

A Framework for Universally Composable Diffie-Hellman Key Exchange

Ralf Küsters and Daniel Rausch

University of Stuttgart

Stuttgart, Germany

Email: {ralf.kuesters, daniel.rausch}@informatik.uni-stuttgart.de

Abstract—The analysis of real-world protocols, in particular key exchange protocols and protocols building on these protocols, is a very complex, error-prone, and tedious task. Besides the complexity of the protocols itself, one important reason for this is that the security of the protocols has to be reduced to the security of the underlying cryptographic primitives for every protocol time and again.

We would therefore like to get rid of reduction proofs for real-world key exchange protocols as much as possible and in many cases altogether, also for higher-level protocols which use the exchanged keys. So far some first steps have been taken in this direction. But existing work is still quite limited, and, for example, does not support Diffie-Hellman (DH) key exchange, a prevalent cryptographic primitive for real-world protocols.

In this paper, building on work by Küsters and Tuengerthal, we provide an ideal functionality in the universal composability setting which supports several common cryptographic primitives, including DH key exchange. This functionality helps to avoid reduction proofs in the analysis of real-world protocols and often eliminates them completely. We also propose a new general ideal key exchange functionality which allows higher-level protocols to use exchanged keys in an ideal way. As a proof of concept, we apply our framework to three practical DH key exchange protocols, namely ISO 9798-3, SIGMA, and OPTLS.

Keywords—protocol security, universal composability, Diffie-Hellman key exchange, reduction proofs, IITM model

I. INTRODUCTION

The analysis of security protocols, in particular real-world security protocols is a very complex and challenging task, which has gained a lot of attention in the past few years (see, e.g., [1]–[12]). Several approaches for the analysis of such protocols exist, ranging from manual to tool-supported approaches and from symbolic (Dolev-Yao-style) approaches, which abstract from cryptographic details, to approaches based on cryptographic games and those which perform cryptographic reasoning on implementations directly. In this work, our focus lies on cryptographic approaches.

All such approaches strive to achieve some kind of modularity in order to tame the complexity of the analysis (see, e.g., [3], [9], [13], [14]). But security proofs are typically still very complex, tedious, and error-prone. Besides the complexity of the protocols itself, an important reason for this is that for every protocol one has to carry out reduction proofs from the security notions of the protocols to the cryptographic primitives employed time and again. Even in

universal composability models [15]–[18], for which modularity is the driving force, protocol designers typically have to carry out (tedious, repetitive, and error-prone) reduction proofs.

One important goal of this work is therefore to provide a framework within the setting of universal composability (cf. Section II) which gets rid of reduction proofs as much as possible or ideally even altogether, and which is applicable to a wide range of real-world security protocols. This should lead to proofs that are shorter, without being imprecise, as well as easier to understand and carry out. Being based in the setting of universal composability, the framework should also facilitate modular reasoning, allow for re-using existing results, and of course provide security in arbitrary adversarial environments (universal composition).

The main idea behind our framework, which builds on and extends work by Küsters and Tuengerthal [11], [19] (see below), is as follows. First recall that in models for universal composability security properties are expressed by so-called ideal functionalities, which perform their tasks in an ideal secure way. A real protocol \mathcal{P}' may use an ideal functionality \mathcal{F} (or several such functionalities) as a subroutine to perform its task. Typically one shows that \mathcal{P}' (along with \mathcal{F}) realizes another (higher-level) ideal functionality, say \mathcal{F}' . Composition theorems available in models for universal composability then allow one to replace \mathcal{F} by its realization \mathcal{P} , which then implies that \mathcal{P}' using \mathcal{P} realizes \mathcal{F}' . Now, in our framework we provide an ideal functionality $\mathcal{F}_{\text{crypto}}$ which covers various cryptographic primitives, including standard Diffie-Hellman (DH) key exchanges based on the DDH assumption, symmetric/asymmetric encryption, key derivation, MACing, and signing. We show that $\mathcal{F}_{\text{crypto}}$ can be realized by standard cryptographic assumptions, which is a once and for all effort. In essentially all other approaches for protocol analysis this kind of reduction to the cryptographic assumptions of primitives has to be carried out time and again in the analysis of every single protocol. In contrast, in our framework one can prove the security of a protocol \mathcal{P} using $\mathcal{F}_{\text{crypto}}$ without using any reduction proofs or hybrid arguments (at least not for the primitives supported by $\mathcal{F}_{\text{crypto}}$). In a last step, by composition theorems, $\mathcal{F}_{\text{crypto}}$ can be replaced by its realization so that the ideal cryptographic primitives are replaced by their real counterparts.

All primitives provided by $\mathcal{F}_{\text{crypto}}$ can be used with each

other in an idealized way. For example, a protocol \mathcal{P} using $\mathcal{F}_{\text{crypto}}$ can first exchange a key via an ideal Diffie-Hellman key exchange where some messages are (ideally) signed and then derive a MAC and a symmetric encryption key from the DH key. Importantly, both keys can still be used in an idealized way, i.e., one can perform ideal MACing and encryption using these keys.

In addition to $\mathcal{F}_{\text{crypto}}$, our framework also provides new functionalities for ideal key exchange that allow a higher level protocol to still use a session key in an idealized way.

Altogether, when using these functionalities, the need for reduction proofs is greatly reduced or such proofs are avoided completely in many cases. Protocol designers can argue on an intuitive information theoretic level while being able to analyze a protocol in a very modular way with universally composable security guarantees.

Contributions. More specifically, our contributions are as follows.

- We extend the ideal functionality $\mathcal{F}_{\text{crypto}}$ from [19] to also support standard DH key exchange with two key shares g^a and g^b . This is a crucial step as many real-world protocols support Diffie-Hellman key exchanges and thus could not have been analyzed before using $\mathcal{F}_{\text{crypto}}$. Designing such an extension requires care in order for the extension to, on the one hand, provide all expected properties and, on the other hand, still be realizable under standard cryptographic assumptions.
- Our functionality $\mathcal{F}_{\text{crypto}}$ ensures that the adversary on the network cannot interfere with higher level protocols while they use $\mathcal{F}_{\text{crypto}}$ to perform local computations. While this is expected and natural for such an ideal functionality, it previously was impossible to model this property. Leveraging fundamental results of recent work by Camenisch et al. [20], who have introduced the concept of responsive environments, we can now indeed provide this property for $\mathcal{F}_{\text{crypto}}$, which further simplifies security proofs.
- We propose and prove a realization for $\mathcal{F}_{\text{crypto}}$ based on standard cryptographic assumptions. The proof is quite involved, with several hybrid arguments, as $\mathcal{F}_{\text{crypto}}$ allows for a wide range of operations. But, as explained above, due to the modularity of our framework this is a once and for all effort.
- Inspired by an ideal functionality from [11], we propose two new functionalities for both mutually and unilaterally authenticated key exchange with perfect forward secrecy. Unlike most other key exchange functionalities, which output the key, our functionalities allow higher-level protocols to still use the exchanged key in an ideal way, namely for idealized key derivation, symmetric encryption, and MACing. Hence, as mentioned, one can avoid reductions proofs also for the higher-level protocols, such as secure channel protocols. Further discussion and comparison with other key exchange functionalities is

provided in Section V.

- We illustrate the usefulness of our framework by showing for three different real-world key exchange protocols that they realize our key exchange functionalities with mutual or unilateral authentication. Due to the use of $\mathcal{F}_{\text{crypto}}$, none of the security proofs require any reductions, hybrid arguments, or even probabilistic reasoning.
 - We provide the first analyses of unaltered versions of the ISO 9798-3 [21] and the SIGMA [22] key exchange protocols in an universal composability model (see also Section VII).
 - We analyze the 1-RTT non-static mode of the OPTLS key exchange protocol [23] and find a subtle bug in the original reduction proof. We show that, under the original security assumptions, a slight variation of the protocol is a secure unilaterally authenticated and universally composable key exchange protocol.

Structure of the Paper. In Section II, we briefly recall the IITM model, which is the universal composability model we use in this paper. Section III details the ideal functionality $\mathcal{F}_{\text{crypto}}$, with a realization proposed and proven in Section IV. Our ideal key exchange functionalities are presented in Section V. The case studies are carried out in Section VI. We further discuss advantages and limitations of our framework along with related work in Section VII. We conclude in Section VIII. Further details are provided in the appendix. Full details with all proofs are provided in our technical report [24].

II. THE IITM MODEL

In this section, we briefly recall the IITM model with responsive environments from [20]. This is the model for universal composability we use in this paper. This model in turn is based on the IITM model proposed in [16], [25]. We provide a simplified and high level overview only as the details of this model are not important to follow the rest of the paper. In the IITM model, notions of universal composability are defined based on a general computational model. The model also comes with general composition theorems.

Before we recall the IITM model, we first briefly recall the general idea behind universal composability.

The General Idea Behind Universal Composability. In universal composability models, one considers real and ideal protocols. An ideal protocol, also called ideal functionality, specifies the desired behavior of a protocol, and in particular, its intended security properties. The real protocol, which is the protocol one would like to analyze, is supposed to realize the ideal protocol, i.e., it should be at least as secure as the ideal protocol. More specifically, for every adversary on the real protocol, called the real adversary, there should exist an adversary on the ideal protocol, the ideal adversary (or simulator), such that no environment can distinguish the real

from the ideal setting. Now, since, by definition, there exists no successful attack on the ideal protocol, attacks on the real protocol cannot be successful either.

The General Computational Model. The general computational model of the IITM model is defined in terms of systems of interactive Turing machines. An interactive Turing machine (machine, for short) is a probabilistic polynomial-time Turing machine with named bidirectional tapes. The names determine how different machines are connected in a system of machines.

A system \mathcal{S} of IITMs is of the form $\mathcal{S} = M_1 \mid \cdots \mid M_k \mid !M'_1 \mid \cdots \mid !M'_k$, where the M_i and M'_j are machines. The operator ‘!’ indicates that in a run of a system an unbounded number of (fresh) copies of a machine may be generated; for machines without this operator there is at most one instance of this machine in every run of the system. Systems in which multiple copies of machines may be generated are often needed, e.g., for multi-party protocols or for systems describing the concurrent execution of multiple instances of a protocol. In a run of a system \mathcal{S} , at any time only one machine is active and all other machines wait for new input. A (copy of a) machine M can trigger another (copy of a) machine M' by sending a message on a tape that connects both machines. Identifiers, e.g., session and/or party identifiers, contained in the message can be used to address a specific copy of M' .¹ If a new identifier is used, a fresh copy of M' will be generated (if M' is prefixed with ‘!’). The first machine to be triggered in a run of a system is the so-called master machine. This machine is also triggered if a machine does not produce output. In this paper, the environment (see below) will always play the role of the master machine. A run stops if the master machine does not produce output or a machine outputs a message on a special tape named *decision*. Such a message is considered to be the overall output of the system. Systems will always have polynomial runtime in the security parameter (and possibly the length of auxiliary input).

Two systems \mathcal{P} and \mathcal{Q} are called indistinguishable ($\mathcal{P} \equiv \mathcal{Q}$) if the difference between the probability that \mathcal{P} outputs 1 (on the decision tape) and the probability that \mathcal{Q} outputs 1 is negligible in the security parameter η .

Types of Systems. We need the following terminology. For a system \mathcal{S} , the tapes of machines in \mathcal{S} that do not have a matching tape (belonging to another machine in \mathcal{S}) are called *external*. These tapes are grouped into *I/O* and *network tapes*. We consider three different types of systems, modeling i) *real and ideal protocols/functionalities*, ii) *adversaries and simulators*, and iii) *environments*: Protocol systems and environmental systems are systems which have an I/O and network interface, i.e., they may have I/O and network

¹The IITM model contains a general addressing mechanisms. In this paper, we use a specific instantiation of this mechanism as will be clear from the subsequent sections.

tapes. Adversarial systems only have a network interface. Environmental systems may contain a master machine and may produce output on the decision tape.

Environmental systems and adversarial systems are called *responsive* if they answer so-called *restricting* messages on the network immediately. Restricting messages are of the form $(\text{Respond}, id, m)$ where id and m are arbitrary bit strings. When a responsive environment/adversary receives such a message from a system \mathcal{Q} on some network tape t , it has to ensure that the next message that \mathcal{Q} receives is of the form (id, m') , for some bit string m' , and that this message is received on tape t (except for a negligible set of runs). In this sense, responsive environments/adversaries have to respond immediately to restricting messages, i.e., if an environment wants to continue its interaction with \mathcal{Q} it first has to send the expected response m' . Restricting messages are useful for exchanging purely modeling related information with the adversary without letting the adversary interfere with the protocol in-between. For example, one can use a restricting message to ask the adversary whether he wants to corrupt a new protocol instance. Note that such a request does not actually exist in reality and thus no real adversary can abuse it to disrupt the protocol execution. Consequently, in a security model, the adversary should also not have this ability. Restricting messages allow us enforce this very natural expectation. Overall, restricting messages are a very convenient mechanism that adds extra expressivity to universal composability models and that allows for a natural modeling when adversaries and protocols have to exchange meta information (see [20] for an in depth discussion). In the rest of the paper, we always assume that environmental and adversarial systems are responsive.

Notions of Simulation-Based Security. We can now define strong simulatability; other equivalent security notions, such as (dummy) UC, can be defined in a similar way.

Definition 1. Let \mathcal{P} and \mathcal{F} be protocol systems with the same I/O interface, the real and the ideal protocol, respectively. Then, \mathcal{P} realizes \mathcal{F} ($\mathcal{P} \leq_R \mathcal{F}$) if there exists an adversarial system \mathcal{S} (a simulator or an ideal adversary) such that the systems \mathcal{P} and $\mathcal{S} \mid \mathcal{F}$ have the same external interface and for all environmental systems \mathcal{E} , connecting only to the external interface of \mathcal{P} (and hence, $\mathcal{S} \mid \mathcal{F}$), it holds true that $\mathcal{E} \mid \mathcal{P} \equiv \mathcal{E} \mid \mathcal{S} \mid \mathcal{F}$.²

Composition Theorems. The IITM model provides several composition theorems. One theorem (see Theorem 1 below) handles concurrent composition of a fixed number of protocol systems. Other theorems guarantee secure composition of an unbounded number of copies of a protocol system.

Theorem 1. Let $\mathcal{P}_1, \mathcal{P}_2, \mathcal{F}_1, \mathcal{F}_2$ be protocol systems such that \mathcal{P}_1 and \mathcal{P}_2 as well as \mathcal{F}_1 and \mathcal{F}_2 only connect via their

²Note that strong simulatability omits the adversary in the real world as he can be subsumed by the environment.

I/O interfaces with each other and $\mathcal{P}_i \leq_R \mathcal{F}_i$, for $i \in \{1, 2\}$. Then, $\mathcal{P}_1 | \mathcal{P}_2 \leq_R \mathcal{F}_1 | \mathcal{F}_2$.

Other composition theorems provided by the IITM model can be found in [11], [16]. These theorems allow one to analyze a single session of a protocol in isolation in order to conclude security of an unbounded number of sessions. All composition theorems of the IITM can be combined and applied iteratively to construct more and more complex systems.

III. IDEAL FUNCTIONALITY FOR CRYPTOGRAPHIC PRIMITIVES

We now present our ideal functionality $\mathcal{F}_{\text{crypto}}$ for cryptographic primitives. As already mentioned in the introduction, a higher-level protocol \mathcal{P} can use $\mathcal{F}_{\text{crypto}}$ for its cryptographic operations. Then, in order to show that $\mathcal{P} | \mathcal{F}_{\text{crypto}} \leq_R \mathcal{F}$, i.e., that \mathcal{P} (using $\mathcal{F}_{\text{crypto}}$ for its cryptographic operations) realizes some ideal functionality \mathcal{F} (e.g., a key exchange functionality), one can argue on a purely information theoretic level, without resorting to reductions or hybrid arguments (at least for those primitives supported by $\mathcal{F}_{\text{crypto}}$). For example, $\mathcal{F}_{\text{crypto}}$ guarantees that only the (honest) owner of a Diffie-Hellman key can get access to keys that are derived from it, and only parties with access to these keys can, e.g., create a MAC with such keys. In all other cryptographic approaches for security protocols, one has to reduce these properties to the security assumptions for Diffie-Hellman key exchange, key derivation, and MAC schemes. Once $\mathcal{P} | \mathcal{F}_{\text{crypto}} \leq_R \mathcal{F}$ has been proven, using the composition theorems of the IITM model one can replace $\mathcal{F}_{\text{crypto}}$ with its realization $\mathcal{P}_{\text{crypto}}$ (see Section IV) by which the ideal operations provided by $\mathcal{F}_{\text{crypto}}$ are replaced by the real counterparts.

As mentioned in the introduction, in [19] a first version of $\mathcal{F}_{\text{crypto}}$ was proposed, which, however, does not support DH key exchange, a fundamental primitive for most real-world key exchange protocols. We also improve $\mathcal{F}_{\text{crypto}}$ in various other ways in order to overcome shortcomings of the previous version, as discussed below. Our extension of $\mathcal{F}_{\text{crypto}}$, in particular the treatment of DH key exchange, is non-trivial and needs care in order for it to be widely usable and realizable. In the following, we first recall the version of $\mathcal{F}_{\text{crypto}}$ from [19] and then present our extension.

A. The ideal functionality $\mathcal{F}_{\text{crypto}}$

On a high level, the ideal functionality $\mathcal{F}_{\text{crypto}}$ allows its users to perform the following operations in an ideal way: i) generate symmetric keys, including pre-shared keys, ii) generate public/private keys, iii) derive symmetric keys from other symmetric keys, iv) encrypt and decrypt messages and ciphertexts, respectively (public-key encryption and both unauthenticated and authenticated symmetric encryption are supported), v) compute and verify MACs and digital signatures, and vi) generate fresh nonces. All symmetric and

public keys can be part of plaintexts to be encrypted under other symmetric and public keys. Derived keys can be used just as freshly generated symmetric keys.

Formally, the ideal functionality $\mathcal{F}_{\text{crypto}}$ is a machine with n I/O tapes, representing different roles in a higher level protocol, and a network tape for communicating with the adversary. In runs of a system which contains $\mathcal{F}_{\text{crypto}}$ there will always be one instance of $\mathcal{F}_{\text{crypto}}$ only. This instance handles all requests.

A user of $\mathcal{F}_{\text{crypto}}$ is identified by a tuple $(pid, lsid, r)$, where r is the role/tape which connects the user to $\mathcal{F}_{\text{crypto}}$, pid is a party identifier (PID), and $lsid$ is a local session identifier (local SID). The local session ID is chosen and managed by higher level protocols and not further interpreted by $\mathcal{F}_{\text{crypto}}$. For example, it could be some session identifier that was established during a protocol run. All messages on I/O tapes are prefixed with $(pid, lsid)$ so $\mathcal{F}_{\text{crypto}}$ can identify the user who sent/receives a message.

Users of $\mathcal{F}_{\text{crypto}}$, and its realization, do not get their hands on the actual (private) keys but rather get pointers to such keys which can then be used to perform several cryptographic operations (see below).

The adversary can statically corrupt asymmetric (signing/encryption) keys,³ i.e., he can corrupt them before they are used for the first time but not afterwards. The corruption status of asymmetric keys determines whether operations with these keys are performed ideally or without ideal security guarantees. Similarly, the adversary can statically corrupt symmetric keys when they are generated or, in the case of pre-shared keys, when they are retrieved for the first time. In the case of symmetric keys, the functionality keeps track of whether a key might be *known* to the adversary/environment (e.g., because it was explicitly corrupted or because it was encrypted with a corrupted key). For this purpose, $\mathcal{F}_{\text{crypto}}$ maintains a set Keys of all symmetric keys and a set $\text{Keys}_{\text{known}} \subseteq \text{Keys}$ which contains all keys that might be known to the environment. The known/unknown status of symmetric keys is then used to determine whether symmetric operations are performed ideally or without ideal security guarantees. In the following, we will call a key *known* if it is in $\text{Keys}_{\text{known}}$ and *unknown* if it is in $\text{Keys} \setminus \text{Keys}_{\text{known}}$.

Symmetric keys in $\mathcal{F}_{\text{crypto}}$ are equipped with a key type that determines their usage. That is, a key k is of the form (t, k') where k' is the actual bit string used in algorithms while t is the key type. Keys of type *pre-key* are used to derive other keys, keys of type *unauthenc-key* and *authenc-key* are used for (un)authenticated encryption and decryption, and keys of type *mac-key* are used to create and verify MACs. This models the practice of using keys for a single purpose only.

³In our extension of $\mathcal{F}_{\text{crypto}}$, corruption of asymmetric signing keys is dynamic. That is, the adversary can corrupt signing keys at any point in time.

The ideal functionality $\mathcal{F}_{\text{crypto}}$ is parameterized with a leakage algorithm \mathcal{L} that is used to determine the information that is leaked when a plaintext x is encrypted ideally. For example $\mathcal{L}(x, 1^n) = 1^{|x|}$ can be used to leak exactly the length of x . We call such a leakage algorithm *length preserving*. The adversary is supposed to provide algorithms for authenticated and unauthenticated symmetric encryption, MACing, public key encryption, and signing. The adversary also provides the actual bit strings of all keys generated by $\mathcal{F}_{\text{crypto}}$. The functionality $\mathcal{F}_{\text{crypto}}$ ensures only that a new unknown symmetric key k is fresh (i.e., $k \notin \text{Keys}$) and prevents key guessing of unknown keys when receiving a new known key k (i.e., $k \notin \text{Keys} \setminus \text{Keys}_{\text{known}}$). Note that, as the adversary provides the keys, he knows the actual value of symmetric keys that are marked as unknown in $\mathcal{F}_{\text{crypto}}$. This is not a contradiction as the known/unknown status determines only whether operations are performed ideally; of course, in the realization a key that is marked unknown will indeed be unknown to the environment.

The functionality $\mathcal{F}_{\text{crypto}}$ offers the following list of commands to a user ($pid, lsid, r$) (see [19] for a detailed definition of every command):

- **Generating fresh, symmetric keys** [(New, t)]. A user can generate a new symmetric key of type t .
- **Establishing pre-shared keys** [(GetPSK, $t, name$)]. A user can ask for a pointer to a pre-shared symmetric key of type t , which can be used for modeling setup assumptions. If another user creates a key with the same input $name$, then this means that the two users share the created key.
- **Store** [(Store, t, k)]. A user can manually store a (known) key k of type t in $\mathcal{F}_{\text{crypto}}$.
- **Retrieve** [(Retrieve, ptr)]. A user can retrieve the key k a pointer ptr refers to, by which k is marked as known.
- **Equality test** [(Equal?, ptr, ptr')]. A user can test whether two of her pointers refer to the same key (same type and same bit string).
- **Public key requests** [(GetPubKeyPKE, p') or (GetPubKeySig, p')]. A user can ask for the public encryption/verification key, if any, of another party p' .
- **Key derivation** [(Derive, ptr, t', s)]. A user can derive a new symmetric key of type t' from salt s and key k of type ptr pre-key to which ptr points.
- **Encryption/decryption under symmetric keys** [(Enc, ptr, x) and (Dec, ptr, y)]. A user can encrypt a plaintext x and decrypt a ciphertext y using a key k of type $t \in \{\text{unauthenc-key}, \text{authenc-key}\}$ to which ptr points. The plaintext x may contain (pointers to) symmetric keys. As the result of the decryption of y , a user may learn symmetric keys. The exact operations depend, among others, on whether or not k is known. For example, if k is unknown, encryption is ideal, i.e., a ciphertext is produced which depends on $\mathcal{L}(x, 1^n)$ only.
- **Encryption and Decryption under public keys**

[(PKEnc, p', pk, x) and (PKDec, y)]. Asymmetric encryption/decryption works just as symmetric encryption/decryption, with the main difference being that the encryption command takes as input the PID p' and the public key pk of the intended recipient.

- **Creating and verifying MACs** [(Mac, ptr, x) and (MacVerify, ptr, x, σ)]. A user can create a MAC for or verify a MAC σ on a message x with key k of type mac-key to which ptr points.
- **Creating and verifying signatures** [(Sign, x) and (SigVerify, p', pk, x, σ)]. A user can create or verify a signature σ on a message x using his own private signing key or the public verification key pk of party p' , respectively.
- **Generating fresh nonces** [(NewNonce)]. A user can ask for a fresh nonce that does not collide with any previously generated nonces.
- **Corruption status request.** A user can ask whether one of her symmetric keys, or a public key of some party p' was corrupted by the adversary. This is used for modeling corruption: the environment can make sure that the corruption status of a key is the same in the real and ideal worlds.

B. Diffie-Hellman KE in $\mathcal{F}_{\text{crypto}}$

We now present our extension to $\mathcal{F}_{\text{crypto}}$ that supports Diffie-Hellman key exchange. On a high level, the extension lets users generate secret Diffie-Hellman exponents (e) and the corresponding public key shares (g^e), called *DH shares* in what follows. Exponents can be combined with arbitrary DH shares, not necessarily generated by $\mathcal{F}_{\text{crypto}}$, to produce a new symmetric key. If an exponent is combined with a DH share created by $\mathcal{F}_{\text{crypto}}$, then the resulting key will only be accessible by the owners of the two exponents that were used to create the key. The resulting key can then be used to derive other keys, e.g., for encryption or MACing. Whether or not this key derivation is performed ideally depends on several factors, such as whether any of the exponents is known to the environment/adversary (see below). Furthermore, $\mathcal{F}_{\text{crypto}}$ guarantees that new exponents/DH shares are fresh, i.e., no other user has access to the same exponent and no keys were already created from the share.

Before we describe our extension in detail, we first have to explain how we use restricting messages (cf. Section II). There are many situations where $\mathcal{F}_{\text{crypto}}$ needs to retrieve some information from the adversary, such as cryptographic algorithms or values of fresh keys. The adversary might use such requests to interfere with the run of $\mathcal{F}_{\text{crypto}}$ in an unintended way by, e.g., never responding to some of the requests. Importantly, such attacks do not relate to anything in reality: $\mathcal{F}_{\text{crypto}}$ models local computations that always succeed in reality. Our extension of $\mathcal{F}_{\text{crypto}}$ leverages the power of restricting messages to guarantee that an adversary cannot interfere with local computations, while still being able to

provide cryptographic values to $\mathcal{F}_{\text{crypto}}$. In the following, for brevity, we will say that a message m is restricting when we mean that the message $(\text{Respond}, \perp, m)$ is sent on the network. Recall from Section II that an environment has to respond to such a message immediately. We will implicitly assume that $\mathcal{F}_{\text{crypto}}$ repeats these messages until an expected response is received (e.g., when the response needs to be a value within a certain range).

We can now detail our extension. Formally, we parameterize $\mathcal{F}_{\text{crypto}}$ with a $\text{GroupGen}(1^\eta)$ algorithm that is used to generate the Diffie-Hellman group. This algorithm takes as input the current security parameter η , runs in polynomial time in η (except with negligible probability), and outputs a description (G, n, g) of a cyclic group G where $|G| = n$ and g is a generator of G . We require that it is possible in polynomial time (in η) to check whether a bit string encodes a group member of such a group, and that the group operation is efficiently computable.

Diffie-Hellman exponents are modeled analogously to keys in $\mathcal{F}_{\text{crypto}}$. That is, a user gets pointers to her exponents, never the actual exponent, and can use these pointers to perform Diffie-Hellman key exchange. However, users *do* get the DH share g^e belonging to an exponent e . The actual values of exponents are stored in two sets, Exp and $\text{Exp}_{\text{known}} \subseteq \text{Exp}$. An exponent in $\text{Exp}_{\text{known}}$ is called *known*, while an exponent in $\text{Exp} \setminus \text{Exp}_{\text{known}}$ is called *unknown*. The known/unknown status of exponents is used to determine whether keys created from them are considered known/unknown. Just as for keys, the environment provides the actual values of exponents, even if they are considered unknown. Of course, an exponent that is marked unknown in $\mathcal{F}_{\text{crypto}}$ will in fact be unknown to the environment in the realization $\mathcal{P}_{\text{crypto}}$. $\mathcal{F}_{\text{crypto}}$ prevents exponent collisions (i.e., if a new unknown exponent e is created, then $e \notin \text{Exp}$) and exponent guessing (i.e., if a new known exponent e is created, then $e \notin \text{Exp} \setminus \text{Exp}_{\text{known}}$). Additionally, $\mathcal{F}_{\text{crypto}}$ also maintains a set BlockedElements of blocked DH shares that contains group elements h that may not be generated when a new exponent e is created, i.e., $g^e \neq h$. In particular, this set contains all DH shares that have been used to create a Diffie-Hellman key (see Section III-C for an explanation).

We add another symmetric key type dh-key to $\mathcal{F}_{\text{crypto}}$ which represents Diffie-Hellman keys. Keys of this type may only be generated via a new GenDHKey command (see below) or be inserted into $\mathcal{F}_{\text{crypto}}$ via the existing Store command; they may not be created by any other commands. These keys can be used to derive new symmetric keys of arbitrary types, but they may not be used for encryption or creating MACs directly. Furthermore, just as all other key types, they can be encrypted as part of a plaintext.

Upon the first activation of $\mathcal{F}_{\text{crypto}}$, we now let $\mathcal{F}_{\text{crypto}}$ execute $\text{GroupGen}(1^\eta)$ and store the generated group (G, n, g) . Then, $\mathcal{F}_{\text{crypto}}$ sends both the group (G, n, g) and a request for cryptographic algorithms to the adversary via a restricting

message. When this initialization is complete, $\mathcal{F}_{\text{crypto}}$ either continues to process the original message that activated it (if the first message was received on an I/O tape) or returns control to the adversary (if the first message was received on a network tape).

Our extension of $\mathcal{F}_{\text{crypto}}$ provides the following additional commands to a user $(pid, lsid, r)$ on the I/O interface:

- **Get generated group** $[(\text{GetDHGroup})]$. The user can request the group (G, n, g) that was generated by $\mathcal{F}_{\text{crypto}}$ during initialization. $\mathcal{F}_{\text{crypto}}$ responds by sending $(\text{DHGroup}, (G, n, g))$ to the user.
- **Generate a fresh exponent** $[(\text{GenExp})]$. The user can request a pointer to a new unknown exponent e . This request is forwarded to the adversary via a restricting message, who is supposed to provide an exponent $e \in \{1, \dots, n\}$. The functionality $\mathcal{F}_{\text{crypto}}$ then ensures that this exponent e is fresh, i.e., it does not collide with existing exponents ($e \notin \text{Exp}$), and that g^e is not blocked from being generated (i.e., $g^e \notin \text{BlockedElements}$). If the freshness check fails, $\mathcal{F}_{\text{crypto}}$ asks the adversary again for another e until the check succeeds. Then, $\mathcal{F}_{\text{crypto}}$ adds e to Exp , stores a pointer ptr for user $(pid, lsid, r)$ pointing to e , and returns $(\text{ExpPointer}, ptr, g^e)$ to the user.
- **Mark group element as used** $[(\text{BlockGroupElement}, h)]$. The user can instruct $\mathcal{F}_{\text{crypto}}$ to manually block a group element h from being generated during a GenExp request. This is useful for higher level protocols to ensure that, if they receive some DH share h , no future GenExp request will output the same DH share even if h was not originally created by $\mathcal{F}_{\text{crypto}}$. Upon receiving this command, $\mathcal{F}_{\text{crypto}}$ checks that h is a valid group element and, if so, adds it to BlockedElements . In any case, $\mathcal{F}_{\text{crypto}}$ returns OK. See Section III-C for a discussion of this command.
- **Retrieve an exponent** $[(\text{RetrieveExp}, ptr)]$. The user can retrieve the exponent e a pointer ptr refers to. In this case, $\mathcal{F}_{\text{crypto}}$ adds e to $\text{Exp}_{\text{known}}$ and outputs $(\text{Exponent}, e)$ to the user.
- **Store an exponent** $[(\text{StoreExp}, e)]$. The user can also insert a new (known) exponent $e \in \{1, \dots, n\}$ into $\mathcal{F}_{\text{crypto}}$. Upon receiving this request, $\mathcal{F}_{\text{crypto}}$ prevents guessing of unknown exponents by ensuring that $e \notin \text{Exp} \setminus \text{Exp}_{\text{known}}$. If the check succeeds, e is added to $\text{Exp}_{\text{known}}$, a new pointer ptr for this exponent is created, and $(\text{ExpPointer}, ptr)$ is returned to the user. If the check fails, $(\text{ExpPointer}, \perp)$ is returned to the user.
- **Generate a new Diffie-Hellman key** $[(\text{GenDHKey}, ptr, h)]$. A user can ask $\mathcal{F}_{\text{crypto}}$ to create a new key of type dh-key from some group element h and the exponent e to which ptr points. When receiving this request, $\mathcal{F}_{\text{crypto}}$ first ensures that h actually is a group element and returns $(\text{Pointer}, \perp)$ to the user otherwise. If the check succeeds, $\mathcal{F}_{\text{crypto}}$ adds h to the set BlockedElements

to ensure that h will not be output by future GenExp requests (cf. Section III-C for a discussion). Furthermore, if $h = g^e$, then e is marked as known, i.e., is added to the set $\text{Exp}_{\text{known}}$ (cf. Section III-C). A new pointer ptr' to the DH key is created as follows:

First, $\mathcal{F}_{\text{crypto}}$ checks whether a key has already been generated by the group elements g^e and h . If so, then the pointer ptr' is set to this key. Otherwise, a new key is generated as follows.

If h belongs to an unknown exponent (i.e., there exists $d \in \text{Exp} \setminus \text{Exp}_{\text{known}}$ such that $h = g^d$) and e is marked unknown, then the adversary is asked via a restricting message to provide a fresh unknown key $k \in G$ of type dh-key. Formally, this is done by sending the restricting message (`ProvideDHKey`, `unknown`, e , d) on the network.⁴ The functionality $\mathcal{F}_{\text{crypto}}$ ensures that k is fresh, i.e., $k \notin \text{Keys}$ (and keeps asking for a new k if this is not the case), and then sets the pointer ptr' to k .

If the checks regarding the exponents fail, i.e., there is no $d \in \text{Exp} \setminus \text{Exp}_{\text{known}}$ such that $h = g^d$ or e is marked known, then the adversary is asked via a restricting message to provide a known key $k \in G$ of type dh-key. Formally, this is done by sending the restricting message (`ProvideDHKey`, `known`, e , h) on the network. The functionality $\mathcal{F}_{\text{crypto}}$ prevents key-guessing of unknown keys, i.e., if $k \in \text{Keys} \setminus \text{Keys}_{\text{known}}$, the functionality asks for another key. The pointer ptr' is then set to k . Furthermore, if there is no $d \in \text{Exp}$ such that $h = g^d$, then the exponent e is marked known by adding it to $\text{Exp}_{\text{known}}$ even if it was unknown before (we explain this in the remarks below).

In any case, $\mathcal{F}_{\text{crypto}}$ records that g^e and h have been used to create a key k and returns (`Pointer`, ptr') to the user.

In addition to these commands, we improve the overall expressivity and usability of $\mathcal{F}_{\text{crypto}}$ as follows:

- In [19], the adversary was allowed to corrupt a fresh key generated via the `New` command. As this command models a local computation performed by honest parties, we remove this ability. Keys generated by this command are now always uncorrupted and thus unknown.
- Every time $\mathcal{F}_{\text{crypto}}$ adds a symmetric key k to $\text{Keys}_{\text{known}}$, it sends a restricting message (`AddedKnownKey`, k) to the adversary and waits for any response on the network before continuing. This makes explicit that $\mathcal{F}_{\text{crypto}}$ does not provide any guarantees on the secrecy of actual values or the status of keys. As the adversary is already asked to provide unknown keys, there is no need to also leak them. While this change is not necessary for realizing $\mathcal{F}_{\text{crypto}}$ (see Section IV), it reduces the burden imposed

⁴We note that it is important to tell the adversary the known/unknown status for our realization as this determines whether our simulator responds with g^{ed} or g^c , $c \leftarrow_{\mathcal{S}} \{1, \dots, n\}$. Also note that the adversary knows the actual values of e and d anyway, so there is no security loss by directly sending these values on the network.

on simulators when using $\mathcal{F}_{\text{crypto}}$ as part of a higher-level protocol.

- As mentioned in Section III-A, the adversary may statically corrupt private keys. We now allow the adversary to corrupt signing keys adaptively, i.e., these keys can be corrupted by the adversary at any time.
- As mentioned above, our extension uses the power of restricting messages to guarantee that the environment/adversary cannot interfere with a higher level protocol while using $\mathcal{F}_{\text{crypto}}$ (for DH related and other operations) by defining all network messages to be restricting if they are sent while some operation is in progress.

C. Remarks

The ideal functionality $\mathcal{F}_{\text{crypto}}$ marks DH keys as unknown only if they were generated from two unknown exponents. In particular, if an unknown exponent e is used with a group element h which was not created by $\mathcal{F}_{\text{crypto}}$, then the resulting key is marked known and hence no security guarantees are given for this key. Otherwise, $\mathcal{F}_{\text{crypto}}$ would not be realizable: In a realization of $\mathcal{F}_{\text{crypto}}$, an environment might know the exponent d such that $h = g^d$, in which case it is trivial to compute the DH key g^{ed} . Hence, if $\mathcal{F}_{\text{crypto}}$ used such a key to derive other keys ideally, an environment could easily distinguish $\mathcal{F}_{\text{crypto}}$ and its realization $\mathcal{P}_{\text{crypto}}$.

We want to use the Decisional Diffie-Hellman (DDH) assumption for realizing $\mathcal{F}_{\text{crypto}}$. However, $\mathcal{F}_{\text{crypto}}$ provides operations that are not covered by the DDH experiment. To be more precise, the environment can use $\mathcal{F}_{\text{crypto}}$ to compute $(g^e)^e = g^{e^2}$ and h^e (where e is a secret exponent stored in $\mathcal{F}_{\text{crypto}}$ and h is an arbitrary group element not generated by $\mathcal{F}_{\text{crypto}}$). By the DDH assumption, we cannot guarantee that the environment does not learn anything about e itself or keys created with e in these cases. Indeed, if an adversary is able to calculate the function $f_a(h) \rightarrow h^a$ or the function $f'(g^e) \rightarrow g^{(e^2)}$ (where a is one of the exponents from the DDH experiment and h, g^e are arbitrary group elements) he can break the DDH assumption (see, e.g., [26], [27] for details). Thus, we have to consider e to be known in these cases.

The need for the `BlockedElements` set and the `BlockGroupElement` command might seem surprising at first: Typically, cryptographic libraries in real world protocols do not keep track of “seen” DH shares and then block them from being generated. However, this set and the corresponding command are necessary to lift an important property from the realization $\mathcal{P}_{\text{crypto}}$ to the case of the idealization. In the realization, it happens with negligible probability only that g^e for some fresh exponent e equals some DH share h which might already have been used to create a key. However, $\mathcal{F}_{\text{crypto}}$ allows the adversary to choose the actual value of e , i.e., he might choose the exponent such that $g^e = h$. To get the same guarantees as or even stronger guarantees than in the realization, $\mathcal{F}_{\text{crypto}}$ uses the

set `BlockedElements` to record all DH shares it has seen so far. With this set $\mathcal{F}_{\text{crypto}}$ makes sure that when creating a new exponent the corresponding DH share is “fresh”, i.e., does not belong to `BlockedElements`. The command `BlockGroupElement` allows higher level protocols to notify $\mathcal{F}_{\text{crypto}}$ about DH shares they obtain such that $\mathcal{F}_{\text{crypto}}$ can add these shares to `BlockedElements`. For example, when a responder in a DH-based key exchange protocol receives a DH share h , she would first add this share to $\mathcal{F}_{\text{crypto}}$ using the command `BlockGroupElement` and then create her own share. By this, $\mathcal{F}_{\text{crypto}}$ can make sure that the responder’s share is indeed fresh, and in particular, different from h . The responder can then use the `GenDHKey` command to derive a fresh DH key from h and her own DH share. We note that the `BlockedElements` set does not exist in $\mathcal{P}_{\text{crypto}}$ while the `BlockGroupElement` command in fact does nothing. Thus, after replacing $\mathcal{F}_{\text{crypto}}$ with $\mathcal{P}_{\text{crypto}}$, every call of the `BlockGroupElement` command can be omitted entirely, yielding a natural protocol implementation.

While we opted for a definition of $\mathcal{F}_{\text{crypto}}$ with a single DH key type for simplicity, it is trivial to extend $\mathcal{F}_{\text{crypto}}$ to multiple DH key types to model two or more groups that are used simultaneously. Such an extension uses one set `Exp` and `Expknown` and separate pointers to exponents for every DH key type. All results presented in the following carry over to this setting.

IV. REALIZATION

In this section, we construct a realization $\mathcal{P}_{\text{crypto}}$ of $\mathcal{F}_{\text{crypto}}$. This realization, which we describe in Section IV-A in detail, implements all operations of $\mathcal{F}_{\text{crypto}}$ via common cryptographic schemes in a natural and expected way. In Section IV-B, we then prove that $\mathcal{P}_{\text{crypto}}$ indeed realizes $\mathcal{F}_{\text{crypto}}$ under standard cryptographic assumptions. This proof is quite involved and includes several reductions and hybrid arguments, but due to the composition theorems this is a once and for all effort. As mentioned in the introduction, protocol designers can use $\mathcal{F}_{\text{crypto}}$ for their security analysis and then replace it with $\mathcal{P}_{\text{crypto}}$ without re-doing any proofs.

A. Formal Definition of $\mathcal{P}_{\text{crypto}}$

Formally, the machine $\mathcal{P}_{\text{crypto}}$ has the same network and I/O interface as $\mathcal{F}_{\text{crypto}}$. It is parameterized with three schemes Σ_{authenc} , $\Sigma_{\text{unauthenc}}$, Σ_{pub} for (un-)authenticated symmetric and public key encryption, a MAC scheme Σ_{MAC} , an algorithm `GroupGen`(η) with the same properties as for $\mathcal{F}_{\text{crypto}}$, and two families of pseudo-random functions (PRF) $F = \{F_\eta\}_{\eta \in \mathbb{N}}$ and $F' = \{F'_\eta\}_{\eta \in \mathbb{N}}$ that take as input a key and a salt and output a key (see our technical report [24] for formal definitions of these primitives). When activated for the first time by some message m , $\mathcal{P}_{\text{crypto}}$ initializes itself by executing `GroupGen` and storing the result before processing m . Just as $\mathcal{F}_{\text{crypto}}$, the machine $\mathcal{P}_{\text{crypto}}$ keeps track of symmetric key types and uses them to decide which

primitives may be executed with a given key (the family F is used for deriving keys from keys of type `pre-key`, while F' is used for key derivation from keys of type `dh-key`). The realization keeps track of the corruption status of keys in order to answer corruption status requests from the environment, but its behavior is independent of the corruption status otherwise. In particular, it does not maintain the sets `Keys`, `Keysknown`, `Exp`, `Expknown` and does not include any checks on freshness or key/exponent collisions.

We now give a detailed description of how each of the DH related commands is implemented in $\mathcal{P}_{\text{crypto}}$; see [19] for the remaining commands.

- **Get generated group** [`(GetDHGroup)`]. Outputs the group description (G, n, g) that was generated during the initialization of $\mathcal{P}_{\text{crypto}}$.
- **Generate a fresh exponent** [`(GenExp)`]. $\mathcal{P}_{\text{crypto}}$ chooses $e \xleftarrow{\$} \{1, \dots, n\}$, creates a pointer to e , and outputs (ptr, g^e) to the user.
- **Mark group element as used** [`(BlockGroupElement, h)`]. $\mathcal{P}_{\text{crypto}}$ returns OK.
- **Retrieve an exponent** [`(RetrieveExp, ptr)`]. $\mathcal{P}_{\text{crypto}}$ outputs the exponent e to which ptr points.
- **Store an exponent** [`(StoreExp, e)`]. $\mathcal{P}_{\text{crypto}}$ stores $e \in \{1, \dots, n\}$, creates a new pointer ptr for this exponent, and returns ptr to the user.
- **Generate a new Diffie-Hellman key** [`(GenDHKey, ptr, h)`]. $\mathcal{P}_{\text{crypto}}$ ensures that $h \in G$ and returns $(\text{Pointer}, \perp)$ to the user if this is not the case. Then, $\mathcal{P}_{\text{crypto}}$ calculates the key $k := h^e$ where e is the exponent to which ptr points to. A new pointer ptr' pointing to k is created and returned to the user.

The realization $\mathcal{P}_{\text{crypto}}$ also adopts all usability improvements from $\mathcal{F}_{\text{crypto}}$ described at the end of Section III-B.

B. Showing that $\mathcal{P}_{\text{crypto}}$ realizes $\mathcal{F}_{\text{crypto}}$

In this section, we state and prove our core theorem, namely, that $\mathcal{P}_{\text{crypto}}$ realizes $\mathcal{F}_{\text{crypto}}$. We want to use standard cryptographic assumptions for this, but these assumptions provide security only in a certain context. For example, standard assumptions for symmetric encryption do not provide any security guarantees in the presence of key cycles where a key is (indirectly) encrypted by itself. This is why reasonable higher level protocols generally avoid situations that are not covered by cryptographic assumptions; in contrast, environments in universal composability models are free to use $\mathcal{P}_{\text{crypto}}$ and $\mathcal{F}_{\text{crypto}}$ in any way they want and, in particular, they may create settings where the assumptions fail. In order to capture the expected use of $\mathcal{P}_{\text{crypto}}/\mathcal{F}_{\text{crypto}}$ as a subroutine of a reasonable higher level protocol, we thus slightly restrict environments such that they expose certain natural properties of higher level protocols. We note that this approach is established in the literature (see, e.g., [28]). The next paragraphs describe and discuss our restriction in more detail.

Recall that we want to use the DDH assumption in order to prove $\mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}_{\text{crypto}}$. The general idea is that the simulator in the proof of this statement will provide g^{ab} when asked for a known DH key, and g^c (for $c \xleftarrow{\$} \{1, \dots, n\}$) in case of an unknown DH key. However, this leads to the so-called commitment problem: Once the simulator has committed to g^c for an unknown key, neither a nor b may become known; otherwise the environment could calculate g^{ab} on its own and distinguish the real from the ideal world. We note that the commitment problem is not specific to our modeling of DH keys, but rather is a general issue in universal composability models (see, e.g., [29]). To address this problem, we restrict the environment (the higher level protocol that uses $\mathcal{F}_{\text{crypto}}$) to not cause the commitment problem. That is, once an unknown exponent e has been used to create an unknown DH key g^e , the environment may no longer manually retrieve e from $\mathcal{F}_{\text{crypto}}$, create a DH key from e and the corresponding DH share g^e (yielding g^{e^2}), or use e with a DH share h where $\mathcal{F}_{\text{crypto}}$ does not know the secret exponent of h . Observe, however, that most real-world protocols meant to achieve perfect forward secrecy fulfill this restriction: In such protocols, an exponent e is generated, used exactly once to generate a DH key, and then deleted from memory. Hence, after a key was created, the protocol will never access the exponent in any way, and thus, also never cause the commitment problem. For example, this holds true for all protocols analyzed in Section VI. It might be possible to relax these restrictions, enabling an analysis of protocols that re-use the same exponent, by using the non-standard PRF-ODH assumption⁵ [5], [6] instead of the DDH assumption. We want to explore a formulation of $\mathcal{F}_{\text{crypto}}$ based on this assumption in future work.

A similar commitment problem exists for encryption and key derivation. However, again most real-world protocols do not cause this problem (see also [19]). This leads us to the following formal restriction of environments:

We say that an environment \mathcal{E} *does not cause the commitment problem* (is *non-committing*), if the following happens with negligible probability only: i) in a run of $\mathcal{E} | \mathcal{F}_{\text{crypto}}$, after an unknown key k has been used to encrypt a message or derive a new key, k becomes known later on in the run, i.e., is marked known by $\mathcal{F}_{\text{crypto}}$, and ii) in a run of $\mathcal{E} | \mathcal{F}_{\text{crypto}}$, after an unknown exponent e or the corresponding group element g^e has been used to create an unknown DH key k , e becomes known later on in the run, i.e., is marked known by $\mathcal{F}_{\text{crypto}}$.

Besides the commitment problem, we also have to take care of key cycles. As mentioned, standard security definitions such as IND-CCA2, which we want to use for

⁵Informally, the PRF-ODH assumption states that, given a Diffie-Hellman key g^{ab} which is used to key a pseudo random function $f(g^{ab}, s)$, no adversary that knows g^a and g^b can distinguish a challenge output of the PRF from random, even when given access to an oracle $O(h, s) := f((h)^a, s)$ (where h is a group element and s is a salt).

our realization, do not provide any security in this case. Indeed, security in the presence of key cycles is usually not required: real-world protocols generally do not encrypt keys anymore once these keys have been used for the first time. Obviously, such protocols also do not produce key cycles. This observation leads to the following natural restriction of environments:

An environment \mathcal{E} is called *used-order respecting* if the following happens with negligible probability only: in a run of $\mathcal{E} | \mathcal{F}_{\text{crypto}}$ an unknown key k (i.e., k is marked unknown in $\mathcal{F}_{\text{crypto}}$) which has been used for encryption or key derivation at some point is encrypted itself by an unknown key k' used for the first time later than k .

We call an environment *well-behaved* if it is used-order respecting and does not cause the commitment problem. For such well-behaved environments, we can show that $\mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}_{\text{crypto}}$ if all cryptographic primitives fulfill the standard cryptographic assumptions. As explained above, many real world protocols fulfill the requirements of well-behaved environments, and hence, if they are analyzed using $\mathcal{F}_{\text{crypto}}$, one can replace $\mathcal{F}_{\text{crypto}}$ with its realization afterwards.

In the following theorem, formally, instead of considering a specific set of environments, we use a machine \mathcal{F}^* to manually enforce the properties of well-behaved environments for all environments. The machine \mathcal{F}^* is plugged between the environment and the I/O interface of $\mathcal{P}_{\text{crypto}}/\mathcal{F}_{\text{crypto}}$ and forwards all messages while checking that the conditions of well-behaved environments are fulfilled.⁶ If at some point one of the conditions is violated, instead of forwarding the current message, \mathcal{F}^* stops and blocks all future communication. We obtain the following theorem:

Theorem 2. *Let $\Sigma_{\text{unauth-enc}}$, $\Sigma_{\text{auth-enc}}$, Σ_{pub} be encryption schemes, Σ_{mac} be a MAC scheme, Σ_{sig} be a signature scheme, GroupGen be an algorithm as above, F be a family of pseudo-random functions, and F' be a family of pseudo-random functions for GroupGen. Let $\mathcal{P}_{\text{crypto}}$ be parameterized with these algorithms. Let $\mathcal{F}_{\text{crypto}}$ be parameterized with GroupGen and a leakage algorithm L which leaks exactly the length of a message. Then,*

$$\mathcal{F}^* | \mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}^* | \mathcal{F}_{\text{crypto}}$$

if $\Sigma_{\text{unauth-enc}}$ and Σ_{pub} are IND-CCA2 secure, $\Sigma_{\text{auth-enc}}$ is IND-CPA and INT-CTXT secure, Σ_{mac} and Σ_{sig} are UF-CMA secure, GroupGen always outputs groups with $n \geq 2$ and such that the DDH assumption holds true for GroupGen.⁷

⁶Note that this can be done by observing the I/O traffic and asking $\mathcal{F}_{\text{crypto}}$ about the corruption status of keys.

⁷We refer the reader to our technical report [24] for the formal definitions of these security notions. We have to require $n \geq 2$ because the trivial group which contains only the neutral element fulfills the DDH assumption, but is not suitable for realizing $\mathcal{F}_{\text{crypto}}$. In particular, collisions of randomly chosen exponents do not happen with a negligible probability if there is only one element.

As mentioned, the proof of this theorem is quite involved. It consists of a series of hybrid systems where we replace parts of $\mathcal{P}_{\text{crypto}}$ with the ideal versions used in $\mathcal{F}_{\text{crypto}}$ and then show that no environment can distinguish these replacements. Each of these steps involves several reductions and hybrid arguments itself. In particular, some of these reductions are intertwined with each other, as, e.g., the security of symmetric encryption and key derivation rely on each other. We provide a proof sketch in Appendix A; a full proof with all details can be found in our technical report [24].

V. IDEAL FUNCTIONALITIES FOR KEY EXCHANGE WITH KEY USABILITY

In this section, we present our ideal functionalities for key exchange, one functionality for mutual authentication, denoted by $\mathcal{F}_{\text{key-use}}^{\text{MA}}$, and one for unilateral authentication, $\mathcal{F}_{\text{key-use}}^{\text{UA}}$. These functionalities are of general interest and should be widely applicable. In Section VI, we use them in our case studies. In the following, we first present $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and then describe how $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ differs.

The functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$. The ideal functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ is inspired by an ideal key exchange functionality from [11], but has important differences, which among others makes it more widely applicable (see the comparison at the end of this section). In particular, neither unilateral authentication nor perfect forward secrecy were considered in [11].

Similar to other exchange functionalities (e.g., [30]), $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ guarantees that an uncorrupted instance that outputs a session key is in a session with an instance of its intended communication partner and only uncorrupted instances from the same session will have access to the session key. However, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ (and also $\mathcal{F}_{\text{key-use}}^{\text{UA}}$) has several features that distinguish it from key exchange functionalities typically considered in the literature.

First, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ never directly outputs session keys to users. Instead it provides a user with a pointer and allows the user to perform ideal cryptographic operations with it (among others, symmetric encryption, MACing, deriving new keys from the session key which can then be used further). This is an important feature as higher level protocols that use $\mathcal{F}_{\text{key-use}}^{\text{MA}}$, such as secure channel protocols, can use the session key still in an ideal way, which simplifies the analysis of higher level protocols and avoids reduction proofs.

Second, unlike most other formulations of key exchange functionalities in the literature, the above feature also makes it possible to realize $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ by key exchange protocols that use the session key during the key exchange. Most key exchange functionalities simply output a session key that was chosen uniformly at random, and thus, a realization must ensure that the session key is indistinguishable from a random one. However, this is not the case if the key was used during the actual key exchange, e.g., to encrypt a message, as then the environment can check whether the key that

is output after a successful key exchange can decrypt said message. In contrast, our functionality does not output the session key but only gives access to idealized cryptographic operations. As long as a key exchange protocol ensures separate domains of messages that are, e.g., encrypted with the session key during and after the key establishment phase, it can realize $\mathcal{F}_{\text{key-use}}^{\text{MA}}$.

Third, almost all formulations of functionalities (including key exchange functionalities) in the universal composability literature use so-called pre-established session IDs [11]: users somehow, outside of the protocol, agree on a (global) unique session ID and then use that session ID to access the same ideal functionality. As argued in [11], this hinders the faithful analysis of real-world protocols where such global session IDs are not a priori available; session IDs are often rather implicitly established during the protocol run. In fact, as illustrated in [11], an *insecure* key establishment protocol can be transformed into a *secure* one by assuming that global session IDs have been established prior to the actual protocol run. Therefore, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ does not rely on pre-established session IDs. Instead, just as $\mathcal{F}_{\text{crypto}}$, it uses local session IDs that are chosen and managed by the higher level instances. Local sessions (of an initiator and a responder) are combined by the adversary/simulator into a global session sharing one key during the protocol run.

Formally, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ is a machine that has two I/O tapes t_I and t_R (initiator and responder role), one network tape, and two I/O tapes t'_I and t'_R that connect to $\mathcal{F}_{\text{crypto}}$, which is used as a subroutine by $\mathcal{F}_{\text{key-use}}^{\text{MA}}$. $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ is parameterized with a symmetric key type $t_{\text{key}} \in \{\text{pre-key}, \text{unauthenc-key}, \text{authenc-key}, \text{mac-key}\}$ which determines the type of the keys that are output after a successful key exchange. Similarly to $\mathcal{F}_{\text{crypto}}$, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ handles all (local) sessions for all users. Messages from/to any I/O tape are expected to be prefixed with $(pid, lsid)$ where $pid \in \{0, 1\}^*$ is a party ID and $lsid \in \{0, 1\}^*$ is a local session ID managed by the higher level protocol. Thus, a user participating in a key exchange can be fully identified by $(pid, lsid, r)$, where $r \in \{I, R\}$ specifies the role of that user (and the tape she is using).

The functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ maintains a mapping state $: \{0, 1\}^* \rightarrow \{\perp, \text{started}, \text{inSession}, \text{exchangeFinished}, \text{sessionClosed}, \text{corrupted}\}$, initially set to \perp for every input which stores the current state for every user $(pid, lsid, r)$. The functionality also stores the PID of the intended partner of a user $(pid, lsid, r)$ via a mapping partner $: \{0, 1\}^* \rightarrow \{0, 1\}^*$. The functionality provides the following operations to higher level protocols:

- A user $(pid, lsid, r)$ with state $(pid, lsid, r) = \perp$ can start a key exchange by sending $m = (\text{InitKE}, pid', m')$, where pid' denotes the party ID of the intended partner and $m' \in \{0, 1\}^*$ is an arbitrary bit string which the realization might use in the key exchange protocol. Upon receiving this message, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ sets state $(pid, lsid, r) :=$

- started, sets $\text{partner}(pid, lsid, r) := pid'$, and forwards $(m, (pid, lsid, r))$ to the adversary.
- A user $(pid, lsid, r)$ with $\text{state}(pid, lsid, r) = \text{exchangeFinished}$ can use $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ to access symmetric operations of the subroutine $\mathcal{F}_{\text{crypto}}$. To be more precise, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ forwards the commands `New`, `Equal?`, `Enc`, `Dec`, `Mac`, `MacVerify`, and `Derive` to $\mathcal{F}_{\text{crypto}}$ on tape $t'_r, r \in \{I, R\}$. Upon receiving a response of $\mathcal{F}_{\text{crypto}}$, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ forwards this response to the user while internally keeping track of all pointers that the user has access to.
 - A user $(pid, lsid, r)$ with $\text{state}(pid, lsid, r) = \text{exchangeFinished}$ can close her session in $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ by sending `CloseSession`, by which she loses access to all of her keys. $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ sets $\text{state}(pid, lsid, r) := \text{sessionClosed}$, notifies the adversary with a restricting message `(CloseSession, (pid, lsid, r))`,⁸ and, after receiving any response from the adversary, returns `OK` to the user.

Corruption is modeled in such a way that the adversary may corrupt instances before a key exchange and after they have closed a session, but not while a session is active (see the discussion below). More precisely:

- The adversary can send `(Corrupt, (pid, lsid, r))` to corrupt a user where $\text{state}(pid, lsid, r) \in \{\perp, \text{sessionClosed}\}$. The user's state is updated accordingly.
- $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ forwards messages to/from corrupted users (in role r) between the I/O tape t_r and the network tape. It does not give the adversary access to the subroutine $\mathcal{F}_{\text{crypto}}$. This models perfect forward secrecy as the adversary should not gain access to any keys after the session is closed, even if he corrupts one of the parties.
- A user $(pid, lsid, r)$ may at any time ask for its corruption status by sending `Corrupt?`. $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ answers this request immediately without contacting the adversary. However, if $\text{state}(pid, lsid, r) = \perp$, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ first asks the adversary whether he wants to corrupt the user by sending him the restricting message `(CorruptUser?, (pid, lsid, r))`, expects a response b and, if $b = \text{true}$, sets $\text{state}(pid, lsid, r) := \text{corrupted}$. In any case, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ then returns the corruption status of $(pid, lsid, r)$ to the user.

The adversary can also declare two local sessions to belong to a global session and he decides when a user has successfully established a key:

- The adversary may send the message `(GroupSession, (pidI, lsidI), (pidR, lsidR))` to $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ if the following holds true: $\text{state}(pid_I, lsid_I, I) \in \{\text{started}, \text{corrupted}\}$, $\text{state}(pid_R, lsid_R, R) \in \{\text{started}, \text{corrupted}\}$, and both users are not yet part

⁸This models that one can usually observe whether some session is still active by monitoring the network of a party. Keeping this information secret is typically not a goal of secure key exchange protocols.

of a global session. The functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ then sets the state of uncorrupted users to `inSession` and stores that $(pid_I, lsid_I, I)$ and $(pid_R, lsid_R, R)$ are in the same global session. It then uses the `GetPSK` command of $\mathcal{F}_{\text{crypto}}$ to get pointers to an unknown key k of type t_{key} for the two users (if the received key is corrupted, then $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ asks for another key until its gets an uncorrupted one). Finally, it sends `OK` to the adversary. We note that, while we allow the adversary to pair an uncorrupted user with a corrupted one, the corrupted user will not get access to the session key in $\mathcal{F}_{\text{crypto}}$ (as already explained above).

- The adversary may send `(FinishKE, (pid, lsid, r))` where $\text{state}(pid, lsid, r) = \text{inSession}$ to complete the key exchange for an uncorrupted user. This message is accepted only if the user $(pid, lsid, r)$ is in a session with its intended partner, i.e., he is in a session with a user $(pid', lsid', r')$ such that $pid' = \text{partner}(pid, lsid, r)$. The functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ then sets $\text{state}(pid, lsid, r) := \text{exchangeFinished}$ and outputs `(Established, ptr)`, where ptr is the pointer to the previously established session key k .

The functionality $\mathcal{F}_{\text{key-use}}^{\text{UA}}$. The functionality $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ is similar but models unilateral authentication of the responder only. That is, it gives an initiator the same guarantees as $\mathcal{F}_{\text{key-use}}^{\text{MA}}$, while a responder may accept any connection without authentication. More formally, $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ differs from $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ as follows:

- Responders no longer indicate an intended session partner when starting a key exchange.
- The adversary may instruct $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ to output a key (`FinishKE`) for an uncorrupted instance of a responder that has already started a key exchange even if that instance is not yet part of a global session.
- If an honest responder instance is instructed to output a session key, no checks regarding the identity of the session partner are performed. Furthermore, unless the responder is in a global session with an honest initiator, the session key may be corrupted/known.
- Responder instances that have already output a key may still be mapped into a global session if i) they are not yet part of a global session and ii) their session key is uncorrupted/unknown. Their session partner will receive the same session key.

Discussion. The functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ assumes that responders know the identity of the initiator at the start of the key exchange. One could easily define a variant $\mathcal{F}_{\text{key-use}}^{\text{MA}'}$ where the responder learns the identity of the initiator only at the end of the key exchange. Note, however, that an environment for $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ is free to choose the expected identities of peers of the responder instances anyway, so it can always choose the identities at the start of a run appropriately.

The corruption model of both $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ requires

the corruption status of instances to stay unchanged during the key exchange. This is not strictly necessary for the ideal functionalities themselves (we could easily define them to model full dynamic corruption). But due to the commitment problem realizations typically have to adopt the same corruption model anyway. Therefore, we chose to also restrict the corruption model of $\mathcal{F}_{\text{key-use}}^{\text{MA}}/\mathcal{F}_{\text{key-use}}^{\text{UA}}$ as this makes these functionalities easier to use by higher level protocols. We note that this is not a strong restriction compared to full adaptive corruption, as session keys from key exchange protocols are usually very short lived, and hence, the window for corruption is small.

While $\mathcal{F}_{\text{key-use}}^{\text{MA}}/\mathcal{F}_{\text{key-use}}^{\text{UA}}$ are inspired by a functionality proposed in [11], the functionalities differ in several important aspects: As mentioned before, unilateral authentication is not considered in [11]. Also, $\mathcal{F}_{\text{key-use}}^{\text{MA}}/\mathcal{F}_{\text{key-use}}^{\text{UA}}$ model perfect forward secrecy, unlike the functionality in [11]. The functionality in [11] supports only symmetric encryption as an operation for higher-level protocols, and hence, is insufficient for modeling the cryptographic operations of most higher-level protocols. Furthermore, most common ideal functionalities for key exchange in the literature, including the functionality of [11] but also, e.g., the one from the CK model [30], impose overly strict security requirements. Thus, there are some reasonable protocols that cannot realize these functionalities. To be more precise, those functionalities require that the views of both parties are identical when the first party outputs its key. In other words, if, e.g., Alice wants to talk to Bob and outputs a session key, then the protocol must not only ensure that Alice’s session partner is indeed Bob, but also that Bob believes he is talking to Alice (even if Bob has not even finished his part of the protocol yet). However, this is not the case in protocols such as the SIGMA protocol family. While the initiator knows that she is talking to her intended communication partner when she outputs a key, the responder has not yet confirmed the identity of the initiator, and thus their views may differ. Even though these protocols cannot realize the functionality in [11] and the like, the SIGMA protocol family is still reasonable as this protocol family ensures that the responder learns the correct identity of the initiator before outputting her own session key (as we show in Section VI-B). By relaxing the requirements on establishing a global session and instead performing additional checks when a session key is output, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ allows for the analysis of a wider variety of protocols.

VI. CASE STUDIES

In this section, we carry out several case studies to illustrate the usefulness of our framework. We analyze one of the ISO 9798-3 protocols [21] and the SIGMA protocol with identity protection [22]. Both protocols are meant to provide mutually authenticated key exchange. We also analyze one mode of OPTLS [23] for unilaterally authenticated key




Figure 1. The ISO 9798-3 protocol for mutual authentication. At the end of the protocol, users share a key g^{xy} that is then used to derive a session key.

exchange that served as the basis for the key exchange protocol in TLS 1.3 draft-09 [31], and point out a subtle bug in the original game-based proof.

We show that these protocols realize $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and $\mathcal{F}_{\text{key-use}}^{\text{UA}}$, respectively. In our modeling of these protocols, we use $\mathcal{F}_{\text{crypto}}$ to perform all cryptographic operations. By Theorem 2, $\mathcal{F}_{\text{crypto}}$ can then be replaced by its realization $\mathcal{P}_{\text{crypto}}$ so that the protocols use the actual cryptographic primitives. Due to the use of $\mathcal{F}_{\text{crypto}}$, the proofs are quite simple as they rely on high level information theoretic arguments only; they do not need a single reduction, not even any probabilistic reasoning. At the same time, we obtain strong universal composability guarantees for the protocols. Moreover, the use of local session IDs in our framework allows for a faithful modeling of the protocols. As discussed at the beginning of Section V, other universal composability approaches impose pre-established (global) session IDs on the protocols, and hence, modify the protocols quite severely.

A. ISO protocol

The ISO 9798-3 [21] protocol for mutual authentication is depicted in Figure 1. It is based on Diffie-Hellman key exchange and uses signatures to ensure mutual authentication.

The modeling of the ISO protocol in our framework is straightforward. We use two machines M_I and M_R to model the initiator and responder role, respectively. These machines provide the same I/O interface as $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and each one has a network tape. They use $\mathcal{F}_{\text{crypto}}$ as a subroutine to perform all cryptographic operations. In every run of the protocol, there is one instance of M_I/M_R per user $(pid, lsid)$, with each instance executing the protocol according to Figure 1. As soon as an instance receives some DH share, it uses the `BlockGroupElement` command to ensure that $\mathcal{F}_{\text{crypto}}$ “knows” this share, and hence, fresh exponents do not collide with it.⁹ At the end of the protocol, instances create a DH key from g^x and g^y and use this to derive the session key of type `unauthenc-key`.¹⁰ They then output a pointer to that session key and subsequently provide the same interface

⁹As mentioned earlier, this operation can be omitted when $\mathcal{F}_{\text{crypto}}$ is replaced with its realization. The resulting protocol is a natural implementation of the ISO protocol.

¹⁰We could also have chosen any other symmetric key type supported by $\mathcal{F}_{\text{key-use}}^{\text{MA}}$. The security proof is independent of this.

as $\mathcal{F}_{\text{key-use}}^{\text{MA}}$, i.e., they allow a user to use $\mathcal{F}_{\text{crypto}}$ to perform (ideal) cryptographic operations with the session key.

Corruption of M_I/M_R is modeled analogously to $\mathcal{F}_{\text{key-use}}^{\text{MA}}$. That is, protocol participants might be corrupted by the adversary (by sending a special message) before the start of the protocol run or after a session has been closed, but not while a key exchange/session is active. While this is more restricted than full adaptive corruption, it is still a reasonable and meaningful modeling, as already discussed in Section V. Besides directly being corrupted by the adversary, an instance of M_I/M_R also considers itself corrupted (even though not directly controlled by the adversary) if its own signing key or the signing key of its intended peer is corrupted. This models that no security guarantees, and in particular no guarantees about authentication, can be given if the adversary has access to the long term secrets. Please refer to our technical report [24] for a detailed definition of the corruption behavior.

The following theorem states that the ISO protocol is a secure universally composable mutually authenticated key exchange protocol. As mentioned before, our modeling allows one to use session keys returned by this protocol to be used by higher level protocol in an ideal way.

Theorem 3. *Let M_I and M_R be machines modeling the ISO protocol as described above, let $\mathcal{F}_{\text{crypto}}$ and $\mathcal{F}'_{\text{crypto}}$ be two versions of the ideal crypto functionality with the same parameters, and let $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ be the ideal functionality for mutually authenticated key exchanges with parameter $t_{\text{key}} = \text{unauthenc-key}$. Then the following holds true:*

$$M_I | M_R | \mathcal{F}_{\text{crypto}} \leq_R \mathcal{F}_{\text{key-use}}^{\text{MA}} | \mathcal{F}'_{\text{crypto}}.$$

As mentioned before, the proof of this theorem does not require any reductions, not even probabilistic reasoning, which greatly simplifies the overall proof. We note that we directly show this theorem in the multi session setting. While there exists a single session theorem for local session IDs [11], in our case the analysis is already simple in the multi session setting.

Proof: In the following, we say that a party pid is corrupted if the signing key of party pid is corrupted. We call an instance $(pid, lsid, r)$ corrupted if it outputs `true` when asked for its corruption status by the environment, and we say that an instance $(pid, lsid, r)$ is explicitly corrupted if the adversary took control of this instance by sending the special `Corrupt` message.

We have to define a simulator \mathcal{S} and show that $\mathcal{E} | M_I | M_R | \mathcal{F}_{\text{crypto}} \equiv \mathcal{E} | \mathcal{S} | \mathcal{F}_{\text{key-use}}^{\text{MA}} | \mathcal{F}'_{\text{crypto}}$ for all environments $\mathcal{E} \in \text{Env}_R(M_I | M_R | \mathcal{F}_{\text{crypto}})$. The simulator \mathcal{S} internally simulates the protocol $M_I | M_R | \mathcal{F}_{\text{crypto}}$ and keeps the corruption statuses of user instances in $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and simulated instances of M_I/M_R synchronized. When \mathcal{S} has to initialize $\mathcal{F}_{\text{crypto}}$, \mathcal{S} first sends a message to $\mathcal{F}'_{\text{crypto}}$ to initialize it and receives a group (G, n, g) in response which is used for the

simulation of $\mathcal{F}_{\text{crypto}}$. \mathcal{S} then asks the environment for the cryptographic algorithms and forwards them to $\mathcal{F}'_{\text{crypto}}$.

If $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ indicates that a user $(pid, lsid, r)$ has started a key exchange, \mathcal{S} does the same in its internal simulation. If an uncorrupted initiator $(pid_I, lsid_I, I)$ accepts a group element g^y and outputs a pointer to a session key, then \mathcal{S} instructs $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ to create a session from $(pid_I, lsid_I, I)$ and the instance $(pid_R, lsid_R, R)$ that created the signature in the second protocol message. The subroutine $\mathcal{F}'_{\text{crypto}}$ of $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ will then ask \mathcal{S} to provide the value for the session key; \mathcal{S} provides the same value that is used in its simulation as session key. Finally, \mathcal{S} instructs $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ to output the session key pointer for $(pid_I, lsid_I, I)$. If an uncorrupted instance $(pid_R, lsid_R, R)$ outputs a pointer to a session key, \mathcal{S} instructs $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ to output the session key pointer for $(pid_R, lsid_R, R)$. While a session key is used, the simulator may be asked by $\mathcal{F}'_{\text{crypto}}$ to provide new unknown keys (e.g., when deriving keys). In this case, \mathcal{S} simulates the same operation in $\mathcal{F}_{\text{crypto}}$ and forwards the keys to $\mathcal{F}'_{\text{crypto}}$. If \mathcal{S} is notified that some instance $(pid, lsid, r)$ has closed its session, \mathcal{S} updates the internal simulation accordingly and responds with `OK`. \mathcal{S} uses the internal simulation to process inputs/outputs for/from corrupted instances.

We now show that \mathcal{S} is a good simulator. As explained in Section III-B, due to the use of restricting messages, we can conveniently assume that all operations performed by $\mathcal{F}_{\text{crypto}}$ are atomic, without any side effects on the machines M_I or M_R . This simplifies the overall proof.

First, observe that \mathcal{S} keeps the key sets of $\mathcal{F}_{\text{crypto}}$ and $\mathcal{F}'_{\text{crypto}}$ “synchronized”, i.e., the set of keys of $\mathcal{F}'_{\text{crypto}}$ is a subset of all keys of $\mathcal{F}_{\text{crypto}}$ and all keys in $\mathcal{F}'_{\text{crypto}}$ have the same known/unknown status in $\mathcal{F}_{\text{crypto}}$. This is easy to see, as the simulator provides all unknown keys for $\mathcal{F}'_{\text{crypto}}$ while it is not possible for the environment to insert any known keys. As both key sets are synchronized, $\mathcal{F}'_{\text{crypto}}$ will accept all keys that have been accepted by the internally simulated $\mathcal{F}_{\text{crypto}}$ and thus the environment cannot use the freshness check on new keys to distinguish real from ideal world.

The following argument is split into four cases, for which we argue that the simulation is perfect: Honest initiator instances during key establishment, honest responder instances during key establishment, honest instances after key establishment, and corrupted instances.

Let $(pid_I, lsid_I, I)$ be an uncorrupted instance of M_I that wants to establish a session with party pid' . It is easy to see that the simulator can perfectly simulate the behavior of such an instance up to the point when it outputs a key as the behavior does not depend on any data present in $\mathcal{F}'_{\text{crypto}}$. In particular, honest instances will use $\mathcal{F}_{\text{crypto}}$ only to create/verify signatures, and exchange Diffie-Hellman keys; both of these operations are unavailable in $\mathcal{F}'_{\text{crypto}}$ and thus can separately be simulated by \mathcal{S} .

We have to argue that \mathcal{S} finds an instance of a responder that can be paired with $(pid_I, lsid_I, I)$: If $(pid_I, lsid_I, I)$

outputs a session key pointer, then it must have accepted the second message of the ISO protocol and the signing key of its intended partner pid' must still be uncorrupted (otherwise, the protocol would block according to our modeling of corruption). Hence, there is some instance belonging to pid' , say $(pid', lsid', r')$, that has signed the message $m = (g^x, g^y, pid_I)$, where x is the secret exponent of $(pid_I, lsid_I, I)$ and y is the secret exponent of $(pid', lsid', r')$. This instance is uncorrupted: On the one hand, it cannot be explicitly corrupted by the adversary as the party pid' is still uncorrupted. On the other hand, as $(pid', lsid', r')$ considers pid_I to be the partner of the key exchange (which is acknowledged in the signature), we know that $(pid', lsid', r')$ also does not consider itself corrupted due to corrupted signing keys. Next, we argue that this instance is a responder, i.e., $r' = R$: If it were an initiator, then the signed message m would imply that this instance received and accepted the second protocol message containing a message $m' = (g^y, g^x, pid_I)$ signed by an uncorrupted instance of pid_I , where x is the secret exponent of the instance of pid_I . However, as x/g^x is created ideally, there is only one honest instance that would sign such a message, namely $(pid_I, lsid_I, I)$, which does not output any signatures before accepting the second message. This implies $r' = R$. We still have to show that $(pid', lsid', r')$ was not yet assigned to a session by \mathcal{S} : The simulator pairs (honest) responder instances with those (honest) initiator instances that accept the second message, but as x/g^x is unique, the only honest initiator instance that accepts this message is $(pid_I, lsid_I, I)$. Hence, we have that $(pid', lsid', r')$ is not yet part of a global session and can be paired with $(pid_I, lsid_I, I)$. Finally, observe that both x/g^x and y/g^y have been created ideally (with $x \neq y$) and thus the key derived from them will be considered unknown in $\mathcal{F}_{\text{crypto}}$. The simulator can provide the exact same key from the simulation to $\mathcal{F}'_{\text{crypto}}$ as the key sets are synchronized. Note in particular that only $(pid_I, lsid_I, I)$ and $(pid', lsid', r')$ can get a pointer to this key, which matches the behavior of $\mathcal{F}_{\text{key-use}}^{\text{MA}}$.

The remaining cases are similar. We provide them in Appendix B. ■

By Theorem 2, we can now replace $\mathcal{F}_{\text{crypto}}$ by its realization $\mathcal{P}_{\text{crypto}}$ which yields that the ISO protocol (when using the actual cryptographic operations) is a universally composable mutual authenticated key exchange protocol.

Corollary 1. *Let M_I, M_R as defined above, let $\mathcal{F}_{\text{crypto}}, \mathcal{P}_{\text{crypto}}$, and \mathcal{F}^* as in Theorem 2, in particular, we have that $\mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}_{\text{crypto}}$ and \mathcal{F}^* enforces well-behaved environments. Then the following holds true:*

$$\mathcal{F}^* | M_I | M_R | \mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}^* | \mathcal{F}_{\text{key-use}}^{\text{MA}} | \mathcal{F}_{\text{crypto}}.$$

Proof: This statement follows easily from Theorem 1, Theorem 2, Theorem 3, and transitivity of the \leq_R relation




Figure 2. The SIGMA protocol with identity protection. The keys k_e and k_m are derived from g^{ab} , where k_e is used to encrypt and k_m is used to mac messages during the key exchange. Another key k_s is also derived from g^{ab} and used as session key.

as well as the fact that the machines M_I and M_R constitute a well-behaved environment when combined with \mathcal{F}^* and any another environment \mathcal{E} : corrupted instances do not have access to unknown keys, so they cannot violate the well-behaved property. Uncorrupted instances during the key usage phase are well-behaved due to \mathcal{F}^* . Uncorrupted instances during the key establishment phase can violate the well-behaved property only by causing the commitment problem for Diffie-Hellman keys, i.e., set an unknown exponent to known after it was used to create an unknown key. This case does not occur as exponents are never accessed/used after one key has been created with them. ■

B. SIGMA Protocol

The SIGMA protocol with identity protection [22] is depicted in Figure 2. Unlike the ISO protocol, it uses the exchanged DH key to derive three other keys, two of which are used during the key exchange to ensure authentication and confidentiality of party IDs, while the third is used as session key.

We model the SIGMA protocol analogously to the ISO protocol. We use unauthenticated encryption to encrypt messages in the protocol; authenticated encryption is not necessary. The following theorem states that the SIGMA protocol is a secure universally composable mutually authenticated key exchange protocol.

Theorem 4. *Let M_I and M_R be the machines modeling the SIGMA protocol, let $\mathcal{F}_{\text{crypto}}$ and $\mathcal{F}'_{\text{crypto}}$ be two versions of the ideal crypto functionality with the same parameters, and let $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ be the ideal functionality for mutually authenticated key exchanges with parameter $t_{\text{key}} = \text{unauthenc-key}$. Then the following holds true:*

$$M_I | M_R | \mathcal{F}_{\text{crypto}} \leq_R \mathcal{F}_{\text{key-use}}^{\text{MA}} | \mathcal{F}'_{\text{crypto}}.$$

We provide the proof of this theorem in our technical report [24]. Again, it does not need any reductions or probabilistic reasoning. Just as for the ISO protocol, by Theorem IV-A we can replace $\mathcal{F}_{\text{crypto}}$ by its realization $\mathcal{P}_{\text{crypto}}$.

Corollary 2. *Let M_I, M_R as defined above, let $\mathcal{F}_{\text{crypto}}, \mathcal{P}_{\text{crypto}}$, and \mathcal{F}^* as in Theorem 2, in particular, we have*




Figure 3. The 1-RTT non-static mode of OPTLS. Both chello and shello are arbitrary bit strings that are exchanged during the protocol (they can be used to negotiate parameters for a higher level protocol). The message sfin is a MAC on the whole key exchange, i.e., $\text{sfin} = \text{MAC}_{k_m}(\text{chello}, g^x, \text{shello}, g^y, \text{SIG}_B(g^y))$. The keys k_e (for encryption), k_m (for MACing) and the session key k_s are derived from the DH key g^{xy} as shown in Figure 4.




Figure 4. Key derivation in the 1-RTT non-static mode of OPTLS. HKDF [32] is a key derivation function that takes as input a key (arrows on the left), context information (upper left arrows), and a salt (upper right arrows). It outputs a variable number of keys (bottom arrows).

that $\mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}_{\text{crypto}}$ and \mathcal{F}^* enforces well-behaved environments. Then the following holds true:

$$\mathcal{F}^* \mid M_I \mid M_R \mid \mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}^* \mid \mathcal{F}_{\text{key-use}}^{\text{MA}} \mid \mathcal{F}_{\text{crypto}}.$$

C. OPTLS

The OPTLS protocol family [23] specifies several key exchange protocols with unilateral authentication. It was built to meet the specific requirements of TLS 1.3 for key exchange; a slightly modified version was included in draft-09 of TLS 1.3 [31]. In Figure 3, we show the so-called non-static mode of OPTLS. Unlike the ISO and SIGMA protocols, OPTLS also specifies the exact key derivation procedure, which we depict in Figure 4.

We model OPTLS in the same way as the ISO and SIGMA protocols, but with the following changes: The machines M_I and M_R execute the protocol from Figure 3 to exchange a key. Instances of responders do not specify an intended session partner at the beginning (as the protocol does not authenticate the initiator to the responder) and thus also do not consider themselves to be corrupted if their session partner is corrupted. We use the optional bit string m' , which is part of the `InitKE` message expected by $\mathcal{F}_{\text{key-use}}^{\text{UA}}$, to provide instances of M_I with the chello message, and instances of M_R with the shello message.

We model HKDF via the `Derive` command of $\mathcal{F}_{\text{crypto}}$. As $\mathcal{F}_{\text{crypto}}$ provides a single argument for key derivation, we concatenate both context information and salt and use the resulting string as salt for $\mathcal{F}_{\text{crypto}}$. This models that HKDF should provide independent keys if either salt or context

information is changed. Another technical difference is that HKDF outputs a variable number of keys, while $\mathcal{F}_{\text{crypto}}$ outputs a single key for every salt. It is easy to extend $\mathcal{F}_{\text{crypto}}$ to also support deriving multiple keys from a single salt and then realize it with a secure variable length output PRF. Nevertheless, for simplicity, we use the current formulation of $\mathcal{F}_{\text{crypto}}$ and instead call the `Derive` command twice to obtain two keys. Formally, we use two different salts which are obtained by prefixing the original salt with 0 or 1, depending on whether the first or second key is derived.

Surprisingly, OPTLS does *not* realize $\mathcal{F}_{\text{key-use}}^{\text{UA}}$. To see this, consider the following setting: an honest initiator outputs a session key which was generated from its own DH share g^x and the responders DH share g^y . The responder instance that signed g^y might have received a different group element, say $h \neq g^x$, in the first protocol message. If h was not honestly generated by $\mathcal{F}_{\text{crypto}}$, then y will be marked known after the calculation of h^y because the DDH assumption does not guarantee that an attacker learns nothing from y in this case. As y is marked known, the key g^{xy} and all keys derived from it will also be marked known. Thus, we have no security guarantees for the MAC and an attacker can easily let the initiator instance accept, even though there is no instance of a responder that can be paired with it (the responder that signed g^y outputs a different session key).

We note that this is not a direct attack against the protocol but rather shows that assuming hardness of DDH and security of the PRF family is not sufficient to prove the security of this protocol mode. Indeed, we found that the original game-based security proof of this protocol from [23] is flawed: In the proof, where the authors use the same cryptographic assumptions, g^{xy} is replaced by $g^z, z \stackrel{\$}{\leftarrow} \{1, \dots, n\}$ during a hybrid argument (cf. game 2). The authors claim that this can be reduced to the DDH assumption. But a simulator in the reduction to DDH would have to simulate the game where g^x of the initiator and g^y of the responder are replaced with the challenges from the DDH game. Now, the simulator might have to calculate h^y (for some group element h) and derive keys from it, if the responder received h in its first message. This is impossible with just the DDH assumption as the simulator neither knows y nor has an oracle to compute h^y , i.e., he cannot simulate the game faithfully.

To fix this problem both in the original paper and in our setting, one can use stronger assumptions. For example, one could use the PRF-ODH assumption [5], [6], where the adversary additionally gets access to an oracle for calculating keys derived from h^y (where y is one of the secret exponents and h is provided by the adversary). As mentioned earlier, we leave a formulation of $\mathcal{F}_{\text{crypto}}$ based on the PRF-ODH assumption for future work. An alternative fix for this problem (again for both settings) is to have g^x signed as well, i.e., signing (g^x, g^y) as in the SIGMA protocol. This

allows for an analysis using the DDH assumption, as now the signature guarantees that the responder paired g^y with g^x only. The following theorem states that this variant is a secure universally composable unilaterally authenticated key exchange.

Theorem 5. *Let M_I and M_R be machines modeling the variant of the 1-RTT non-static mode of OPTLS that signs both g^x and g^y . Let $\mathcal{F}_{\text{crypto}}$ and $\mathcal{F}'_{\text{crypto}}$ be two versions of the ideal crypto functionality with the same parameters, and let $\mathcal{F}_{\text{key-use}}^{\text{UA}}$ be the ideal functionality for unilaterally authenticated key exchanges with parameter $t_{\text{key}} = \text{unauthenc-key}$. Then the following holds true:*

$$M_I \mid M_R \mid \mathcal{F}_{\text{crypto}} \leq_R \mathcal{F}_{\text{key-use}}^{\text{MA}} \mid \mathcal{F}'_{\text{crypto}}.$$

The proof of this theorem can be found in our technical report [24]. Again, it does not need any reductions or probabilistic reasoning. As before, we can again replace $\mathcal{F}_{\text{crypto}}$ by its realization $\mathcal{P}_{\text{crypto}}$.

Corollary 3. *Let M_I, M_R as defined above, let $\mathcal{F}_{\text{crypto}}, \mathcal{P}_{\text{crypto}}$, and \mathcal{F}^* as in Theorem 2, in particular, we have that $\mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}_{\text{crypto}}$ and \mathcal{F}^* enforces well-behaved environments. Then the following holds true:*

$$\mathcal{F}^* \mid M_I \mid M_R \mid \mathcal{P}_{\text{crypto}} \leq_R \mathcal{F}^* \mid \mathcal{F}_{\text{key-use}}^{\text{MA}} \mid \mathcal{F}_{\text{crypto}}.$$

VII. DISCUSSION AND RELATED WORK

There are several different approaches for analyzing security protocols, mainly symbolic, game-based, implementation-based, and universal composability approaches. All of these approaches have different advantages and shortcomings; there is no silver bullet, as can be seen, for example, by the fact that real-world protocols, such as TLS, have been studied in the literature using all of these approaches (often computer-aided), taking different views and making use of the specific merits thereof (see, e.g., [1]–[9]).

- Symbolic (Dolev-Yao-style) approaches abstract from low level cryptographic details in order to offer a very high degree of automation (see, e.g., [33]–[35]).
- Implementation-based analysis captures details of the actual implementations of protocols, which is very desirable, but of course also makes the analysis more involved (see, e.g., [36]–[39]).
- Game-based models are very expressive and flexible in defining security properties of a protocol (see, e.g., [40]–[42] and [43], [44] for tools). While they do not enjoy built in modularity, efforts have been made to improve the modularity provided by these models (see, e.g., [13], [14]).
- Universal composability approaches come with modularity built in and allow one to show that protocols are secure in arbitrary (polynomially bounded) environments (see, e.g. [15]–[18]). But due to the commitment problem, they

can be more limited in their corruption modeling. In some cases, instead of allowing for full adaptive corruption, one might have to model corruption in a more restricted, but still reasonable way (see also the discussions in Sections V and VI).

Our framework adds the feature of avoiding or limiting the need for tedious and error-prone reductions, while at the same time allowing to establish universally composable security guarantees. In particular, proofs are simplified and results can easily be re-used and built upon.

In the remainder of this section, we discuss closely related work in more detail. The works [11], [19] have been discussed in detail before already.

In [45], Canetti and Gajek abstract Diffie-Hellman key exchange via an ideal key encapsulation functionality. There are two key differences to our framework. First, unlike $\mathcal{F}_{\text{crypto}}$ and our key usability functionalities, the ideal key encapsulation functionality does not allow a user/higher-level protocol to use the exchanged key in an idealized way or to use it with other primitives, which entails reductions proofs. Second, a large class of protocols cannot be analyzed with their key encapsulation functionality: in order to prove the realization, they impose a very strong restriction on the environment/higher-level protocols, namely, an initiator may use her secret exponent a only with DH shares g^b that have been honestly generated by a responder. Many protocols, including all case studies considered in this paper, do not fulfill this requirement: If, for example, the responder is corrupted, then the environment may sign arbitrary DH shares that were not honestly generated. These DH shares will be accepted by the initiator, which violates the requirement.

Our case studies have not yet been faithfully analyzed in a universal composability setting (see [23], [41], [46] for game-based analyses). Variants of the ISO 9798-3 and SIGMA protocols have been analyzed in the UC model in [30], [45], [46]. These variants assume that protocol participants have already established a global, unique SID prior to running the actual protocol, and then either use different signing keys for every new session (which is unrealistic) or, if they re-use the same key across different sessions, prefix all signed messages with the SID. The latter is a so-called joint-state realization, which, however, results in a protocol that differs from the actual protocol. As illustrated in [11], such modifications can potentially create a secure protocol from an insecure one.

Moreover, the analysis of the (variant of the) SIGMA protocol in [46] needed a modified version of the ideal key exchange functionality \mathcal{F}_{ke} where initiators and responders *cannot* specify their intended peers. Our functionality $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ is the first that allows for proving security of the SIGMA protocol in a setting where the initiator and the responder can specify their intended peers.

VIII. CONCLUSION

In this paper, we have proposed an ideal functionality $\mathcal{F}_{\text{crypto}}$ that models various cryptographic primitives which can be combined with each other and can be used in an idealized way. Importantly, $\mathcal{F}_{\text{crypto}}$ supports Diffie-Hellman key exchange, a widely and extensively used primitive in real-world protocols. We also provided new functionalities, $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ and $\mathcal{F}_{\text{key-use}}^{\text{UA}}$, for ideal mutual and unilateral authenticated key exchange which lift the properties of $\mathcal{F}_{\text{crypto}}$ to the next protocol level. Notably, these functionalities allow for analyzing a wider range of key exchange protocols than traditional formulations of ideal key exchange functionalities.

Altogether, our approach gets rid of reductions and hybrid arguments for primitives that are supported by $\mathcal{F}_{\text{crypto}}$. Instead, proofs rely on simpler information theoretic arguments only, which facilitates proofs and makes it easier to uncover subtle problems that otherwise might get lost in sequences of reductions. At the same time, our approach offers very high modularity and strong universal composable security guarantees.

We have illustrated the usefulness of our framework in three case studies. In the case of OPTLS, we uncovered a subtle problem in the original reduction due to the simplicity of $\mathcal{F}_{\text{crypto}}$, which makes very explicit in which cases security can be guaranteed.

In future work, we will apply our framework to other real world protocols and extend the framework to further facilitate their cryptographic analysis.

IX. ACKNOWLEDGMENT

his project was in part funded by the *Deutsche Forschungsgemeinschaft* (DFG) through Grant KU 1434/9-1.

REFERENCES

- [1] C. Cremers, M. Horvat, S. Scott, and T. van der Merwe, "Automated Analysis and Verification of TLS 1.3: 0-RTT, Resumption and Delayed Authentication," in *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*. IEEE Computer Society, 2016, pp. 470–485.
- [2] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P. Strub, and J. K. Zinzindohoue, "A Messy State of the Union: Taming the Composite State Machines of TLS," in *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. IEEE Computer Society, 2015, pp. 535–552.
- [3] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P. Strub, and S. Z. Béguelin, "Proving the TLS Handshake Secure (As It Is)," in *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 8617. Springer, 2014, pp. 235–255.
- [4] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y. Strub, "Implementing TLS with Verified Cryptographic Security," in *IEEE Symposium on Security and Privacy (S&P 2013)*. IEEE Computer Society, 2013.
- [5] H. Krawczyk, K. G. Paterson, and H. Wee, "On the Security of the TLS Protocol: A Systematic Analysis," in *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Proceedings*, ser. Lecture Notes in Computer Science, vol. 8042. Springer, 2013, pp. 429–448.
- [6] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk, "On the Security of TLS-DHE in the Standard Model," in *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Proceedings*, ser. Lecture Notes in Computer Science, R. Safavi-Naini and R. Canetti, Eds., vol. 7417. Springer, 2012, pp. 273–293.
- [7] B. Dowling, M. Fischlin, F. Günther, and D. Stebila, "A Cryptographic Analysis of the TLS 1.3 Handshake Protocol Candidates," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*. ACM, 2015, pp. 1197–1210.
- [8] C. Badertscher, C. Matt, U. Maurer, P. Rogaway, and B. Tackmann, "Augmented Secure Channels and the Goal of the TLS 1.3 Record Layer," in *Provable Security - 9th International Conference, ProvSec 2015, Kanazawa, Japan, November 24-26, 2015, Proceedings*, ser. Lecture Notes in Computer Science, vol. 9451. Springer, 2015, pp. 85–104.
- [9] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi, "(De-)Constructing TLS 1.3," in *Progress in Cryptology - INDOCRYPT 2015 - 16th International Conference on Cryptology in India, Bangalore, India, December 6-9, 2015, Proceedings*, ser. Lecture Notes in Computer Science, vol. 9462. Springer, 2015, pp. 85–102.
- [10] S. C. Williams, "Analysis of the SSH Key Exchange Protocol," in *13th IMA International Conference of Cryptography and Coding (IMACC 2011)*, ser. Lecture Notes in Computer Science, L. Chen, Ed., vol. 7089. Springer, 2011, pp. 356–374.
- [11] R. Küsters and M. Tuengerthal, "Composition Theorems Without Pre-Established Session Identifiers," in *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS 2011)*, Y. Chen, G. Danezis, and V. Shmatikov, Eds. ACM, 2011, pp. 41–50.
- [12] K. Cohn-Gordon, C. Cremers, B. Dowling, L. Garratt, and D. Stebila, "A Formal Security Analysis of the Signal Messaging Protocol," *IACR Cryptology ePrint Archive*, vol. 2016, p. 1013, 2016.
- [13] C. Brzuska, M. Fischlin, N. P. Smart, B. Warinschi, and S. C. Williams, "Less is more: relaxed yet composable security notions for key exchange," *Int. J. Inf. Sec.*, vol. 12, no. 4, pp. 267–297, 2013.
- [14] C. Brzuska, M. Fischlin, B. Warinschi, and S. C. Williams, "Composability of Bellare-Rogaway Key Exchange Protocol," in *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS 2011)*. ACM, 2011, pp. 51–62.

- [15] R. Canetti, “Universally Composable Security: A New Paradigm for Cryptographic Protocols,” in *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*. IEEE Computer Society, 2001, pp. 136–145.
- [16] R. Küsters, “Simulation-Based Security with Inexhaustible Interactive Turing Machines,” in *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW-19 2006)*. IEEE Computer Society, 2006, pp. 309–320, see <http://eprint.iacr.org/2013/025/> for a full and revised version.
- [17] D. Hofheinz and V. Shoup, “GNUC: A New Universal Composability Framework,” *J. Cryptology*, vol. 28, no. 3, pp. 423–508, 2015.
- [18] U. Maurer, “Constructive Cryptography - A New Paradigm for Security Definitions and Proofs,” in *Theory of Security and Applications - Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31 - April 1, 2011, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 6993. Springer, 2011, pp. 33–56.
- [19] R. Küsters and M. Tuengerthal, “Ideal Key Derivation and Encryption in Simulation-based Security,” in *Topics in Cryptology - CT-RSA 2011, The Cryptographers’ Track at the RSA Conference 2011, Proceedings*, ser. Lecture Notes in Computer Science, A. Kiayias, Ed., vol. 6558. Springer, 2011, pp. 161–179.
- [20] J. Camenisch, R. R. Enderlein, S. Krenn, R. Küsters, and D. Rausch, “Universal Composition with Responsive Environments,” in *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security*, ser. Lecture Notes in Computer Science, J. H. Cheon and T. Takagi, Eds., vol. 10032. Springer, 2016, pp. 807–840.
- [21] “ISO/IEC IS 9798-3, Entity authentication mechanisms — Part 3: Entity authentication using asymmetric techniques,” 1993.
- [22] H. Krawczyk, “SIGMA: The ‘SIGn-and-MAC’ Approach to Authenticated Diffie-Hellman and Its Use in the IKE-Protocols,” in *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, D. Boneh, Ed., vol. 2729. Springer, 2003, pp. 400–425.
- [23] H. Krawczyk and H. Wee, “The OPTLS Protocol and TLS 1.3,” in *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, 2016, pp. 81–96.
- [24] R. Küsters and D. Rausch, “A Framework for Universally Composable Diffie-Hellman Key Exchange,” *Cryptology ePrint Archive*, Tech. Rep. 2017/256, 2017, available at <http://eprint.iacr.org/2017/256>.
- [25] R. Küsters and M. Tuengerthal, “The IITM Model: a Simple and Expressive Model for Universal Composability,” *Cryptology ePrint Archive*, Tech. Rep. 2013/025, 2013, available at <http://eprint.iacr.org/2013/025>.
- [26] U. M. Maurer and S. Wolf, “Diffie-Hellman Oracles,” in *Advances in Cryptology - CRYPTO ’96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, ser. Lecture Notes in Computer Science, vol. 1109. Springer, 1996, pp. 268–282.
- [27] M. Abdalla, M. Bellare, and P. Rogaway, “The Oracle Diffie-Hellman Assumptions and an Analysis of DHIES,” in *Topics in Cryptology - CT-RSA 2001, The Cryptographer’s Track at RSA Conference 2001, San Francisco, CA, USA, April 8-12, 2001, Proceedings*, ser. Lecture Notes in Computer Science, vol. 2020. Springer, 2001, pp. 143–158.
- [28] M. Backes, M. Dürmuth, D. Hofheinz, and R. Küsters, “Conditional Reactive Simulatability,” *International Journal of Information Security (IJIS)*, vol. 7, no. 2, pp. 155–169, April 2008.
- [29] R. Canetti and M. Fischlin, “Universally Composable Commitments,” in *Advances in Cryptology—CRYPTO 2001, 21st Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, vol. 2139. Springer, 2001, pp. 19–40.
- [30] R. Canetti and H. Krawczyk, “Universally Composable Notions of Key Exchange and Secure Channels,” in *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings*, ser. Lecture Notes in Computer Science, vol. 2332. Springer, 2002, pp. 337–351.
- [31] E. Rescorla, “The transport layer security (TLS) protocol version 1.3 (draft 09),” October 2015, <https://tools.ietf.org/html/draft-ietf-tls-tls13-09>.
- [32] H. Krawczyk, “Cryptographic Extraction and Key Derivation: The HKDF Scheme,” in *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Proceedings*, ser. Lecture Notes in Computer Science, T. Rabin, Ed., vol. 6223. Springer, 2010, pp. 631–648.
- [33] B. Blanchet, “An Efficient Cryptographic Protocol Verifier Based on Prolog Rules,” in *Proceedings of the 14th IEEE Computer Security Foundations Workshop (CSFW-14)*. IEEE Computer Society, 2001, pp. 82–96.
- [34] S. Meier, B. Schmidt, C. Cremers, and D. A. Basin, “The TAMARIN Prover for the Symbolic Analysis of Security Protocols,” in *Computer Aided Verification - 25th International Conference (CAV 2013)*, ser. Lecture Notes in Computer Science, N. Sharygina and H. Veith, Eds., vol. 8044. Springer, 2013, pp. 696–701.
- [35] C. J. F. Cremers, “The Scyther Tool: Verification, Falsification, and Analysis of Security Protocols,” in *Computer Aided Verification, 20th International Conference, CAV 2008, Princeton, NJ, USA, July 7-14, 2008, Proceedings*, ser. Lecture Notes in Computer Science, vol. 5123. Springer, 2008, pp. 414–418.
- [36] N. Swamy, C. Hritcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P. Strub, M. Kohlweiss, J. K. Zinzindohoue, and S. Z. Béguelin, “Dependent types and multi-monadic effects in F,” in *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium*

A. Proof Sketch of Theorem 2

As mentioned before, one of the key ideas for the definition of the simulator \mathcal{S} is to provide g^c for unknown Diffie-Hellman keys, where c is chosen uniformly at random from $\{1, \dots, n\}$, and g^{ab} for known ones. The proof itself consists of series of hybrid systems where we replace parts of the realization with the version used in the ideal protocol and then show that no environment can distinguish this replacement with more than a negligible probability.

In the first step, one defines a hybrid system $\mathcal{P}_{\text{crypto}}^1$ where all asymmetric operations and nonce generation is handled as in $\mathcal{F}_{\text{crypto}}$ while all other operations are performed as in $\mathcal{P}_{\text{crypto}}$. Because we did not modify any of these operations, the original proof still holds, which reduced this step to the security of the asymmetric operations.

Next, one defines a hybrid system $\mathcal{P}_{\text{crypto}}^2$ where also exponent handling is replaced with the ideal version. In particular, $\mathcal{P}_{\text{crypto}}^2$ prevents exponent guessing and collisions. Any distinguishing environment on this system can be reduced to the DDH assumption: If the environment manages to guess an exponent, or an unknown exponent is generated that is not fresh, then this can be used by an attacker on the DDH assumption to calculate the secret exponent a from the experiment. We note that this reduction requires a lot of attention to details and is more involved than usual reductions to the DDH assumption. This is because $\mathcal{P}_{\text{crypto}}^2$ can be used by the environment to perform several operations with a and g^a that are not available in the DDH experiment; an adversary must be able to simulate all of these operations without actually knowing a .

In the the third hybrid system $\mathcal{P}_{\text{crypto}}^3$ one replaces real with ideal Diffie-Hellman key generation, however, without preventing key collisions or key guessing. That is, the simulator provides the Diffie-Hellman keys as described above. This step requires a hybrid argument itself, as we have to replace a polynomial number of unknown keys in the order of their creation. We can then reduce the distinguishing advantage of an environment for the r -th and $r + 1$ -th hybrid system to the DDH assumption. Importantly, we have to establish a *single* negligible bound for the distinguishing advantage that is independent of r , as the sum of polynomially many different negligible functions is not necessarily negligible. Just as in the previous step, the reduction in this step requires a lot of care for details as there are several operations in the hybrid systems that an adversary on the DDH assumption has to simulate without knowing the secret exponents a and b of the DDH experiment.

In the fourth hybrid system $\mathcal{P}_{\text{crypto}}^4$, symmetric encryption and key derivation are replaced with their ideal versions, and key guessing and key collision are prevented. Again, this step requires a hybrid argument which is quite involved as we have to consider symmetric encryption and

- [37] N. Swamy, J. Chen, C. Fournet, P. Strub, K. Bhargavan, and J. Yang, “Secure distributed programming with value-dependent types,” *J. Funct. Program.*, vol. 23, no. 4, pp. 402–451, 2013.
- [38] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei, “Refinement Types for Secure Implementations,” in *Proceedings of the 21st IEEE Computer Security Foundations Symposium, CSF 2008*. IEEE Computer Society, 2008, pp. 17–32.
- [39] R. Küsters, T. Truderung, and J. Graf, “A Framework for the Cryptographic Verification of Java-like Programs,” in *25th IEEE Computer Security Foundations Symposium (CSF 2012)*. IEEE Computer Society, 2012, pp. 198–212.
- [40] M. Bellare and P. Rogaway, “Entity Authentication and Key Distribution,” in *Advances in Cryptology – Crypto ’93, 13th Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, D. Stinson, Ed., vol. 773. Springer-Verlag, 1993, pp. 232–249.
- [41] R. Canetti and H. Krawczyk, “Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels,” in *Advances in Cryptology – EUROCRYPT 2001*, ser. Lecture Notes in Computer Science, B. Pfitzmann, Ed., vol. 2045. Springer, 2001, pp. 453–474.
- [42] C. J. F. Cremers and M. Feltz, “Beyond eCK: Perfect Forward Secrecy under Actor Compromise and Ephemeral-Key Reveal,” in *Computer Security - ESORICS 2012 - 17th European Symposium on Research in Computer Security*, ser. Lecture Notes in Computer Science, S. Foresti, M. Yung, and F. Martinelli, Eds., vol. 7459. Springer, 2012, pp. 734–751.
- [43] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin, “Computer-Aided Security Proofs for the Working Cryptographer,” in *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, ser. Lecture Notes in Computer Science, vol. 6841. Springer, 2011, pp. 71–90.
- [44] B. Blanchet, “A Computationally Sound Mechanized Prover for Security Protocols,” in *IEEE Symposium on Security and Privacy (S&P 2006)*. IEEE Computer Society, 2006, pp. 140–154.
- [45] R. Canetti and S. Gajek, “Universally Composable Symbolic Analysis of Diffie-Hellman based Key Exchange,” *Cryptology ePrint Archive*, Tech. Rep. 2010/303, 2010, available at <http://eprint.iacr.org/2010/303>.
- [46] R. Canetti and H. Krawczyk, “Security Analysis of IKE’s Signature-Based Key-Exchange Protocol,” in *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, M. Yung, Ed., vol. 2442. Springer, 2002, pp. 143–161.

key derivation simultaneously: All symmetric keys can be encrypted, thus the security of symmetric keys depends on the security of the encryption scheme. However, Diffie-Hellman keys and key derivation keys can be used to create new symmetric keys, i.e., the security of the encryption scheme in turn depends on the security of the key derivation schemes. In the hybrid argument, we track the order in which unknown keys are used for the first time. The r -th hybrid system performs operations with the first r unknown keys ideally, and all other operations as in the realization. One can then reduce the distinguishing advantage of an environment for the r -th and $r + 1$ -th hybrid system to the security games of the encryption and key derivation schemes. Again, it is important to establish a negligible bound for the distinguishing advantage that is independent of r .

In the final step, we have to replace MACs with their ideal versions. As this step is unaffected by our extension, just as the first step, the original proof still holds, which reduced this step to the security of the MAC scheme. \square

B. Postponed cases of the proof of Theorem 3

We still have to show that the simulation is perfect in case of an uncorrupted instance of a responder during the key establishment phase, in case of uncorrupted instances during the key usage phase, and in case of corrupted instances.

Let $(pid_R, lsid_R, R)$ be an uncorrupted instance of M_R that wants to establish a session with pid' . We only have to show that $(pid_R, lsid_R, R)$ is already part of a global session in $\mathcal{F}_{\text{key-use}}^{\text{MA}}$ when it outputs a pointer to the session key, as every action up to that point can be simulated perfectly. Observe that, if $(pid_R, lsid_R, R)$ outputs such a pointer, then it has accepted the third protocol message and pid' must still be uncorrupted. In other words, there is an instance of pid' , say $(pid', lsid', r')$, that has signed the message $m = (g^y, g^x, pid_R)$, where y is the secret exponent of $(pid_R, lsid_R, R)$ and x is the secret exponent of $(pid', lsid', r')$. This instance is uncorrupted by the same argument as above. We now argue that this instance is an initiator, i.e., $r' = I$: Suppose by contradiction that $r' = R$, i.e., the message was signed by a responder whose secret exponent is x and who has received the group element g^y . Recall that, when such an instance receives g^y , it first uses the `BlockGroupElement` command on g^y . Thus, afterwards no instance will be able to generate g^y via a `GenExp` command. Hence, the instance $(pid', lsid', r')$ cannot have received its first protocol message *before* the instance $(pid_R, lsid_R, R)$ has received its first protocol message, as in this case $(pid_R, lsid_R, R)$ would no longer be able to create the exponent y . By the same argument, $(pid', lsid', r')$ also cannot have received its first protocol message *after* $(pid_R, lsid_R, R)$ has received its first protocol message, as in this case $(pid', lsid', r')$ would not be able to create the secret exponent x . Of course, we also have that

$(pid_R, lsid_R, R) \neq (pid', lsid', r')$ as $x \neq y$ (g^x is blocked when g^y is generated). Thus, we conclude that $r' = I$. We still have to argue that $(pid_R, lsid_R, R)$ is already in a global session with $(pid', lsid', r')$: As $(pid', lsid', r')$ has signed a message, it has already completed its part of the key exchange and thus is in a session with some responder. By the definition of \mathcal{S} , this will be the honest instance of a responder that signed the message $m' = (g^x, g^y, pid')$. However, the instance $(pid_R, lsid_R, R)$ is the only one that would sign such a message as y/g^y is unique, so $(pid', lsid', r')$ is in a session with $(pid_R, lsid_R, R)$. Note that both instances use the same unknown exponents x and y to derive a session key, with $x \neq y$, and they are never paired with any other DH shares. Thus both instances will output pointers to the same unknown session key.

Now consider an honest instance in the key usage phase. As shown above, such an instance in the real world/internal simulation will have a pointer to an unknown session key in $\mathcal{F}_{\text{crypto}}$. Furthermore, no instance besides the two instances in the same session will have access to this pointer as no other instances have a pointer to x or y . Thus, instances in the real world behave just like instances in the ideal world that use $\mathcal{F}'_{\text{crypto}}$, i.e., the simulation is perfect also in this case.

Finally consider a corrupted instance. The simulator has full control over the I/O interface of such an instance. If the instance was explicitly corrupted by the adversary (i.e., it is under the control of the adversary) either before or after the key exchange, the adversary gets access only to known keys which do not exist in $\mathcal{F}'_{\text{crypto}}$. Thus, the simulator is able to simulate the exact behavior of $\mathcal{F}_{\text{crypto}}$ for such explicitly corrupted instances. In the case of a corrupted instance that was not explicitly corrupted by the adversary (i.e., where one of the signing keys is corrupted), the simulator also has to simulate unknown keys. However, these unknown keys will never be inserted into/used in $\mathcal{F}'_{\text{crypto}}$ as no honest instance will complete a KE with a corrupted instance (as shown above). Thus, the simulator can also easily simulate this case.

We note that \mathcal{S} is a responsive simulator as it fulfills the runtime conditions and it responds immediately to restricting messages as long as the environment does the same, which happens with overwhelming probability. This concludes the proof. \square