\lambda}) = \{\text{GET}\}$.

Lemma 19 (Single Key). *Keying games K and keyed games D are compatible when they have the same key length λ and they define disjoint oracles, i.e., $\text{out}(K) \cap \text{out}(D) = \emptyset$. For all compatible keying and keyed games, with the notations above, we have*

$$(a) \quad \frac{\text{CK}^0}{\text{CD}^0} \circ \text{KEY}^\lambda \approx_a \frac{\text{CK}^1}{\text{CD}^1} \circ \text{KEY}^\lambda, \quad (b) \quad \frac{\text{CK}^0}{\text{CD}^0} \circ \text{KEY}^\lambda \approx_b \frac{\text{CK}^0}{\text{CD}^1} \circ \text{KEY}^\lambda,$$

where, for all adversaries \mathcal{A} ,

$$\begin{aligned} \epsilon_a(\mathcal{A}) &\leq \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^0} \right) + \epsilon_D \left(\mathcal{A} \circ \frac{\text{CK}^1}{\text{ID}_{\text{out}(\text{CD})}} \right), \\ \epsilon_b(\mathcal{A}) &\leq \epsilon_a(\mathcal{A}) + \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^1} \right). \end{aligned}$$




Figure 3: Reduction to the keying game (left) and the keyed game (right).

Version (b) of the single key lemma is needed when proving theorems in which the ideal version of the top-level game exposes cryptographic material that depends on the concrete key, e.g. an encryption of zero.

Proof. Fig. 3 gives the proof outline using graphs: To show (a), we first idealize the core keying package, switching from SET to GEN (left); and then we idealize the core keyed package (Fig. 3, right). To show (b), we also de-idealize the core keying package, switching back from GEN to SET (left).

We give a more detailed proof below, using the algebraic rules of Section 2 to rewrite packages in order to apply Def. 17 and 18.

(1) Idealizing the core keying package. The first intermediate goal is to bring the package into a shape where we can use Def. 17 to change CK^0 into CK^1 . Below, for all adversaries \mathcal{A} , we have $\epsilon_1(\mathcal{A}) = \epsilon_K \left(\mathcal{A} \circ \frac{ID_{out(CK)}}{CD^0} \right)$.

$$\begin{aligned} \frac{CK^0}{CD^0} \circ KEY^\lambda &\stackrel{\text{code}}{\equiv} \frac{ID_{out(CK)}}{CD^0} \circ \frac{CK^0}{ID_{\{GET\}}} \circ KEY^\lambda \quad (\text{identity \& interchange}) \\ &\stackrel{\epsilon_1}{\approx} \frac{ID_{out(CK)}}{CD^0} \circ \frac{CK^1}{ID_{\{GET\}}} \circ KEY^\lambda \stackrel{\text{code}}{\equiv} \frac{CK^1}{CD^0} \circ KEY^\lambda \end{aligned}$$

(2) Idealizing the core keyed package. As a second step, we want to use Def. 18 to move from CD^0 to CD^1 and thus need to make $ID_{\{GEN\}}$ appear. Note that we can use $ID_{\{GEN\}}$ because $\{GEN\}$ is equal to the input interface of CK^1 . This was not possible before idealizing to CK^1 , since $in(CK^0) = \{SET\}$. Below, for all adversaries \mathcal{A} , we have $\epsilon_2(\mathcal{A}) = \epsilon_D \left(\mathcal{A} \circ \frac{CK^1}{ID_{out(CD)}} \right)$.

$$\begin{aligned} \frac{CK^1}{CD^0} \circ KEY^\lambda &\stackrel{\text{code}}{\equiv} \frac{CK^1}{ID_{out(CD)}} \circ \frac{ID_{\{GEN\}}}{CD^0} \circ KEY^\lambda \quad (\text{identity \& interchange}) \\ &\stackrel{\epsilon_2}{\approx} \frac{CK^1}{ID_{out(CD)}} \circ \frac{ID_{\{GEN\}}}{CD^1} \circ KEY^\lambda \stackrel{\text{code}}{\equiv} \frac{CK^1}{CD^1} \circ KEY^\lambda \end{aligned}$$

(3) De-idealizing the core keying package. Finally, we move back from CK^1 to CK^0 , taking the inverse steps of idealizing the core keying package. We obtain $\epsilon_3(\mathcal{A}) = \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^1} \right)$. Below, for all adversaries \mathcal{A} , we have $\epsilon_3(\mathcal{A}) = \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^1} \right)$.

$$\begin{aligned} \frac{\text{CK}^1}{\text{CD}^1} \circ \text{KEY}^\lambda &\stackrel{\text{code}}{\equiv} \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^1} \circ \frac{\text{CK}^1}{\text{ID}_{\{\text{GET}\}}} \circ \text{KEY}^\lambda \quad (\text{identity \& interchange}) \\ &\stackrel{\epsilon_3}{\approx} \frac{\text{ID}_{\text{out}(\text{CK})}}{\text{CD}^1} \circ \frac{\text{CK}^0}{\text{ID}_{\{\text{GET}\}}} \circ \text{KEY}^\lambda \stackrel{\text{code}}{\equiv} \frac{\text{CK}^0}{\text{CD}^1} \circ \text{KEY}^\lambda \end{aligned}$$

□

4 KEM-DEMs

Cramer and Shoup [CS03, §7] show that composing a CCA-secure key encapsulation mechanism (KEM) and a CCA-secure data encapsulation mechanism (DEM) yields a CCA-secure public-key encryption (PKE). Using the **KEY** package composition introduced in Section 3, we give a new formulation of their KEM-DEM proof.

Schemes are function definitions that do not employ state variables. We write M^β for a package calling functions of the scheme β in its parameters. Formally, for a package M with oracles Ω , M^β denotes the package with oracles $\text{inline}(\Omega, \beta)$.

We denote the set of functions defined by a PKE scheme with ciphertext expansion $\text{clen}(|m|)$ by $\zeta = \{\text{kgen}, \text{enc}, \text{clen}, \text{dec}\}$. We denote the set of functions of a DEM scheme with key length λ and ciphertext expansion $\text{clen}(|m|)$ by $\theta = \{\lambda, \text{enc}, \text{clen}, \text{dec}\}$, where we recall that enc is a deterministic, one-time encryption algorithm. We prepend function names by ζ and θ for disambiguation. We denote a KEM scheme with output key length λ and encapsulation length elen by $\eta = \{\text{kgen}, \text{encap}, \text{elen}, \text{decap}, \lambda\}$, where kgen produces a key pair (pk, sk) , $\text{encap}(pk)$ generates a symmetric key k of length $\eta \cdot \lambda$ and a key encapsulation c of length $\eta \cdot \text{elen}$, while $\text{decap}(sk, c)$ given sk and an encapsulation c returns a key k . For all three schemes, we consider perfect correctness. Throughout this section, we consider a single symmetric-key length λ that corresponds to the length of the symmetric key used by the DEM scheme as well as the length of the symmetric key produced by the encapsulation mechanism $\eta \cdot \text{encap}$. We now turn to the security notions which are IND-CCA security notions for all three primitives.

Definition 20 (PKE-CCA Security). *Let ζ be a PKE-scheme. We define its IND-CCA advantage $\epsilon_{\text{PKE-CCA}}^\zeta$, where $\text{PKE-CCA}^{b,\zeta}$ defines the following oracles, i.e., $\text{out}(\text{PKE-CCA}^\zeta) = \{\text{PKGEN}, \text{PKENC}, \text{PKDEC}\}$.*

PKGEN()	PKENC(m)	PKDEC(c')
assert $sk = \perp$ $pk, sk \leftarrow_{\S} \zeta.kgen()$ return pk	assert $pk \neq \perp$ assert $c = \perp$ if $b = 0$ then $c \leftarrow_{\S} \zeta.enc(pk, m)$ else $c \leftarrow_{\S} \zeta.enc(pk, 0^{ m })$ return c	assert $sk \neq \perp$ assert $c' \neq c$ $m \leftarrow \zeta.dec(sk, c')$ return m

We model the KEM as a keying and the DEM as a keyed package. We will use the KEY^λ package as specified in Def. 16. Note that we additionally require that encapsulations are indistinguishable from random.

Definition 21 (KEM-CCA Security). *Let η be a KEM. We define its IND-CCA advantage $\epsilon_{\text{KEM-CCA}}^\eta$ using a keying game whose core keying package $\text{KEM}^{b,\eta}$ defines the following oracles, so that $\text{out}(\text{KEM-CCA}^\eta) = \{\text{KEMGEN}, \text{ENCAP}, \text{DECAP}, \text{GET}\}$:*

KEMGEN()	ENCAP()	DECAP(c')
assert $sk = \perp$ $pk, sk \leftarrow_{\S} \eta.kgen()$ return pk	assert $pk \neq \perp$ assert $c = \perp$ if $b = 0$ then $k, c \leftarrow_{\S} \eta.encap(pk)$ $\text{SET}(k)$ else $k, c \leftarrow_{\S} \eta.encap(pk)$ $\text{GEN}()$ return c	assert $sk \neq \perp$ assert $c' \neq c$ $k \leftarrow \eta.decrap(sk, c')$ return k

Note that the adversary queries GET to obtain the challenge key. Encoding the standard KEM notion in this way enables the following algebraic reasoning:

$$\text{KEM-CCA}^{0,\eta} \stackrel{\text{code}}{\equiv} \frac{\text{KEM}^{0,\eta}}{\text{ID}_{\{\text{GET}\}}} \circ \text{KEY}^{\eta,\lambda} \stackrel{\epsilon_{\text{KEM-CCA}}^\eta}{\approx} \frac{\text{KEM}^{1,\eta}}{\text{ID}_{\{\text{GET}\}}} \circ \text{KEY}^{\eta,\lambda} \stackrel{\text{code}}{\equiv} \text{KEM-CCA}^{1,\eta}$$

Definition 22 (DEM-CCA Security). *Let θ be a DEM. We define its IND-CCA advantage $\epsilon_{\text{DEM-CCA}}^\theta$ using a keying game with output interface $\text{out}(\text{DEM-CCA}^\theta) = \{\text{GEN}, \text{ENC}, \text{DEC}\}$, where the oracles of the core keyed packages $\text{DEM}^{b,\theta}$ are defined as follows:*

ENC(m)	DEC(c')
assert $c = \perp$	assert $c \neq c'$
$k \leftarrow \text{GET}()$	$k \leftarrow \text{GET}()$
if $b = 0$ then	$m \leftarrow \theta.\text{dec}(k, c')$
$c \leftarrow \theta.\text{enc}(k, m)$	return m
else	
$c \leftarrow \theta.\text{enc}(k, 0^{ m })$	
return c	

Note that DEM security justifies the following equational reasoning

$$\text{DEM-CCA}^{0,\theta} \stackrel{\text{code}}{\equiv} \frac{\text{DEM}^{0,\theta}}{\text{ID}_{\{\text{GEN}\}}} \circ \text{KEY}^{\theta,\lambda} \stackrel{\epsilon_{\text{DEM-CCA}}^\theta}{\approx} \frac{\text{DEM}^{1,\theta}}{\text{ID}_{\{\text{GEN}\}}} \circ \text{KEY}^{\theta,\lambda} \stackrel{\text{code}}{\equiv} \text{DEM-CCA}^{1,\theta}$$

4.1 Composition and Proof

We prove that the PKE scheme obtained by composing a KEM-CCA secure KEM and a DEM-CCA secure DEM is PKE-CCA secure.

Construction 23 (KEM-DEM Construction). *Let η be a KEM and θ be a DEM. We define the PKE scheme ζ with ciphertext expansion $\zeta.\text{clen}(\ell) = \eta.\text{elen} + \theta.\text{clen}(\ell)$ as follows:*

$\zeta.\text{ngen}()$	$\zeta.\text{enc}(pk, m)$	$\zeta.\text{dec}(sk, c)$
return $\eta.\text{gen}()$	$k, c_1 \leftarrow \eta.\text{encap}(pk)$	$c_1 c_2 \leftarrow c$
	$c_2 \leftarrow \theta.\text{enc}(k, m)$	$k \leftarrow \eta.\text{decap}(sk, c_1)$
	return $c_1 c_2$	$m \leftarrow \theta.\text{dec}(k, c_2)$
		return m

Theorem 24 (PKE Security of the KEM-DEM Construction). *Let ζ be the PKE scheme in Construction 23. For adversaries \mathcal{A} , we have that*

$$\begin{aligned} \epsilon_{\text{PKE-CCA}}^\zeta(\mathcal{A}) &\leq \epsilon_{\text{KEM-CCA}}^\eta \left(\mathcal{A} \circ \text{MOD-CCA} \circ \frac{\text{ID}_{\text{out}(\text{KEM}^\eta)}}{\text{DEM}^{0,\theta}} \right) + \\ &\quad \epsilon_{\text{DEM-CCA}}^\theta \left(\mathcal{A} \circ \text{MOD-CCA} \circ \frac{\text{KEM}^{1,\eta}}{\text{ID}_{\text{out}(\text{DEM}^\theta)}} \right) + \epsilon_{\text{KEM-CCA}}^\eta \left(\mathcal{A} \circ \text{MOD-CCA} \circ \frac{\text{ID}_{\text{out}(\text{KEM}^\eta)}}{\text{DEM}^{1,\theta}} \right) \end{aligned}$$

where the oracles of MOD-CCA are defined in Fig. 4.

PKGEN()	PKENC(m)	PKDEC(c')
assert $pk = \perp$	assert $pk \neq \perp$	assert $pk \neq \perp$
$pk \leftarrow \text{KEMGEN}()$	assert $c = \perp$	assert $c \neq c'$
return pk	$c_1 \leftarrow \text{ENCAP}()$	$c'_1 c'_2 \leftarrow c'$
	$c_2 \leftarrow \text{ENC}(m)$	if $c'_1 = c_1$ then
	$c \leftarrow c_1 c_2$	$m \leftarrow \text{DEC}(c'_2)$
	return (c)	else
		$k' \leftarrow \text{DECAP}(c'_1)$
		$m \leftarrow \theta.\text{dec}(k', c'_2)$
		return m

Figure 4: MOD-CCA construction.

In Appendix D, we prove via code comparison that for $b \in \{0, 1\}$, $\text{PKE-CCA}^{b,\zeta}$ is perfectly indistinguishable from $\text{MOD-CCA} \circ \frac{\text{KEM}^{0,\eta}}{\text{DEM}^{b,\theta}} \circ \text{KEY}^\lambda$. Thus, for all adversaries \mathcal{A} , we can now apply the single key lemma, Lemma 19.(a), to the adversary $\mathcal{B} = \mathcal{A} \circ \text{MOD-CCA}$, as KEM-CCA^η is a keying game, DEM-CCA^θ is a keyed game, and the two are compatible. For all adversaries \mathcal{B} , we have

$$\mathcal{B} \circ \frac{\text{KEM}^{\eta,0}}{\text{DEM}^{\theta,0}} \circ \text{KEY}^\lambda \stackrel{\epsilon(\mathcal{B})}{\approx} \mathcal{B} \circ \frac{\text{KEM}^{\eta,0}}{\text{DEM}^{\theta,1}} \circ \text{KEY}^\lambda.$$

and the value $\epsilon(\mathcal{B})$ is less or equal to

$$\epsilon_{\text{KEM-CCA}}^\eta \left(\mathcal{B} \circ \frac{\text{ID}_{\text{out}(\text{KEM}^\eta)}}{\text{DEM}^{0,\theta}} \right) + \epsilon_{\text{DEM-CCA}}^\theta \left(\mathcal{B} \circ \frac{\text{KEM}^{1,\eta}}{\text{ID}_{\text{out}(\text{DEM}^\theta)}} \right) + \epsilon_{\text{KEM-CCA}}^\eta \left(\mathcal{B} \circ \frac{\text{ID}_{\text{out}(\text{KEM}^\eta)}}{\text{DEM}^{1,\theta}} \right).$$

5 Multi-Instance Packages and Composition

Definition 25 (Indexed Packages). *For a command c with free names $\text{fn}(c)$ we denote by c_i the command in which every function name $f \in \text{fn}(c)$ is replaced by a name f_i with the additional index i . For function definition $\text{O} = f(x) \mapsto c$, we denote by O_{i-} the definition $f_i(x) \mapsto c$ and by O_i the definition $f_i(x) \mapsto c_i$.*

Let D be a package with function definitions Ω . We denote by D_{i-} and D_i packages with definitions $\{\text{O}_{i-} | \text{O} \in \Omega\}$ and $\{\text{O}_i | \text{O} \in \Omega\}$ respectively. This means that $\text{in}(\text{D}_{i-}) = \text{in}(\text{D})$ and $\text{in}(\text{D}_i) = \{f_i | f \in \text{in}(\text{D})\}$.

Definition 26 (Multi-Instance Operator). *For a package D and $n \in \mathbb{N}$, we define $\prod_{i=1}^n \text{D}_{i-} := (\text{D}_{1-} | \dots | \text{D}_{n-})$ and $\prod_{i=1}^n \text{D}_i := (\text{D}_1 | \dots | \text{D}_n)$.*

Note that using a product sign $\prod_{i=1}^n D_i$ to denote multi-instance parallel composition ($D_1 \mid \dots \mid D_n$) is convenient, since it allows to emphasize the multi-instance notation via a prefix which is more prominent than merely a special subscript or index, it reduces the number of brackets per expression, and it allows to avoid dots. While common in arithmetics and, notably, the π -calculus, product notation might be a bit unusual for cryptographers. Also note that including indices in oracle names assures that instances of the same package have disjoint output interfaces which is necessary for their parallel composition. The following lemma states that the multi-instance operator $\prod_{i=1}^n$ commutes with parallel composition, sequential composition and ID.

Lemma 27 (Multi-Instance Interchange). *Let M and N be packages such that M matches the output interface of N . Let P be a packages such that $\text{out}(M)$ and $\text{out}(P)$ are disjoint. Then, for any number n of instances, the following hold:*

$$\begin{aligned} \prod_{i=1}^n (M \circ N)_i &\stackrel{\text{code}}{\equiv} \prod_{i=1}^n M_i \circ \prod_{i=1}^n N_i & \text{ID}_{\text{out}(\prod_{i=1}^n M_i)} &\stackrel{\text{code}}{\equiv} \prod_{i=1}^n (\text{ID}_{\text{out}(M)})_i \\ \prod_{i=1}^n \left(\frac{M}{P} \right)_i &\stackrel{\text{code}}{\equiv} \frac{\prod_{i=1}^n M_i}{\prod_{i=1}^n P_i} & M_{i-} &\stackrel{\text{code}}{\equiv} \text{ID}_{\text{out}(M),i-} \circ M \end{aligned}$$

Proof. Firstly, note that the package $\prod_{i=1}^n M_i \circ \prod_{i=1}^n N_i$ is well-defined, since $\prod_{i=1}^n M_i$ matches the input interface of $\prod_{i=1}^n N_i$ due to Definition 25. Using the interchange rule, we obtain that it is code equivalent to $\prod_{i=1}^n (M \circ N)_i$. Note that $(\prod_{i=1}^n M_i \mid \prod_{i=1}^n P_i)$ is well-defined due to the disjointness condition on the output interfaces. The term is equal to $\prod_{i=1}^n \left(\frac{M}{P} \right)_i$ by associativity of parallel composition. The last two equations follow by inspection of the ID definitions. \square

5.1 Multi-Instance Lemma

We introduce a multi-instance lemma that allows us to turn arbitrary games using symmetric keys into multi-instance games.

Lemma 28 (Multi-Instance). *Let M be a game pair with distinguishing advantage ϵ_M . Then for any number n of instances, adversaries \mathcal{A} , and reduction \mathcal{R} that samples $j \leftarrow_{\$} \{0, \dots, n-1\}$ and runs*

$$\left(\prod_{i=1}^j M_i^1 \mid \text{ID}_{\text{out}(M), (j+1)-} \mid \prod_{i=j+2}^n M_i^0 \right)$$

we have that the game pair $\text{MI}^b \stackrel{\text{code}}{\equiv} \prod_{i=1}^n M_i^b$ satisfies $\epsilon_{\text{MI}}(\mathcal{A}) \leq n \cdot \epsilon_M(\mathcal{A} \circ \mathcal{R})$.

In Appendix B we provide a systematic recipe for hybrid arguments and instantiate it for the proof of this lemma.

5.2 Multiple Keys Lemma

We now combine key composition and multi-instance lemmas. For this purpose, we use a multi-instance version of the following single-instance package **CKEY**. In contrast to the simpler **KEY** package, **CKEY** allows for corrupted keys (whence the name **CKEY**) and, consequently, needs to allow the symmetric-key protocol to check whether keys are honest.

Definition 29 (**CKEY** Package). *For $\lambda \in \mathbb{N}$, **CKEY** is the package that defines the oracles below, i.e., $\text{out}(\text{CKEY}) = \{\text{GEN}, \text{SET}, \text{CSET}, \text{GET}, \text{HON}\}$.*

GEN()	SET(k')	CSET(k')	GET()	HON()
assert $k = \perp$	assert $k = \perp$	assert $k = \perp$	assert $k \neq \perp$	assert $h \neq \perp$
$k \leftarrow_{\$} \{0, 1\}^\lambda$	$k \leftarrow k'$	$k \leftarrow k'$		
$h \leftarrow 1$	$h \leftarrow 1$	$h \leftarrow 0$	return k	return h

A corruptible keying game is composed of a core keying package and the multi-instance version of **CKEY** $^\lambda$. The core keying package can set corrupt keys via the **CSET** oracle. A corruptible keyed game is single-instance but will be turned into a multi-instance game later. Its core keyed package can access the honesty status of keys via the **HON** oracle.

Definition 30 (Corruptible Keying Game). *A corruptible keying game \mathbb{K} is composed of a core keying packages **CK** and the **CKEY** package as follows:*

$$\mathbb{K}^{b,\lambda} \stackrel{\text{code}}{\equiv} \frac{\text{CK}^{b,\lambda}}{\prod_{i=1}^n (\text{ID}_{\{\text{GET}, \text{HON}\}})_i} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda.$$

where $n, \lambda \in \mathbb{N}$, $b \in \{0, 1\}$, $\text{in}(\text{CK}^{0,\lambda}) = \{\text{SET}_i, \text{CSET}_i\}_{i=1}^n$, and $\text{in}(\text{CK}^{1,\lambda}) = \{\text{GEN}_i, \text{CSET}_i\}_{i=1}^n$.

Definition 31 (Corruptible Keyed Game). *A corruptible keyed game \mathbb{D} is composed of a core keyed package **CD** and the **CKEY** package as follows:*

$$\mathbb{D}^{b,\lambda} \stackrel{\text{code}}{\equiv} \frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}}{\text{CD}^{b,\lambda}} \circ \text{CKEY}^\lambda.$$

where $\lambda \in \mathbb{N}$, $b \in \{0, 1\}$, and $\text{in}(\text{CD}^{0,\lambda}) = \text{in}(\text{CD}^{1,\lambda}) = \{\text{GET}, \text{HON}\}$.

Lemma 32 (Multiple Keys). *Keying and keyed games \mathbb{K} and \mathbb{D} are compatible when they have the same key length λ and they define disjoint oracles $\text{out}(\mathbb{K}) \cap \text{out}(\prod_{i=1}^n \mathbb{D}_i)$. For all compatible corruptible keying and keyed games, with the notation above, we have that*

$$\frac{\text{CK}^0}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda \stackrel{\epsilon}{\approx} \frac{\text{CK}^0}{\prod_{i=1}^n \text{CD}_i^1} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda,$$

where for all adversaries \mathcal{A} , $\epsilon(\mathcal{A})$ is less or equal to

$$\epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\prod_{i=1}^n \text{CD}_i^0} \right) + n \cdot \epsilon_D \left(\mathcal{A} \circ \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \mathcal{R} \right) + \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\prod_{i=1}^n \text{CD}_i^1} \right).$$

where reduction \mathcal{R} samples $j \leftarrow_s \{0, \dots, n-1\}$ and implements the package

$$\left(\prod_{i=1}^j \text{M}_i^1 \middle| (\text{ID}_{\text{out}(\text{M})})_{(j+1)-} \middle| \prod_{i=j+2}^n \text{M}_i^0 \right),$$

where $\text{M}^b \stackrel{\text{code}}{\equiv} \frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}}{\text{CD}^b} \circ \text{CKEY}^\lambda$.

Proof. The proof proceeds analogously to the 3 steps in the proof of Lemma 19.(b)², i.e., idealizing the corruptible keying game, then the corruptible keyed game and then de-idealizing the corruptible keying game. For the algebraic proof steps, we use the multi-instance variants of the identity rule and the interchange rule, as given in Lemma 27.

Multi-instance Lemma. We invoke the multi-instance lemma (Lemma 28) on game pair M with $\text{M}^b \stackrel{\text{code}}{\equiv} \frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}}{\text{CD}^b} \circ \text{CKEY}^\lambda$. By applying the lemma, we obtain that for all adversaries \mathcal{B} , we have

$$\epsilon_{\text{MI}}(\mathcal{B}) \leq n \cdot \epsilon_D(\mathcal{B} \circ \mathcal{R}), \quad (1)$$

where $\text{MI}^b \stackrel{\text{code}}{\equiv} \prod_{i=1}^n \text{M}_i^b$ and reduction \mathcal{R} samples $j \leftarrow_s \{0, \dots, n-1\}$ and implements the package $\left(\prod_{i=1}^j \text{M}_i^1 \middle| (\text{ID}_{\text{out}(\text{M})})_{(j+1)-} \middle| \prod_{i=j+2}^n \text{M}_i^0 \right)$.

Idealizing the keying core package. For the second part of the proof, the steps that idealize the corruptible keying game are analogous to the single-instance key composition proof, and we obtain

$$\frac{\text{CK}^0}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda \stackrel{\epsilon_1}{\approx} \frac{\text{CK}^1}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda,$$

where $\epsilon_1(\mathcal{A}) = \epsilon_K \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\prod_{i=1}^n \text{CD}_i^0} \right)$.

²We could have stated a variant (a) version of these lemma. Key-exchange, however, typically requires variant (b), e.g., because the MAC used for authentication is not idealized as a random strings in the top-level security definition.

Idealizing the multi-instance version of CD^0 . We discuss ϵ_2 after presenting the transformations.

$$\begin{aligned}
& \frac{\text{CK}^1}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda \\
\stackrel{\text{code}}{\equiv} & \frac{\text{CK}^1 \circ \prod_{i=1}^n \text{ID}_{\{\text{GEN}, \text{CSET}\}}\}_i}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)} \circ \prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda && \text{(identity rule)} \\
\stackrel{\text{code}}{\equiv} & \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \frac{\prod_{i=1}^n (\text{ID}_{\{\text{GEN}, \text{CSET}\}}\}_i)}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda && \text{(interchange rule)} \\
\stackrel{\text{code}}{\equiv} & \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \prod_{i=1}^n \left(\frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}\}_i}{\text{CD}_i^0} \circ \text{CKEY}_i^\lambda \right) && \text{(interchange rule)} \\
\stackrel{\approx \epsilon}{\approx} & \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \prod_{i=1}^n \left(\frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}\}_i}{\text{CD}_i^1} \circ \text{CKEY}_i^\lambda \right) && (2)
\end{aligned}$$

We have $\epsilon_2(\mathcal{A}) = \epsilon_{\text{MI}} \left(\mathcal{A} \circ \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \right)$. Moreover, plugging in Inequality 1, we obtain

$$\epsilon_2(\mathcal{A}) \leq n \cdot \epsilon_{\text{CD}} \left(\mathcal{A} \circ \frac{\text{CK}^1}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \mathcal{R} \right).$$

De-idealizing the keying core package. In turn to transform Term (2) into $\frac{\text{CK}^1}{\prod_{i=1}^n \text{CD}_i^1} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda$, we perform the first 3 transformation steps above in reverse order. We de-idealizing analogous to Lemma 19.(b) and obtain

$$\frac{\text{CK}^1}{\prod_{i=1}^n \text{CD}_i^1} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda \stackrel{\epsilon_3}{\approx} \frac{\text{CK}^0}{\prod_{i=1}^n \text{CD}_i^1} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda,$$

where for all adversaries \mathcal{A} , we have $\epsilon_3(\mathcal{A}) = \epsilon_{\text{CK}} \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{CK})}}{\prod_{i=1}^n \text{CD}_i^1} \right)$. □

6 Composition of Forward-Secure Key Exchange

In this section, we apply the multiple keys lemma (Lemma 32) to forward-secure key exchange. We start with a short definition of authenticated key exchange (AKE) protocols with forward security based on the definition of forward security by Bellare, Rogaway and Pointcheval [BPR00] adapted from password authentication to the setting with asymmetric long-term keys. Like them, we use partnering functions as a partnering mechanism. Unlike

them, we do not encode security against passive adversaries via an `Execute` query but rather require the existence of an origin-session, as suggested by Cremers and Feltz [CF15]. Note that we encode the existence of an origin-session also via partnering functions. Brzuska, Fischlin, Warinschi and Williams [BFWW11] essentially use the same security definition as in the present paper, except that they did not encode passivity and used session identifiers instead of partnering functions. We explain our definitional choices at the end of this section.

Definition 33 (Key Exchange Protocol). *A key exchange protocol π consists of a key generation function $\pi.kgen$ and a protocol function $\pi.run$. $\pi.kgen$ returns a pair of keys, i.e., $(sk, pk) \leftarrow_{\$} \pi.kgen$. $\pi.run$ takes as input a state and an incoming message and returns a state and an outgoing message, i.e., $(state', m') \leftarrow_{\$} \pi.run(state, m)$.*

Each party holds several sessions and the function $\pi.run$ is executed locally on the *session* state. We use indices i for sessions and indices u, v for parties. For the i th session of party u , we denote the state by $\Pi[u, i].state$. The state contains at least the following variables. For a variable a , we denote by $\Pi[u, i].a$ the variable a stored in $\Pi[u, i].state$.

- (pk, sk) : the party's own public-key and corresponding private key
- $peer$: the public-key of the intended peer for the session
- $role$: determines whether the session runs as an initiator or responder
- α : protocol status that is either *running* or *accepted*.
- k : the symmetric session key derived by the session

Upon initialization of each session, the session state is initialized with pair (pk, sk) , the public-key $peer$ of the intended peer of a session, a value $role \in \{I, R\}$, $\alpha = running$ and $k = \perp$. The first three variables cannot be changed. The variables α and k can be set only once. We require that

$$\Pi[u, i].\alpha = accepted \implies \Pi[u, i].k \neq \perp.$$

The game that we will define soon will run $(state', m') \leftarrow_{\$} \pi.run(state, \perp)$ on the initial state $state$ and an empty message \perp . For initiator roles, this first *run* returns $m' \neq \perp$, and for responder roles, it outputs $m' = \perp$.

Protocol correctness. For all pairs of sessions which are initialized with (pk_I, sk_I) , pk_R , $role = I$, $\alpha = running$ and $k = \perp$ for one session, and (pk_R, sk_R) , pk_I , $role = R$, $\alpha = running$ and $k = \perp$ for the other session, the following holds: When the messages produced by $\pi.run$ are faithfully transmitted to the other session, then eventually, both sessions have $\alpha = accepted$ and hold the same key $k \neq \perp$.

Partnering. As a partnering mechanism, we use sound partnering functions, one of the partnering mechanisms suggested by Bellare and Rogaway [BR95]. Discussing the specifics, advantages and disadvantages of partnering mechanisms is beyond the scope of this work, we provide a short discussion as well as a definition and the soundness requirement for partnering functions in Appendix C. For the sake of the AKE definition presented in this section, the reader may think of the partnering function $f(u, i)$ as indicating the (first) session (v, j) which derived the same key as (u, i) , has a different role than (u, i) , and is the intended peer of (u, i) . On accepted sessions, it is a symmetric function, thus partners of sessions, if they exist, are unique.

Session key handles. Upon acceptance the SEND oracle returns the index of the CKEY package from which the session key can be retrieved using GET. This index is an administrative identifier that is set when the first of two partnered sessions accept. The second accepting session is then assigned the same identifier as its partner session.

Definition 34 (IND-AKE Security). *For a key exchange protocol $\pi = (\textit{kgen}, \textit{run})$, a symmetric, monotonic, sound partnering function f , and a number of instances $n \in \mathbb{N}$, we define IND-AKE advantage $\epsilon_{\text{IND-AKE}}^{\pi, f, n}$ using a keying game $\text{IND-AKE}^{\pi, f, n}$ with corruptible keying package $\text{AKE}^{b, \pi, f}$ whose oracles are defined in Fig. 5 yielding output interface $\text{out}(\text{IND-AKE}^{\pi, f, n}) = \{\text{NEWPARTY}, \text{NEWSESSION}, \text{SEND}, \text{CORRUPT}, \text{GET}\}$.*

Theorem 35 (BR-Secure Key Exchange is Composable). *Let π be a key exchange protocol with partnering function f such that for $n, \lambda \in \mathbb{N}$, their IND-AKE advantage is $\epsilon_{\text{IND-AKE}}^{\pi, f, n}$. Let D be a corruptible keyed game that is compatible with the corruptible keying game $\text{IND-AKE}^{\pi, f, n}$. Then it holds that*

$$\frac{\text{AKE}^{0, \pi, f}}{\prod_{i=1}^n \text{CD}_i^0} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda \stackrel{\epsilon_{\text{BR}}}{\approx} \frac{\text{AKE}^{0, \pi, f}}{\prod_{i=1}^n \text{CD}_i^1} \circ \prod_{i=1}^n \text{CKEY}_i^\lambda,$$

where

$$\begin{aligned} \epsilon_{\text{BR}}(\mathcal{A}) \leq & \epsilon_{\text{IND-AKE}}^{\pi, f, n} \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{AKE})}}{\prod_{i=1}^n \text{CD}_i^0} \right) + n \cdot \epsilon_{\text{CD}} \left(\mathcal{A} \circ \frac{\text{AKE}^{1, \pi, f}}{\text{ID}_{\text{out}(\prod_{i=1}^n \text{CD}_i)}} \circ \mathcal{R} \right) \\ & + \epsilon_{\text{IND-AKE}}^{\pi, f, n} \left(\mathcal{A} \circ \frac{\text{ID}_{\text{out}(\text{AKE})}}{\prod_{i=1}^n \text{CD}_i^1} \right), \end{aligned}$$

and where reduction \mathcal{R} samples $j \leftarrow_{\text{s}} \{1, \dots, n\}$ and implements the package

$$\left(\prod_{i=1}^{j-1} \text{M}_i^0 \middle| (\text{ID}_{\text{out}(\text{M})})_{j-} \middle| \prod_{i=j+1}^n \text{M}_i^1 \right),$$

where $\text{M}^b \stackrel{\textit{code}}{\equiv} \frac{\text{ID}_{\{\text{GEN}, \text{CSET}\}}}{\text{CD}^0} \circ \text{CKEY}^\lambda$.

<p><u>NEWSSESSION(u, i, r, v)</u></p> <p>assert $PK[u] \neq \perp, PK[v] \neq \perp, \Pi[u, i] = \perp$</p> <p>$\Pi[u, i] \leftarrow ($ $\quad (pk, sk) \leftarrow (PK[u], SK[u]),$ $\quad peer \leftarrow v,$ $\quad role \leftarrow r,$ $\quad \alpha \leftarrow running,$ $\quad k \leftarrow \perp)$</p> <p>$(\Pi[u, i], m) \leftarrow \pi.run(\Pi[u, i], \perp)$</p> <p>return m</p>	<p><u>NEWPARTY(u)</u></p> <p>assert $PK[u] = \perp$</p> <p>$(SK[u], PK[u]) \leftarrow \pi.kgen$</p> <p>$H[u] \leftarrow 1$</p> <p>return $PK[u]$</p>
<p><u>SEND(u, i, m)</u></p> <p>assert $\Pi[u, i].\alpha = running$</p> <p>$(\Pi[u, i], m') \leftarrow \pi.run(\Pi[u, i], m)$</p> <p>if $\Pi[u, i].\alpha \neq accepted$ then \quad return $(m', \perp).$</p> <p>if $\Pi[f(u, i)].\alpha = accepted$ then \quad return $(m', ID[f(u, i)])$</p> <p>$ID[u, i] \leftarrow cntr$</p> <p>if $H[\Pi[u, i].peer] = 1 \vee f(u, i) \neq \perp$ then \quad if $b = 0$ then $\quad\quad$ $SET_{cntr}(\Pi[u, i].k)$ \quad else $\quad\quad$ $GEN_{cntr}()$</p> <p>else \quad $CSET_{cntr}(\Pi[u, i].k)$</p> <p>$cntr \leftarrow cntr + 1$</p> <p>return $(m', ID[u, i])$</p>	<p><u>CORRUPT(u)</u></p> <p>$H[u] \leftarrow 0$</p> <p>return $SK[u]$</p>

Figure 5: Oracles of the core keying package AKE. $cntr$ is initialized to 0.

Proof. We observe that Theorem 35 is a direct application of the multiple keys lemma (Lemma 32). Firstly, AKE is a corruptible core keying package as we have that $\text{in}(\text{AKE}^{0,\pi,f}) = \{\text{SET}, \text{CSET}\}$ and $\text{in}(\text{AKE}^{1,\pi,f}) = \{\text{GEN}, \text{CSET}\}$. Also, by definition, D is a corruptible keyed game that is compatible with the corruptible keying game $\text{IND-AKE}^{\pi,f,n}$. \square

Discussion of definitional choices. Forward secrecy usually requires a notion of time that cryptographic games are not naturally endowed with and that we have no tools to handle in hand-written proofs. In the miTLS work and also in our notation of key exchange security, instead, it is decided *upon acceptance* whether a session shall be idealized or not. The advantage is that one can check *in the moment of acceptance* whether the preconditions for freshness are satisfied, and this check does not require a notion of time. In our encoding, the CKEY package then stores either a real or a random key, and when the partner of the session accepts, the partner session inherits these idealization or non-idealization properties. A downside of this encoding is that it is only suitable for protocols with explicit entity authentication (See, e.g., Fischlin, Günther, Schmidt and Warinschi [FGSW16]), as in those, the first accepting session is already idealized. In particular, our model does not capture two-flow protocols such as HMQV [Kra05].

Using partnering functions instead of session identifiers or key partnering has the advantage that the *at most* condition of Match security defined by Brzuska, Fischlin, Smart, Warinschi and Williams [BFS⁺13] holds syntactically. Thus, one does not need to make probabilistic statements that are external to the games. Note that we made another simplification to the model: Currently, the CKEY module and thus CD does not receive information about the timing of acceptance. This can be integrated at the cost of a more complex CKEY module.

Acknowledgements. We are deeply indebted to Cas Cremers for extensive feedback on an early draft of our article. We are grateful to Simon Peyton Jones for pointing out the associativity of Monadic composition as a generalization of function composition to effectful programs. We thank Giorgia Azzurra Marson and Hoeteck Wee for feedback on the presentation of our IND-CPA toy example. We thank Martijn Stam for suggesting to use KEM-DEM composition as one of our application cases. We thank Théo Winterhalter for pointing us toward a subtle but important mistake in our KEM-DEM proof that he and his co-authors discovered while formalizing our paper using SSProve. We are grateful to Håkon Jacobsen for feedback on our key exchange definition. We thank Ueli Maurer for an inspiring and helpful discussion on abstraction. We thank Sabine Oechsner, Frieder Steinmetz, Bogdan Warinschi, Jan Winkelmann, and Santiago Zanella-Béguelin for helpful suggestions and inspiration.

Chris Brzuska is grateful to NXP for the support of his previously held chair of IT Security Analysis at TU Hamburg. Much of the research was done while the first author was at Microsoft Research Cambridge and during internships and research visits supported by Microsoft and the EU COST framework. In particular, this work was supported by an STSM Grant from COST Action IC1306 “Cryptography for Secure Digital Interaction”. This work was supported by Microsoft Research through its PhD Scholarship Programme. Markulf Kohlweiss is grateful for a fellowship from IOHK.

References

- [AHR⁺21] Carmine Abate, Philipp G. Haselwarter, Exequiel Rivas, Antoine Van Muylder, Théo Winterhalter, Catalin Hritcu, Kenji Maillard, and Bas Spitters. Ssprove: A foundational framework for modular cryptographic proofs in coq. *Cryptology ePrint Archive*, Report 2021/397, 2021. <https://eprint.iacr.org/2021/397>.
- [BBM00] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In *EUROCRYPT 2000*. Springer, 2000.
- [BCLS15] Gilles Barthe, Juan Manuel Crespo, Yassine Lakhnech, and Benedikt Schmidt. Mind the gap: Modular machine-checked proofs of one-round key exchange protocols. In *EUROCRYPT*, 2015.
- [BDKL10] Gilles Barthe, Marion Daubignard, Bruce M. Kapron, and Yassine Lakhnech. Computational indistinguishability logic. In *ACM CCS*, pages 375–386, 2010.
- [BFK⁺13] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, and Pierre-Yves Strub. Implementing TLS with verified cryptographic security. In *Security and Privacy*, 2013.
- [BFK⁺14] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Santiago Zanella Béguelin. Proving the TLS handshake secure (as it is). In *CRYPTO*, 2014.
- [BFS⁺13] Christina Brzuska, Marc Fischlin, Nigel P. Smart, Bogdan Warinschi, and Stephen C. Williams. Less is more: relaxed yet composable security notions for key exchange. *Int. J. Inf. Sec.*, 12(4), 2013.
- [BFWW11] Christina Brzuska, Marc Fischlin, Bogdan Warinschi, and Stephen C. Williams. Composability of Bellare-Rogaway key exchange protocols. In *ACM CCS*, 2011.
- [BGHB11] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. Computer-aided security proofs for the working cryptographer. In *CRYPTO*, 2011.
- [BJ17] Chris Brzuska and Håkon Jacobsen. A modular security analysis of EAP and IEEE 802.11. In *PKC*, 2017.
- [BL13] Daniel J. Bernstein and Tanja Lange. Non-uniform cracks in the concrete: The power of free precomputation. In *ASIACRYPT*, 2013.

- [Bla08] Bruno Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Trans. Dependable Sec. Comput.*, 5(4):193–207, 2008.
- [Bla18] Bruno Blanchet. Composition theorems for cryptoverif and application to TLS 1.3. In *31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018*, pages 16–30, 2018.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In *EUROCRYPT*, 2000.
- [BPW04] Michael Backes, Birgit Pfitzmann, and Michael Waidner. A general composition theorem for secure reactive systems. In *TCC*, 2004.
- [BR95] Mihir Bellare and Phillip Rogaway. Provably secure session key distribution: the three party case. In *STOC*, 1995.
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT*, 2006.
- [Brz13] Christina Brzuska. *On the foundations of key exchange*. PhD thesis, Darmstadt University of Technology, Germany, 2013.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
- [CCD⁺17] Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *EuroS&P 2017*, 2017.
- [CF15] Cas J. F. Cremers and Michèle Feltz. Beyond eck: perfect forward secrecy under actor compromise and ephemeral-key reveal. *Des. Codes Cryptography*, 74(1), 2015.
- [CS03] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Comput.*, 2003.
- [DFGS15] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In *ACM CCS*, 2015.
- [DFK⁺17] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Jonathan Protzenko, Aseem Rastogi, Nikhil Swamy, Santiago Zanella Béguelin,

- Karthikeyan Bhargavan, Jianyang Pan, and Jean Karim Zinzindohoue. Implementing and proving the TLS 1.3 record layer. In *Security and Privacy*, 2017.
- [DKO21] François Dupressoir, Konrad Kohbrok, and Sabine Oechsner. Bringing state-separating proofs to easycrypt - a security proof for cryptobox. Cryptology ePrint Archive, Report 2021/326, 2021. <https://eprint.iacr.org/2021/326>.
- [FGSW16] Marc Fischlin, Felix Günther, Benedikt Schmidt, and Bogdan Warinschi. Key confirmation in key exchange: A formal treatment and implications for TLS 1.3. In *Security and Privacy*, 2016.
- [FKS11] Cédric Fournet, Markulf Kohlweiss, and Pierre-Yves Strub. Modular code-based cryptographic verification. In *ACM CCS*, 2011.
- [HS11] Dennis Hofheinz and Victor Shoup. GNUC: A new universal composability framework. Cryptology ePrint Archive, Report 2011/303, 2011. <http://eprint.iacr.org/2011/303>.
- [HS15] Dennis Hofheinz and Victor Shoup. GNUC: A new universal composability framework. *Journal of Cryptology*, 28(3), 2015.
- [Jac17] Håkon Jacobsen. *A Modular Security Analysis of EAP and IEEE 802.11*. PhD thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2017.
- [JKSS12] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In *CRYPTO 2012*, 2012.
- [Jon03] Simon Peyton Jones. Haskell 98 language and libraries: the revised report, 2003.
- [KMO⁺15] Markulf Kohlweiss, Ueli Maurer, Cristina Onete, Björn Tackmann, and Daniele Venturi. De-Constructing TLS 1.3. In *INDOCRYPT*, 2015.
- [KPW13] Hugo Krawczyk, Kenneth G. Paterson, and Hoeteck Wee. On the security of the TLS protocol: A systematic analysis. In *CRYPTO 2013*, 2013.
- [Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In *CRYPTO*. Springer, 2005.
- [KT13] Ralf Küsters and Max Tuengerthal. The IITM model: a simple and expressive model for universal composability. Cryptology ePrint Archive 2013/025, 2013.

- [Mau02] Ueli M. Maurer. Indistinguishability of random systems. In *EUROCRYPT*, 2002.
- [Mau10] Ueli Maurer. Constructive cryptography - a primer (invited paper). In *FC*, 2010.
- [Mau11] Ueli Maurer. Constructive cryptography - A new paradigm for security definitions and proofs. In *TOSCA*, 2011.
- [MPW92] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Inf. Comput.*, 100(1), 1992.
- [MQU07] Jörn Müller-Quade and Dominique Unruh. Long-term security and universal composability. In *TCC*, 2007.
- [MR11] Ueli Maurer and Renato Renner. Abstract cryptography. In *ITCS*, 2011.
- [MRST06] John C. Mitchell, Ajith Ramanathan, Andre Scedrov, and Vanessa Teague. A probabilistic polynomial-time process calculus for the analysis of cryptographic protocols. *Theor. Comput. Sci.*, 353(1-3), 2006.
- [MT13] Daniele Micciancio and Stefano Tessaro. An equational approach to secure multi-party computation. In *Innovations in Theoretical Computer Science, ITCS*, 2013.
- [Rog06] Phillip Rogaway. Formalizing human ignorance. In *VIETCRYPT*, 2006.
- [Ros18] Mike Rosulek. The joy of cryptography. Online Draft, 2018. <http://web.engr.oregonstate.edu/~rosulekm/crypto/>.
- [RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In *EUROCRYPT*, 2011.
- [SGC12] Don Syme, Adam Granicz, and Antonio Cisternino. *Expert F# 3.0*. Springer, 2012.
- [SHK⁺16] Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. Dependent types and multi-monadic effects in F*. In *POPL*, 2016.
- [Tof96] Mads Tofte. Essentials of standard ML modules. In *Advanced Functional Programming, Second International School, Olympia*, 1996.

- [vLW01] Jan van Leeuwen and Jirí Wiedermann. Beyond the turing limit: Evolving interactive systems. In *SOFSEM 2001*. Springer, 2001.
- [Wik16] Douglas Wikström. Simplified universal composability framework. In *TCC*, 2016.

$\text{MOD-CPA}^b.\text{ENC}(m)$	$\text{MOD-CPA}^b.\text{ENC}(m)$	$\text{MOD-CPA}^b.\text{ENC}(m)$	$\text{IND-CPA}^b.\text{ENC}(m)$
		if $k = \perp$ then	if $k = \perp$ then
		$k \leftarrow_{\$} \{0, 1\}^n$	$k \leftarrow_{\$} \{0, 1\}^n$
if $b = 0$ then	if $b = 0$ then	if $b = 0$ then	if $b = 0$ then
$r \leftarrow_{\$} \{0, 1\}^n$	$r \leftarrow_{\$} \{0, 1\}^n$	$r \leftarrow_{\$} \{0, 1\}^n$	$(r, c) \leftarrow_{\$} \zeta.\text{enc}(k, m)$
$\text{pad} \leftarrow \text{EVAL}(r)$	if $k = \perp$ then		
	$k \leftarrow_{\$} \{0, 1\}^n$		
	$\text{pad} \leftarrow \text{prf}(k, r)$	$\text{pad} \leftarrow \text{prf}(k, r)$	
$c \leftarrow m \oplus \text{pad}$	$c \leftarrow m \oplus \text{pad}$	$c \leftarrow m \oplus \text{pad}$	
if $b = 1$ then	if $b = 1$ then	if $b = 1$ then	if $b = 1$ then
$m' \leftarrow_{\$} \{0, 1\}^n$	$m' \leftarrow_{\$} \{0, 1\}^n$	$m' \leftarrow_{\$} \{0, 1\}^n$	$m' \leftarrow_{\$} \{0, 1\}^n$
$r \leftarrow_{\$} \{0, 1\}^n$	$r \leftarrow_{\$} \{0, 1\}^n$	$r \leftarrow_{\$} \{0, 1\}^n$	$(r, c) \leftarrow_{\$} \zeta.\text{enc}(k, m')$
$\text{pad} \leftarrow \text{EVAL}(r)$	if $k = \perp$ then		
	$k \leftarrow_{\$} \{0, 1\}^n$		
	$\text{pad} \leftarrow \text{prf}(k, r)$	$\text{pad} \leftarrow \text{prf}(k, r)$	
$c \leftarrow m' \oplus \text{pad}$	$c \leftarrow m' \oplus \text{pad}$	$c \leftarrow m' \oplus \text{pad}$	
return (r, c)	return (r, c)	return (r, c)	return (r, c)

Figure 6: The left-most column shows the modular game MOD-CPA that uses an oracle EVAL . From the left-most to the second-left column, we inline the code of $\text{PRF}^0.\text{EVAL}$. From the second-left to the second-right column, we use Bellare-Rogaway-like code-comparison to see that the key generation can be moved up, as it is the same in both branches of the program. We get from the second-right column to the right-most column by considering the code of our concrete construction ζ .

A Example for the usefulness of associativity

Given a pseudorandom function (PRF), we construct a symmetric encryption scheme that is indistinguishable under chosen plaintext attacks (IND-CPA). The goal of this example is to showcase the usefulness of *associativity* of algorithm composition for the writing of reductions. We will write the IND-CPA game in a modular way that makes the game-hop which replaces the PRF with a random function immediate and thereby modularizes the proof. As is good cryptographic practice, we proceed as follows:

- (1) Specification of security goal: IND-CPA secure symmetric encryption.
- (2) Specification of cryptographic assumptions: PRF security.

- (3) Construction: We build a symmetric encryption scheme from a PRF.
- (4) Reduction: We prove that, if the assumption holds, then our construction satisfies the security goal. I.e., we build a reduction that simulates the IND-CPA game (instantiated with the construction) given oracle access to the PRF game.

(1) IND-CPA security. In the real-or-ideal formalization of IND-CPA security, the adversary has adaptive access to an encryption oracle ENC to which they can adaptively submit a message m . The adversary receives either an encryption of m , or an encryption of a random string of the same length as m . The adversary then needs to distinguish between the two distributions.³ Note that we operate in the concrete security setting as it is more adequate for practice-oriented cryptography and therefore only define advantages rather than security in line with the critique of Rogaway [Rog06], Bernstein and Lange [BL13]. Our ideas can be transferred analogously to the asymptotic setting.

We denote the interaction of the adversary with the encryption oracle as $\mathcal{A} \circ \text{ENC}$ instead of the common notation \mathcal{A}^{ENC} . Moreover, we use the name of the game rather than the oracle, writing $\mathcal{A} \circ \text{IND-CPA}^b$. These convention are inessential on our example, but will be convenient in more complex settings.

Definition 36 (IND-CPA Security). *Let $\zeta = (\zeta.kgen, \zeta.enc, \zeta.dec)$ be a symmetric encryption scheme. The IND-CPA advantage $\epsilon_{\text{IND-CPA}}^{\zeta}(\mathcal{A})$ of adversary \mathcal{A} is*

$$2 \cdot \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \text{IND-CPA}^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{IND-CPA}^1 \right] \right|.$$

We consider $\epsilon_{\text{IND-CPA}}^{\zeta}$ as a function of the adversary and write, equivalently,

$$\text{IND-CPA}^0 \stackrel{\epsilon_{\text{IND-CPA}}^{\zeta}}{\approx} \text{IND-CPA}^1.$$

The game pair IND-CPA^0 and IND-CPA^1 is specified in the right-most column of Figure 6.

(2) PRF security. Given a pseudorandom function prf , we define a security game where the adversary's task is to distinguish between (a) PRF^0 with an EVAL oracle using the real prf function and (b) PRF^1 with an EVAL oracle implementing a random function. For disambiguation based on the secret bit b , we write $\text{PRF}^b.\text{EVAL}$ for the respective oracles.

Definition 37 (PRF Security). *Given a pseudorandom function prf with key length n , we write $\epsilon_{\text{PRF}}^{\text{prf}}(\mathcal{A})$ for the advantage of an adversary \mathcal{A} distinguishing between PRF^0 and PRF^1 .*

³Note that this definition of IND-CPA security is equivalent (by a factor of 2) to the standard left-or-right IND-CPA security definition. We prefer to use a real-or-ideal definitional style since such definitions tend to ease composition, as already observed by Canetti [Can01].

(3) Construction. We construct our symmetric encryption scheme $\zeta = (kgen, enc, dec)$ in Figure 7.

(4) Reduction. We reduce the IND-CPA security of the encryption scheme ζ to the PRF security of prf . Towards this goal, for both

$b \in \{0, 1\}$, we provide a modularized description of IND-CPA^b by the package MOD-CPA^b (see Figure on the right). The package MOD-CPA^b uses an EVAL oracle such that, when MOD-CPA^b is composed with PRF^0 , the package $\text{MOD-CPA}^b \circ \text{PRF}^0$ is perfectly indistinguishable from IND-CPA^b (See Figure 6 for the code-based perfect indistinguishability proof.).

Let \mathcal{A} be an adversary. In the following game-hops, note that the PRF advantage appears twice, as the games IND-CPA^0 and IND-CPA^1 both use $\zeta.\text{enc}$ and thus employ the actual prf and not a random function. The first and last transformation follow by perfect indistinguishability (See Figure 6) and are proven via inlining the code of the corresponding oracles into MOD-CPA^b . The PRF assumption and associativity of algorithm composition cover all other steps, except for the one

$\zeta.kgen$	$\zeta.enc(k, m)$	$\zeta.dec(k, (r, c))$
$k \leftarrow_{\$} \{0, 1\}^n$	$r \leftarrow_{\$} \{0, 1\}^n$	
return k	$pad \leftarrow \text{prf}(k, r)$	$pad \leftarrow \text{prf}(k, r)$
	$c \leftarrow m \oplus pad$	$m \leftarrow c \oplus pad$
	return (r, c)	return m

$\text{PRF}^0.\text{EVAL}(x)$	$\text{PRF}^1.\text{EVAL}(x)$
if $k = \perp$ then	if $T[x] = \perp$ then
$k \leftarrow_{\$} \{0, 1\}^n$	$T[x] \leftarrow_{\$} \{0, 1\}^n$
$y \leftarrow \text{prf}(k, x)$	$y \leftarrow T[x]$
return y	return y

$\text{MOD-CPA}^b.\text{ENC}(m)$
if $b = 0$ then
$r \leftarrow_{\$} \{0, 1\}^n$
$pad \leftarrow \text{EVAL}(r)$
$c \leftarrow m \oplus pad$
if $b = 1$ then
$m' \leftarrow_{\$} \{0, 1\}^n$
$r \leftarrow_{\$} \{0, 1\}^n$
$pad \leftarrow \text{EVAL}(r)$
$c \leftarrow m' \oplus pad$
return (r, c)

Figure 7: Construction of the IND-CPA secure encryption scheme ζ from the pseudorandom function prf . For simplicity, we assume that k, r, pad, m , and c all have the same length n .

labeled *statistical gap*, on which we focus below.

$$\begin{array}{ll}
\mathcal{A} \circ \text{IND-CPA}^0 & \\
\stackrel{\text{perf}}{\equiv} \mathcal{A} \circ \text{MOD-CPA}^0 \circ \text{PRF}^0 & \text{(Perfect equivalence)} \\
\stackrel{\text{code}}{\equiv} (\mathcal{A} \circ \text{MOD-CPA}^0) \circ \text{PRF}^0 & \text{(Associativity)} \\
\stackrel{\epsilon_1(\mathcal{A})}{\approx} (\mathcal{A} \circ \text{MOD-CPA}^0) \circ \text{PRF}^1 & \text{(PRF security, } \epsilon_1(\mathcal{A}) = \epsilon_{\text{PRF}}(\mathcal{A} \circ \text{MOD-CPA}^0) \text{)} \\
\stackrel{\text{code}}{\equiv} \mathcal{A} \circ \text{MOD-CPA}^0 \circ \text{PRF}^1 & \text{(Associativity)} \\
\stackrel{\epsilon_2(\mathcal{A})}{\approx} \mathcal{A} \circ \text{MOD-CPA}^1 \circ \text{PRF}^1 & \text{(Statistical gap)} \\
\stackrel{\text{code}}{\equiv} (\mathcal{A} \circ \text{MOD-CPA}^1) \circ \text{PRF}^1 & \text{(Associativity)} \\
\stackrel{\epsilon_3(\mathcal{A})}{\approx} (\mathcal{A} \circ \text{MOD-CPA}^1) \circ \text{PRF}^0 & \text{(PRF security, } \epsilon_3(\mathcal{A}) = \epsilon_{\text{PRF}}(\mathcal{A} \circ \text{MOD-CPA}^1) \text{)} \\
\stackrel{\text{code}}{\equiv} \mathcal{A} \circ \text{MOD-CPA}^1 \circ \text{PRF}^0 & \text{(Associativity)} \\
\stackrel{\text{perf}}{\equiv} \mathcal{A} \circ \text{IND-CPA}^1 & \text{(Perfect equivalence)}
\end{array}$$

The statistical gap occurs when the game moves from encrypting the adversary's message to encrypting a random message. In both cases, the padding is created via a random function. However, the ciphertext distributions differ whenever there is a collision on the randomness r . In that case, the padding is repeated and therefore, if $b = 0$, the xor of the two ciphertexts equals the xor of the two messages that the adversary queried. In turn, if $b = 1$, then with overwhelming probability, the xor of the two ciphertexts will yield a uniformly random string. Therefore, we need to perform a bad event analysis to bound the probability of a collision on r . Let $q_{\mathcal{A}}$ be an upper bound on the number of oracle calls by the adversary; by the birthday bound, the probability of the bad event is at most $q_{\mathcal{A}}^2/2^{n-1}$ and thus, $\epsilon_2(\mathcal{A}) \leq q_{\mathcal{A}}^2/2^{n-1}$.

Our suggested writing style splits reduction proofs into different kinds of steps. Simple steps such as code equivalence, associativity and using the assumption are carried out separately and algebraically and allow to make the reduction explicit and precise, first as MOD-CPA^0 then MOD-CPA^1 . In turn, the statistical gap argument needed a more complex analysis that can potentially hide subtleties. We thus think that such steps should be avoided whenever possible. E.g., in the aforementioned example, one could use a second assumption such as the indistinguishability of real nonces r (that do have collisions) and ideal nonces (that do not have collisions). The assumption can then be proven without considering the entire IND-CPA game, and it can be used via an algebraic game-hop. However, using more than one assumptions in a proof requires parallel composition and more algebraic rules than associativity. We refer to the main part of the paper for these techniques.

B Hybrid Argument Recipe

Hybrid arguments can be used in various contexts and are the standard technique to reduce multi-instance games to single-instance games. We here write down a general hybrid argument recipe.

Lemma 38 (Hybrid Argument Recipe Lemma). *Let Game^0 , Game^1 , Multi^0 and Multi^1 be four packages with $\text{in}(\text{Game}^0) = \text{in}(\text{Game}^1) = \emptyset$ and $\text{out}(\text{Game}^0) = \text{out}(\text{Game}^1)$ as well as $\text{in}(\text{Multi}^0) = \text{in}(\text{Multi}^1) = \emptyset$ and $\text{out}(\text{Multi}^0) = \text{out}(\text{Multi}^1)$. Let \mathcal{A} be an adversary. Let n be a natural number. Let $\text{H}^0, \dots, \text{H}^n$ be games with $\text{out}(\text{H}^i) = \text{out}(\text{Multi}^1)$, let \mathcal{R}^i be reduction packages with $\text{out}(\mathcal{R}^i) = \text{out}(\text{Multi}^1)$ and $\text{in}(\mathcal{R}^i) = \text{out}(\text{Game}^1)$, and let \mathcal{R} be a package, which samples $j \leftarrow_s \{0, \dots, n-1\}$ and then behaves like \mathcal{R}^j . Then we need to prove the following:*

Claim 1: *It holds that*

$$\text{Multi}^0 \stackrel{\text{perf}}{\equiv} \text{H}^0 \tag{3}$$

$$\text{and Multi}^1 \stackrel{\text{perf}}{\equiv} \text{H}^n \tag{4}$$

Claim 2: *For all $i \in \{0, \dots, n-1\}$ the following holds*

$$\mathcal{R}^i \circ \text{Game}^0 \stackrel{\text{perf}}{\equiv} \text{H}^i \tag{5}$$

$$\text{and } \mathcal{R}^i \circ \text{Game}^1 \stackrel{\text{perf}}{\equiv} \text{H}^{i+1} \tag{6}$$

If Claim 1 and Claim 2 hold, then the package \mathcal{R} satisfies

$$\epsilon_{\text{Multi}}(\mathcal{A}) \leq n \cdot \epsilon_{\text{Game}}(\mathcal{A} \circ \mathcal{R}).$$

Proof. Let \mathcal{A} be an adversary whose input interface matches $\text{out}(\text{Multi}^0)$ and let

$$\epsilon_{i,i'}(\mathcal{A}) = \left| \Pr \left[0 \leftarrow \mathcal{A} \circ \text{H}^i \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \text{H}^{i'} \right] \right|$$

be the distinguishing advantage between hybrids H^i and $\text{H}^{i'}$ for \mathcal{A} .

1. By definition we have

$$\mathcal{A} \circ \text{H}^0 \stackrel{\epsilon_{0,1}(\mathcal{A})}{\approx} \mathcal{A} \circ \text{H}^1 \stackrel{\epsilon_{1,2}(\mathcal{A})}{\approx} \dots \stackrel{\epsilon_{n-2,n-1}(\mathcal{A})}{\approx} \mathcal{A} \circ \text{H}^{n-1} \stackrel{\epsilon_{n-1,n}(\mathcal{A})}{\approx} \mathcal{A} \circ \text{H}^n.$$

From the equation in Claim 1, it follows that $\epsilon_{0,n} = \epsilon_{\text{Multi}}$, i.e.,

$$\mathcal{A} \circ \text{H}^0 \stackrel{\epsilon_{\text{Multi}}(\mathcal{A})}{\approx} \mathcal{A} \circ \text{H}^n$$

By the triangle inequality of \approx we have that

$$\epsilon_{\text{Multi}}(\mathcal{A}) \leq \epsilon_{0,1}(\mathcal{A}) + \cdots + \epsilon_{n-1,n}(\mathcal{A}) = \sum_{\ell=0}^{n-1} \epsilon_{i,i+1}(\mathcal{A})$$

2. Now, we recall the definition of $\epsilon_{i,i+1}$ and plug in Eq. 5 and 6 from Claim 2:

$$\begin{aligned} \epsilon_{\text{Multi}}(\mathcal{A}) &\leq \sum_{i=0}^{n-1} \epsilon_{i,i+1}(\mathcal{A}) \\ &= \sum_{i=0}^{n-1} \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \mathbb{H}^i \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \mathbb{H}^{i+1} \right] \right| \\ &= \sum_{i=0}^{n-1} \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R}^i \circ \text{Game}^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R}^i \circ \text{Game}^1 \right] \right| \end{aligned}$$

When the value j sampled by \mathcal{R} equals i , we have by inlining that $\mathcal{R} \stackrel{\text{code}}{\equiv} \mathcal{R}^i$. Thus, $\epsilon_{\text{Multi}}(\mathcal{A})$ is smaller or equal to

$$\sum_{i=0}^{n-1} \left| \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^0 \mid j = i \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^1 \mid j = i \right] \right| \quad (7)$$

As the sum iterates over all $i \in \{0, \dots, n-1\}$, we obtain

$$\sum_{i=0}^{n-1} \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^b \mid j = i \right] = \frac{\Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^b \right]}{\frac{1}{n}}. \quad (8)$$

Plugging Eq. 8 into Eq. 7 gives us

$$\begin{aligned} \epsilon_{\text{Multi}}(\mathcal{A}) &\leq n \cdot \left(\Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^0 \right] - \Pr \left[1 \leftarrow \mathcal{A} \circ \mathcal{R} \circ \text{Game}^1 \right] \right) \\ &= n \cdot \epsilon_{\text{Game}}(\mathcal{A} \circ \mathcal{R}). \end{aligned}$$

□

We now use the above recipe to provide a proof of Lemma 28. The lemma states that the game pair $\text{MI}^b \stackrel{\text{code}}{\equiv} \prod_{i=1}^n \text{M}_i^b$ satisfies $\epsilon_{\text{MI}}(\mathcal{A}) \leq n \cdot \epsilon_{\text{M}}(\mathcal{A} \circ \mathcal{R})$ for reduction \mathcal{R} that samples $j \leftarrow_s \{0, \dots, n-1\}$ and runs

$$\left(\prod_{i=1}^j \text{M}_i^1 \mid \text{ID}_{\text{out}(\text{M}), (j+1)} - \prod_{i=j+2}^n \text{M}_i^0 \right)$$

Proof. We instantiate the hybrid argument recipe lemma as follows

$$\begin{aligned} \mathbf{Multi}^b &: \equiv \prod_{j=1}^n \mathbf{M}_j^b, \\ \mathbf{Game}^b &: \equiv \mathbf{M}^b. \end{aligned}$$

We define the hybrids \mathbf{H}^i for $0 \leq i \leq n$ as follows

$$\mathbf{H}^i : \equiv \frac{\prod_{j=1}^i \mathbf{M}_j^1}{\prod_{j=i+1}^n \mathbf{M}_j^0}$$

Observe, that indeed, $\mathbf{H}^0 \stackrel{\text{code}}{\equiv} \mathbf{Multi}^0$ and $\mathbf{H}^n \stackrel{\text{code}}{\equiv} \mathbf{Multi}^1$, so Claim 1 holds. We now specify the reduction package \mathcal{R}^i for $0 \leq i \leq n-1$ with $\text{in}(\mathcal{R}^i) = \text{out}(\mathbf{M})$ and $\text{out}(\mathcal{R}^i) = \text{out}(\prod_{j=1}^n \mathbf{M}_j)$. It behaves just as hybrid \mathbf{H}^i , except for instance $i+1$, where \mathcal{R}^i forwards the calls to the oracles provided through its input interface (i.e. \mathbf{M}). Formally,

$$\mathcal{R}^i : \equiv \left(\prod_{j=1}^i \mathbf{M}_j^1 \middle| \left(\text{ID}_{\text{out}(\mathbf{M})} \right)_{(i+1)-} \middle| \prod_{j=i+2}^n \mathbf{M}_j^0 \right)$$

We now need to show that the reductions \mathcal{R}^i satisfy Claim 2. We show Eq. 5, then Eq. 6 follows analogously.

$$\begin{aligned} \mathcal{R}^i \circ \mathbf{M}^0 &\stackrel{\text{code}}{\equiv} \left(\prod_{j=1}^i \mathbf{M}_j^1 \middle| \left(\text{ID}_{\text{out}(\mathbf{M})} \right)_{(i+1)-} \middle| \prod_{j=i+2}^n \mathbf{M}_j^0 \right) \circ \mathbf{M}^0 \\ &\stackrel{\text{code}}{\equiv} \left(\prod_{j=1}^i \mathbf{M}_j^1 \middle| \mathbf{M}_{(i+1)}^0 \middle| \prod_{j=i+2}^n \mathbf{M}_j^0 \right) && \text{(multi-instance interchange rule)} \\ &\stackrel{\text{code}}{\equiv} \left(\prod_{j=1}^i \mathbf{M}_j^1 \middle| \prod_{j=i+1}^n \mathbf{M}_j^0 \right) \\ &\stackrel{\text{code}}{\equiv} \mathbf{H}^i \end{aligned}$$

Therefore, by Lemma 38, \mathcal{R} satisfies

$$\epsilon_{\prod_{j=1}^n \mathbf{M}_j}(\mathcal{A}) \leq n \cdot \epsilon_{\mathbf{M}}(\mathcal{A} \circ \mathcal{R}),$$

which concludes the proof of Lemma 28. □

C Partner Mechanisms in Key Exchange

Partnering is needed in key exchange protocols to specify the pairs of sessions that derive the same key so that security notions for key exchange can exclude trivial winning strategies, such as revealing the key of a partner session. The original BFWW work showed that for composition, the reduction needs to know the partnering between sessions. In our model, we give the partnering information directly to the adversary (since the game returns the same id for matching sessions) and thus also to the reduction. There are a multitude of ways to define partnering in key exchange, and partnering in key exchange is an interesting area of research that is not yet fully clarified. For simplicity, we here follow Bellare and Rogaway’s formulation of public partnering functions that map sessions merely based on public transcripts [BR95]. While partnering functions have not been very popular over the past 15 years, Brzuska and Jacobsen [BJ17, Jac17] recently re-discovered partnering functions, because properties of partnering functions such as uniqueness can be required to hold syntactically, while they only hold probabilistically for concepts such as session identifiers and key equality. These syntactic properties simplify our composition theorem as we discuss in the end of Section 6. The following definition is a prose variant of the definition of transcript given by Brzuska and Jacobsen [BJ17, Jac17].

partnering functions are used within key exchange security games and yet, at the same time, the definition of partnering functions requires part of the game as already defined. The way out of the circularity is as follows: (1) The partnering function can be defined syntactically on transcripts, and the transcripts are well-defined also without a partnering function. (2) No probabilistic properties on the partnering function are required, so that we can consider all powerful adversaries in the consideration of the partnering function.

Definition 39 (Transcript). *The public transcript T of a key exchange game consists of all calls to NEWPARTY, NEWSESSION and SEND by the adversary as well as their answers, except for the answers of SEND where only the first component of each answer becomes part of the transcript.*

Definition 40 (Partnering Functions). *A symmetric and monotonic partnering function is a function f , parametrized by a transcript T , that maps pairs (U, i) of sessions to other pairs (V, j) of sessions*

$$1. f_T(U, i) = (V, j) \implies f_T(V, j) = (U, i), \quad (\text{symmetric})$$

$$2. f_T(U, i) = (V, j) \implies f_{T'}(U, i) = (V, j) \text{ for all } T \subseteq T'. \quad (\text{monotonic})$$

Partnering soundness. For a security analysis based on partnering functions to be meaningful, the partnering function needs to satisfy certain soundness properties. Briefly, soundness demands that partners should: (1) end up with the same session key, (2) agree

upon who they are talking to, (3) have compatible roles, and (4) be unique. However, since we are limiting our attention to symmetric partnering functions in this paper, the last requirement follows directly so we omit it.

Definition 41 (Partnering Function Soundness). *A partnering function is sound if the following holds for all transcripts T . If sessions $f_{T'}(U, i) = (V, J)$ then:*

1. $\pi[U, i].\alpha = \pi[V, j].\alpha = \text{accepted} \implies \pi[U, i].k = \pi[V, j].k \neq \perp$,
2. $\pi[U, i].\text{peer} = pk[V]$, and $\pi[V, j].\text{peer} = pk[U]$.
3. $(\pi[U, i].\text{role} = I, \text{ and } \pi[V, j].\text{role} = R)$ or $(\pi[U, i].\text{role} = R, \text{ and } \pi[V, j].\text{role} = I)$

D Perfect Equivalence for MOD-CCA in the Proof of Theorem 24

We prove that for $b \in \{0, 1\}$, $\text{PKE-CCA}^{b, \zeta}$ is perfectly equivalent to $\text{MOD-CCA} \circ \frac{\text{KEM}^{0, \eta}}{\text{DEM}^{b, \theta}} \circ \text{KEY}^\lambda$ as defined in Section 4. The proof proceeds by inlining all oracle calls in MOD-CCA and inlining the construction ζ in $\text{PKE-CCA}^{b, \zeta}$. See Figure 8 and its caption for the details of the inlining. Note, that in $\text{PKE-CCA}^{b, \zeta}$ on the right-most column, we have already inlined the construction ζ and moved the running of the decapsulation algorithm in the line preceding the **if**-statement, as the line is used in both **if**-branches. As noted in the caption, the difference in PKDEC between columns 3 and 4 can be resolved by applying the correctness of $\eta.\text{encap}$, which implies that the keys k and k' are identical in both branches, enabling us to remove the **if**-statement entirely and thus proving the perfect indistinguishability.

PKENC(<i>m</i>)	PKENC(<i>m</i>)	PKENC(<i>m</i>)	PKENC(<i>m</i>)
assert $pk \neq \perp$	assert $pk \neq \perp$	assert $pk \neq \perp$	assert $pk \neq \perp$
assert $c = \perp$	assert $c = \perp$	assert $c = \perp$	assert $c = \perp$
$c_1 \leftarrow \text{ENCAP}()$	$k, c_1 \leftarrow \eta.\text{encap}(pk)$	$k, c_1 \leftarrow \eta.\text{encap}(pk)$	if $b = 0$ then
$c_2 \leftarrow \text{ENC}(m)$	SET (k)	SET (k)	$k, c_1 \leftarrow \eta.\text{encap}(pk)$
	$k \leftarrow \text{GET}()$	$k \leftarrow \text{GET}()$	$c_2 \leftarrow \zeta.\text{enc}(k, m)$
	if $b = 0$ then	if $b = 0$ then	else
	$c_2 \leftarrow \theta.\text{enc}(k, m)$	$c_2 \leftarrow \theta.\text{enc}(k, m)$	$k, c_1 \leftarrow \eta.\text{encap}(pk)$
	else	else	$c_2 \leftarrow \zeta.\text{enc}(k, 0^{ m })$
	$c_2 \leftarrow \theta.\text{enc}(k, 0^{ m })$	$c_2 \leftarrow \theta.\text{enc}(k, 0^{ m })$	
$c \leftarrow c_1 c_2$	$c \leftarrow c_1 c_2$	$c \leftarrow c_1 c_2$	$c \leftarrow c_1 c_2$
return c	return c	return c	return c
PKDEC(c')	PKDEC(c')	PKDEC(c')	PKDEC(c')
assert $pk \neq \perp$	assert $pk \neq \perp$	assert $pk \neq \perp$	assert $pk \neq \perp$
assert $c \neq c'$	assert $c \neq c'$	assert $c \neq c'$	assert $c \neq c'$
$c'_1 c'_2 \leftarrow c'$	$c'_1 c'_2 \leftarrow c'$	$c'_1 c'_2 \leftarrow c'$	$c'_1 c'_2 \leftarrow c'$
if $c'_1 = c_1$ then	if $c'_1 = c_1$ then	if $c'_1 = c_1$ then	
$m \leftarrow \text{DEC}(c'_2)$	$m \leftarrow \text{DEC}(c'_2)$	$k \leftarrow \text{GET}()$	
		$m \leftarrow \theta.\text{dec}(k, c'_2)$	
else	else	else	
$k' \leftarrow \text{DECAP}(c'_1)$	$k' \leftarrow \eta.\text{decap}(sk, c'_1)$	$k' \leftarrow \eta.\text{decap}(sk, c'_1)$	$k \leftarrow \eta.\text{decap}(sk, c'_1)$
$m \leftarrow \theta.\text{dec}(k', c'_2)$	$m \leftarrow \theta.\text{dec}(k', c'_2)$	$m \leftarrow \theta.\text{dec}(k', c'_2)$	$m \leftarrow \theta.\text{dec}(k, c'_2)$
return m	return m	return m	return m

Figure 8: Col. 1-to-2: ENCAP and DECAP of $\text{KEM-CCA}^{0,\eta}$, and ENC and DEC of $\text{DEM-CCA}^{b,\theta}$ are inlined, highlighted in gray. Col. 2-to-3: Calls to SET and GET do not modify k . Col. 3-to-4: In the ENC oracle, we pull the line $k, c_1 \leftarrow \eta.\text{encap}(pk)$ into each of the branches of the **if**-statement. Else, the two columns differ only when $c'_1 = c_1$ in the PKDEC oracle. PKENC can only be called once and thus, MOD-CCA.PKDEC decrypts c'_2 with the symmetric key k that was previously encapsulated in the MOD-CCA.PKENC oracle. By correctness of the KEM we have that $k = \eta.\text{decap}(sk, c'_1)$ and $\eta.\text{dec}$ thus uses the same k in both cases.