

Authentication in Key-Exchange: Definitions, Relations and Composition

Cyprien Delpuch de Saint Guilhem^{1,2}, Marc Fischlin³, and Bogdan Warinschi²

¹ imec-COSIC, KU Leuven, Belgium

² Dept Computer Science, University of Bristol, United Kingdom

³ Computer Science, Technische Universität Darmstadt, Germany

cyprien.delpuchdesaintguilhem@kuleuven.be, marc.fischlin@cryptoplexity.de, bogdan@cs.bris.ac.uk

Abstract. We present a systematic approach to define and study authentication notions in authenticated key-exchange protocols. We propose and use a flexible and expressive predicate-based definitional framework. Our definitions capture key and entity *authentication*, in both implicit and explicit variants, as well as key and entity *confirmation*, for authenticated key-exchange protocols. In particular, we capture critical notions in the authentication space such as *key-compromise impersonation* resistance and security against *unknown key-share* attacks. We first present and explore these definitions within the Bellare–Rogaway model and then extend them to Canetti–Krawczyk-style models.

We then show two useful applications of our framework. First, we look at the authentication guarantees of three representative protocols to draw several useful lessons for protocol design. The core technical contribution of this paper is then to formally establish that composition of secure implicitly authenticated key-exchange with subsequent confirmation protocols yields explicit authentication guarantees. Without a formal separation of implicit and explicit authentication from secrecy, a proof of this folklore result could not have been established.

1 Introduction

The commonly expected level of security for authenticated key-exchange (AKE) protocols comprises two aspects. *Authentication* provides guarantees on the identities of the parties involved in the protocol execution. *Secrecy* promises that the key is not known by active adversaries. The two together ensure that the key is only known by the “right” parties.

The apparent simplicity of these informal definitions hides a complex landscape as both secrecy and authentication come in many related but distinct flavours. Authentication can be one-way or mutual; it can refer directly to the entities involved in the protocol run (entity authentication) or indirectly to the identities of the parties that hold the keys (key authentication); it can be explicit (i.e. hold as soon as the protocol has finished) or implicit (i.e. refer to the moment when the key is actually used). In addition, AKE protocols are often expected to guarantee key-confirmation, where a party which derives a key is convinced that it has also been derived by another session, to ensure that only properly shared keys are used in communication. Each of the above properties may also come with different levels of strength due to the asymmetry in most two-party protocols.⁴

Starting with the seminal work of Bellare and Rogaway [BR94], research has made steady progress in formalizing and clarifying what some of these security notions signify. For example, the early definition proposed in [BR94] requires explicit entity authentication for all AKE protocols: when a session finishes its execution, there exists a unique session of the intended peer to which it is partnered.⁵ In particular, this definition is applicable to arbitrary two-party protocols (i.e. it does not require that parties derive keys). Subsequent work resulted in numerous extensions and variations including changes to the

⁴ One party always terminates first after sending the last message and is not able to confirm the message’s correct delivery.

⁵ Partnering was originally defined through matching conversations.

core partnering mechanism [BPR00, BR95, CK01, BFWW11, LS17] and extensions of the adversarial powers [CK01, LLM07, CF12]. Over time, definitions have shifted to guarantees associated only with keys, separated from identities [Kra05, BSWW13], and formalized other aspects not considered by original definitions e.g. various forms of forward secrecy [Kra05, CF12] and key-confirmation properties [FGSW16].

All these developments were guided chiefly by protocol design ideas and intuitive understanding of the security guarantees and therefore happened outside of a complete framework of security notions. Despite some works which attempt to systematize and understand the relative merits of the different models [CBH05a, Ust09, Cre11], emphasis was on protocol design and more pressing aspects (e.g. privacy) rather than on a thorough evaluation of the different guarantees (including authentication). This lack of comprehensive study has led to missing definitions, unclear relations between properties, and the use of folklore results which lack formal support. Perhaps surprisingly, core security notions such as *implicit authentication* have influenced the design of large classes of protocols while not being formally defined. Similarly, large classes of attacks on authentication, such as *unknown key-share* or *key-compromise impersonation*, have led many individual protocols to attempt to avoid them, but no single security definition has been argued to properly defend against such attacks.

The case of implicit authentication. The discussion above is best illustrated by the literature on implicit authentication, a concept which has guided protocol design for decades [MTI86, Kra05, CF12]. This line of research weakens the level of authentication provided. It requires that when a session derives a key, the adversary cannot force an unintended party to derive the same key.

Unfortunately, implicit authentication is tangled together with authentication and secrecy guarantee into single monolithic definitions which leads to several undesirable consequences. First, it makes the definitions themselves difficult to understand. The relation between the trust models for authentication and secrecy is not clear cut. For example, secrecy for some session when the intended partner is adversarially controlled does not make sense, whereas integrity guarantees are still desirable. This makes it non-obvious that the two notions can actually be compounded. Arguing that the definition captures the “right” guarantees requires a rather cumbersome reduction from an authentication attack to an attack against secrecy. It also makes the definition less portable; each change in the underlying execution model, e.g. varying the corruption model, requires dealing, unnecessarily, with the idiosyncrasies associated to secrecy (e.g. specific freshness notions)

Finally, it makes reasoning about properties related strictly to authentication cumbersome. A simple example is that it is difficult to justify that implicit authentication is a consequence of explicit authentication. More importantly, the lack of standalone definitions for authentication, separated from that of secrecy, makes it impossible to justify the folklore result that use of an implicitly authenticated key yields explicit authentication guarantees. This is a highly desirable property which allows for more efficient key-exchange protocol which delay the authentication guarantees to when keys are used rather than when they are agreed. This folklore composition, sometime used as alternative definition for implicit authentication, has no rigorous justification: there is no formal proof which establishes if, and under what conditions, the *use* of implicitly authenticated keys provides further authentication guarantees to parties.⁶

Our results. To address these gaps, we present a comprehensive study of the various forms of authentication in AKE protocols. We discuss how different combinations of properties relate to each other through implications and/or equivalence relations. We detail our results below.

Definitions. We identify and formally define a range of authentication properties. Figure 1 partially summarizes our definitional contributions; it shows that we cover implicit and explicit authentication, for both for keys and for entities, together with key and entity confirmation. Entity authentication and confirmation have not been defined separately from the other notions before in the literature and we do so here for completeness and symmetry, and to guide protocol design.

⁶ Yang [Yan13] provides compiler to add entity authentication to secure key exchange protocols via MACs and PRFs and a key refresh step, but this does not cover the case where the actual session key is used, e.g., in the subsequent

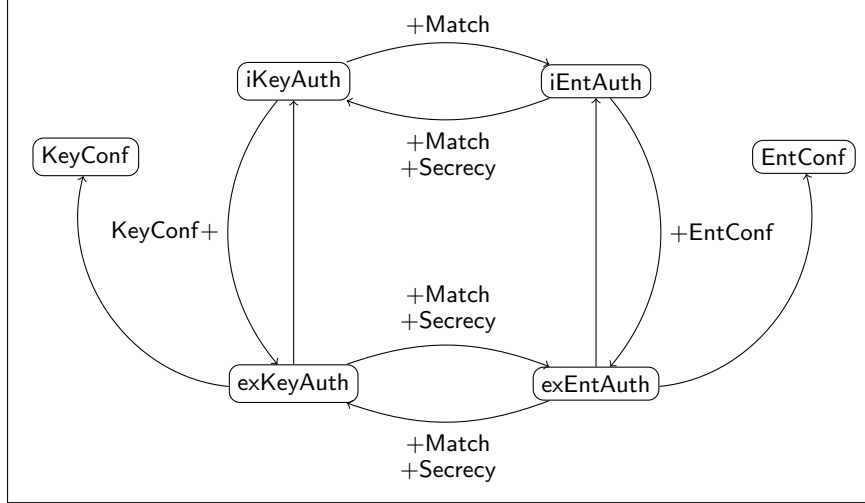


Fig. 1. Relationship of notions for implicit key and entity authentication ($iKeyAuth$ and $iEntAuth$), explicit key and entity authentication ($exKeyAuth$ and $exEntAuth$), key and entity confirmation ($KeyConf$ and $EntConf$), $Match$ security and key $Secrecy$. We note that key confirmation and explicit authentication definitions actually come in two flavours, full and almost-full, depending on which party receives the final protocol message. The implications hold accordingly for these sub versions.

Our definitions use a formalism inspired by the work of Rogaway and Stegers [RS09]. Specifically, we use logical statements to express when a desirable property is satisfied by the overall state of the protocol execution; security then demands that this property is satisfied with overwhelming probability. A key benefit of the approach is that these precise guarantees are unambiguous, unlike the English prose often used to describe authentication properties in AKE [MvOV97], and can be assessed by parsing the logical formalism.

Not evident from Figure 1 is that we cover both one-way and mutual authentication settings using a new and flexible mechanism to unify both versions. Another aspect not illustrated in Figure 1 is that some of the notions come in two flavours: “full” which is the stronger guarantee, offered to the party which receives the last message, and “almost-full” which is offered to the party that sends the last message (and therefore does not know if this message is delivered).

In our work we made several important choices about the definitional framework. The first concerns the adversarial powers and goals. While the rigorous formalization of adversarial goals is one of our main contributions, we are aware that these need to be placed in (and often depend on) the context of the available adversarial powers. A variety of choices exists, especially when deciding which parties the adversary is allowed to corrupt, and we opt for a more conservative scenario. We extend the Bellare–Rogaway (BR) model [BR94] by selectively allowing the adversary to compromise *any* of the two parties that engage in an authenticated session. One implication of this choice is that our definitions fully cover advanced properties like *key-compromise impersonation (KCI)* resistance [BJM97] and security against *unknown key-share (UKS)* attacks [BWM99]. We also fully allow the reveal of session keys as we consider that authentication should already hold upon *derivation* of the key. In a later section, we extend our definitions to the AKE models in the Canetti–Krawczyk style [CK01] which also include the reveal of ephemeral material. We show that our definitions naturally extend to this stronger adversarial model and that we can in fact separate out authentication from secrecy definitions in order to capture our stronger form of implicit authentication.

Another choice is how to capture matching sessions (sometimes called session partnering). Such sessions are deemed by the model to have communicated with one another and, among other implications, they are

channel protocol. Such a composed protocol clearly does not guarantee key secrecy anymore and one thus needs to argue along the authentication property alone to justify this property.

expected to derive the same key. The notion of partnering is crucial for defining secrecy of keys (technically to restrict the powers of the adversary in a meaningful way) but also impacts authentication. When authentication is required, then sessions deemed as matching should agree on identities. In the literature, there are two prevalent mechanisms that capture this idea: matching conversations and session identifiers. In this paper we use session identifiers – some related to keys, some related to identities – and demand that they meet a refined definition of Match security which was first proposed in earlier work [BFWW11] as “partnering security.”

Such careful reflections about security notions and the logical formalism enable the intuitive understanding to be matched almost one-for-one in the security definitions, therefore guaranteeing that each stand-alone definition captures exactly what it is supposed to, something that could not be achieved with a single, secrecy-focused definition.

Relations. A second benefit of our predicate-based definitions is that the use of logic simplifies and facilitates the study of relations between the different properties. It is immediate, for example, that implicit authentication is a consequence of explicit authentication. It is also easy to see that implicit key authentication together with key confirmation yield explicit key authentication.

The logical formulations also make it clear that although key authentication and entity authentication are obviously related, the relation is more subtle than appears at first. A clear separation is that the latter property is applicable to a larger class of protocols. Perhaps less clear is that for AKE protocols, entity authentication does not imply key authentication. Indeed, the authentication guarantee could come from protocol flows that are not involved in key generation; moreover, the key could be non-secret so authentication via keys would then be meaningless. However, we complete the study of relations between the different notions by showing formally that an equivalence holds for all protocols which satisfy Match security and for which the keys are private.

Protocols. To illustrate the different levels of key authentication we briefly discuss how three well-known protocols fare with respect to the notions we put forth. Among other properties, we show that plain Diffie–Hellman formally provides authentication only when no party is required to authenticate themselves. Of course, this case does not imply any meaningful guarantees.

More interestingly, we test our formulation of implicit authentication in the case of the HMQV protocol [Kra05]. Due to the prudent choices in our definitions, we find that we cannot use the original proof of secrecy in a black-box fashion since it holds for different adversarial powers. This indicates that existing security definitions for IAKE protocols do not straightforwardly provide the strongest implicit guarantees. Finally, we show that the much heavier TLS 1.3 protocol satisfies the strongest guarantee: explicit key authentication. Following our analysis, we derive some rules of thumb which can inform the design of protocols for AKE protocols. These analyses with our stand-alone definitions also allow us to hone in on the assumptions that are necessary for authentication which turn out to be significantly weaker than those required for secrecy.

Composition result. As explained above, a protocol with implicit key authentication and key confirmation also satisfies explicit key authentication. This result can be used to prove the folklore (composition) result that *the use* of the key obtained from an implicitly authenticated protocol yields stronger authentication guarantees. Of course, the use of the key has to be meaningful, in the sense that only a party who possesses the key can successfully engage in the task. One can then regard the use of the key as a means to “strengthen” the key-exchange protocol to also provide key confirmation.

We formalize this idea as follows. We define the class of *key-confirming protocols* which, in addition to their basic functionality, provide a key confirmation guarantee upon a “successful” use of a key. Such protocols can be as simple as sending a MAC on some fixed message⁷, or more complex. For example, we argue that an authenticated channel also offers key-confirming guarantees. We then prove that running an implicitly authenticated key-exchange protocol followed by a key-confirming protocol results in an explicitly

⁷ To prevent replay attacks, messages in each direction need to be different.

authenticated key-exchange protocol. The desired result follows by observing that the composition is still an implicitly authenticated protocol in addition to providing key confirmation.

2 Game-based security for AKE protocols

We use the framework of Brzuska et al. [BFWW11] for cryptographic games and describe its formulation of the BR model [BR94] for AKE protocols. We augment it to provide a flexible authentication framework and to capture key-confirmation as formalised by Fischlin et al. [FGSW16].

In the security game, the adversary interacts with parties executing the protocol via queries; these capture its capabilities in a real-world execution. The adversary’s aim is to trigger an event defined as “bad” by the game whilst abiding by the game’s limitations on queries. We use λ to denote the security parameter, let 1^λ be its unary representation and let $\text{negl}(\lambda)$ denote an arbitrary negligible function. We denote by $\{0, 1\}^*$ the set of all finite-length bit-strings.

Identities. We let n_i be the number of parties modelled by the game. Each party has a unique identity i and we denote by $\mathcal{I} = \{i\}$ the identity set of size n_i .

To separate identities who are expected to authenticate we specify a subset $\mathcal{S} \subseteq \mathcal{I}$; this is our first augmentation. This models real-world servers who have to convince clients of their identity by means of a certificate. These clients, modelled by the identities in $\mathcal{I} \setminus \mathcal{S}$, do not have such secret information. This modelling flexibly captures varying forms of authentication; a secure protocol for which $\mathcal{S} = \mathcal{I}$ provides mutual authentication for all sessions, whereas one for which $\mathcal{S} = \emptyset$ provides no authentication whatsoever; for \mathcal{S} a non-empty proper subset of \mathcal{I} , it provides one-way authentication from identities in \mathcal{S} . We leave the specification of \mathcal{S} within \mathcal{I} as a design choice for protocols.

We assume that each party is aware of its own identity in an execution. We sometimes use id_A and id_B to distinguish between the identities of two parties A and B in an execution. We work in the *pre-specified peer model* [MU08] where each session knows its intended partner’s identity from the start. We can restrict to the case that only identities of authenticating partners are known, e.g. if a client remains anonymous in a TLS connection.

Protocols and sessions. A protocol is a pair of algorithms $\pi = (\text{kg}, \zeta)$ where $\text{kg}(1^\lambda)$ is a randomized key generation algorithm and where ζ is the algorithm executed locally by parties engaging in the protocol. Local sessions are identified by a *local session identifier* $\ell \in \mathcal{I} \times \mathcal{I} \times \mathbb{Z}$ where $\ell = (i, j, k)$ refers to the k -th session of identity i with intended peer j . We use $\ell.\text{id}$ and $\ell.\text{pid}$ to refer to i and j respectively. We let n_s be the maximum value of k for any ℓ in a game.

We use the notion of session identifiers as was originally proposed by [BPR00] to “partner” two sessions as having engaged in the same execution. These are computed by the protocol itself, different from the *local session identifier* ℓ which is an artefact of the model. For simplicity we maintain the original approach of considering these session identifiers to be revealed upon acceptance.

2.1 Common game states

Security games maintain states to keep track of the execution of the protocol. This consists of five elements: a list **LSID** of valid local session identifiers, a list **SST** of protocol-related *session states*, a list **LST** of game-related *local session states*, a *game execution state* **EST** which contains global protocol-related information, and a *model state* **MST** which contains global game-related information, relevant to the security notion (e.g. a hidden bit). Our authentication games share the same execution, game and local states. Here we describe and augment the states from [BFWW11].

Game execution state. As in [BFWW11], the execution state **EST** contains a list $\mathcal{L}_{\text{keys}} = \{(i, \text{pk}_i, \text{sk}_i, \delta_i)\}_{i \in \mathcal{I}}$ where $\delta_i \in \{\text{honest}, \text{corrupt}\}$ denotes whether sk_i has been corrupted.

Session state. The state of the local session with identifier $\ell = (i, j, k)$ is composed of the following:

- (pk_i, sk_i) , the long-term key pair of identity i , the “owner” of the session. This is initialised to $\ell.id$ ’s key pair if $i \in \mathcal{S}$ or set to \perp if $i \in \mathcal{I} \setminus \mathcal{S}$. One may think of the public keys as certified, in which case the party also holds a certificate $cert_i$ for pk_i under its identity, but we omit details here.
- pk_j , the long-term public key of identity j , the intended peer of the session. This is initialised to $\ell.pid$ ’s public key if $j \in \mathcal{S}$ or set to \perp if $j \in \mathcal{I} \setminus \mathcal{S}$.
- $crypt \in \{0, 1\}^*$ is some protocol-specific private session state used to maintain secret values from one invocation to the next.
- $accept \in \{\text{true}, \text{false}, \perp\}$ indicates whether the party has *accepted* or *rejected* the session as an succesful execution of π . Initially set to \perp , to signify *running*, $accept$ may change to either *true* or *false* only upon termination. We assume the value of $accept$ is public.
- $sid \in \{0, 1\}^* \cup \{\perp\}$, the session identifier as specified by the protocol. Initially set to \perp , it may be changed once to a non-trivial value. If the sid is different from \perp , then $accept$ must be set to *true*, and vice versa, if $accept$ is set to *true*, then sid must become different from \perp . We assume that the value of sid is made public when $accept$ is set to *true*.
- $key \in \{0, 1\}^* \cup \{\perp\}$ is the (session-)key locally derived during the execution. Initially set to \perp , it may be changed once to a non-trivial value. If the key is different from \perp , then $accept$ must be set to *true*, and vice versa, if $accept$ is set to *true* then, key must become different from \perp . We note that this implies that sessions must terminate with the same call to the protocol as that which sets the key and the sid , they cannot continue once the key is set.
- $kconf \in \{\text{full}, \text{almost}, \text{no}, \perp\}$ indicates the form of key confirmation that the owner expects to receive. This addition to the model captures the fact that one partner of a run always terminates first and therefore may not expect a full confirmation of the final session-key. The value of $kconf$ is initialised to \perp and set when the session is first activated.
- $kcid \in \{0, 1\}^* \cup \{\perp\}$ is a key-confirmation identifier, indicating sessions which will eventually derive the same key. Initially set to \perp , it may be changed once to a non-trivial value and may not be changed again. If key is different from \perp , then $kcid$ must be different from \perp .

We write $SST[\ell] = ((pk_i, sk_i), pk_j, crypt, accept, sid, key, kconf, kcid)$ to denote the session state of ℓ . We use the notation $\ell.sid$ or $\ell.key$ to refer to individual elements and use similar notation for the game, local session or model states.

This session state augments that of [BFWW11] with $kconf$ and $kcid$ from [FGSW16] along with some renaming. These are used to modularly capture the formal definition of key confirmation of [FGSW16], similarly to our addition of \mathcal{S} to capture different authentication directions. We refer to [BFWW11, Section 3] for a discussion on public session identifiers.

Local session state. The local session state consists of:

- $\delta_{ownr} \in \{\text{honest}, \text{corrupt}\}$: denotes whether the owner of the session was corrupted before the session was completed (i.e. while $\ell.accept = \perp$).
- $\delta_{peer} \in \{\text{honest}, \text{corrupt}\}$: denotes whether the intended peer of the session was corrupted before the session was completed.
- $\delta_{sess} \in \{\text{fresh}, \text{revealed}\}$: denotes whether the session-key for this session has been revealed to the adversary.

As in [BFWW11], keeping track of δ_{ownr} separately from δ_i allows sessions that terminated before their owner was corrupted to remain honest; this enables the modelling of *forward secrecy*. We write $LST[\ell] = (\delta_{ownr}, \delta_{peer}, \delta_{sess})$ for the local session state of session ℓ and use the notation $\ell.\delta_{sess}$ to refer to individual elements.

Setup. Modelling protocols using these states requires the following procedures:

- $(SST, EST) \leftarrow \text{setupE}(LSID, kg, 1^\lambda)$: for protocol-relevant components.
- $(LST, MST) \leftarrow \text{setupG}(LSID, SST, EST, 1^\lambda)$: for game-relevant components.

Our setupE is similar to that of [BFWW11] but it only generates long-term keys for the identities in \mathcal{S} and initialises the new elements of the session state.

2.2 Session partnering.

We define the partnering predicate using session identifiers.

Definition 2.1 (Partners). *We say that two sessions ℓ and ℓ' are partners if the predicate $\text{Partner}(\ell, \ell')$ holds true, where*

$$\text{Partner}(\ell, \ell') \iff [(\ell \neq \ell') \wedge (\ell.\text{sid} = \ell'.\text{sid} \neq \perp)].$$

Thus, to be partners, two sessions need to be administratively different and to both have set a non-trivial sid. This does not exclude the possibility that they belong to the same identity, i.e. that $\ell.\text{id} = \ell'.\text{id}$.

For correctness, we require that two sessions executing π without adversarial interaction derive identical sids upon accepting and are therefore partnered. We also require that two such sessions derive identical keys and kcids.

While this definition of partnering appears similar to that of “matching” sessions in CK-like models [CK01,Kra05,LLM07], it differs in two important aspects. The first is that it does not involve the sessions’ identities, thus separating out authentication notions. The second is that our sids are derived by the protocol itself rather than arbitrarily set by a higher layer. We note that our notion of sids supersedes that of matching conversations used for partnering in [Kra05,LLM07].

2.3 Adversarial interaction and common queries

The adversary \mathcal{A} is a probabilistic polynomial-time (PPT) algorithm that interacts with a game through queries specified by a set Q . Upon receiving a query $q \in Q$, the game has a behaviour algorithm χ which takes q together with the state to return a response to \mathcal{A} .

Not every query is always valid; this is captured by the **Valid** predicate which the game evaluates each time a query q is received. Based on q and on the current state, **Valid** returns either **true** or **false** which determines if χ is executed on q .

In addition to the common states, our security games for authentication notions also share a query set Q . We specify here the **Send**, **Reveal** and **Corrupt** queries following the work of Brzuska et al. but, as we model KCI resistance, we modify the **Valid** predicate.

The Send query. Whatever the game, Q always includes the **Send** query. It takes an identifier $\ell \in \text{LSID}$ and a message $m \in \{0, 1\}^*$ as inputs and is processed by χ by running π on $\text{SST}[\ell]$ with input m . This updates SST and returns a response m' which is given to \mathcal{A} together with **accept** and also **sid** if m triggered the termination of the session. This gives control of the network to \mathcal{A} and allows it to forward, alter, delay, create or delete messages.

The Reveal query. When \mathcal{A} submits $\text{Reveal}(\ell)$, this sets $\ell.\delta_{\text{key}} \leftarrow \text{revealed}$ and returns $\ell.\text{key}$.

The Corrupt query. We formalise the $\text{Corrupt}(i)$ query as follows. First, the value of δ_i in $\mathcal{L}_{\text{keys}}$ is set to **corrupt**. Then, for any session of the format $(i, *, *)$ for which **accept** = \perp , we set $\delta_{\text{ownr}} \leftarrow \text{corrupt}$; for any session of the format $(*, i, *)$ which is still running, we set $\delta_{\text{peer}} = \text{corrupt}$. Finally, sk_i is returned to \mathcal{A} .

The Valid predicate. A significant difference to [BFWW11] is that our **Valid** predicate allows for the adversary to submit a **Send** query to a session whose owner has already been corrupted. This is crucial to model KCI resistance as this notion guarantees a security property to sessions whose owner was corrupted *before* they terminated. Furthermore, the **Valid** predicate returns **false** if a **Reveal** query is made to a session whose **key** = \perp .

2.4 Winning condition and formal game definition

A game considers that \mathcal{A} has won a security game, i.e. broken a security property of π , if it succeeds in triggering a “bad” event. This event is defined by a predicate P which is a logical statement evaluated on the state. We denote this by $b \leftarrow P(\text{LSID}, \text{SST}, \text{LST}, \text{EST}, \text{MST})$, where $b \in \{0, 1\}$, and $b = 1$ signifies that \mathcal{A} has successfully triggered the “bad” event. We therefore define a generic security experiment as follows.

Definition 2.2. *A game G maintains a state $(\text{LSID}, \text{SST}, \text{LST}, \text{EST}, \text{MST})$ and is defined by the tuple $(\text{setupE}, \text{setupG}, Q, \text{Valid})$. An experiment is parameterised by a protocol π , an adversary \mathcal{A} and a game G ; it is executed as follows.*

1. *The experiment runs $(\text{SST}, \text{EST}) \leftarrow \text{setupE}(\text{LSID}, \text{kg}, 1^\lambda)$ and $(\text{LST}, \text{MST}) \leftarrow \text{setupG}(\text{LSID}, \text{SST}, \text{EST}, 1^\lambda)$.*
2. *The adversary submits queries from Q to the game which processes them with **Valid** and χ .*
3. *When \mathcal{A} terminates, $b \leftarrow P(\text{LSID}, \text{SST}, \text{LST}, \text{EST}, \text{MST})$ is evaluated by the experiment which finally outputs b .*

We note that our definition of a game G includes more than [BFWW11, Definition 1], namely Q and P , as these also uniquely characterise it. We denote the experiment by $\text{Exp}_{\mathcal{A}, \pi}^G(1^\lambda)$ and we write $\text{Exp}_{\mathcal{A}, \pi}^G(1^\lambda) = b$.

2.5 Match security

Before authentication or secrecy, a “good” AKE protocol should first provide certain correctness and soundness guarantees. A Match-secure AKE protocol should ensure that:

1. Partner sessions derive the same key and kcid (properties 1 and 2 below);
2. at most two sessions derive the same sid (property 3);
3. sessions with the same kcid accept with the same key (property 4).

This guarantees that \mathcal{A} is unable to create disagreements between partnered sessions. Formally, we define the following predicate.

Definition 2.3 (Match predicate). *The Match predicate evaluates to 1 iff*

$$\begin{aligned} \forall \ell, \ell', \ell'' \in \text{LSID}, \\ & (\text{Partner}(\ell, \ell') \wedge \ell.\text{key} \neq \perp \neq \ell'.\text{key}) \implies \text{Samekey}(\ell, \ell') & (1) \\ & \wedge (\text{Partner}(\ell, \ell') \wedge \ell.\text{kcid} \neq \perp \neq \ell'.\text{kcid}) \implies \text{Samekcid}(\ell, \ell') & (2) \\ & \wedge (\text{Partner}(\ell, \ell') \wedge \text{Partner}(\ell, \ell'')) \implies \ell' = \ell'' & (3) \\ & \wedge (\text{Samekcid}(\ell, \ell') \wedge \ell.\text{key} \neq \perp \neq \ell'.\text{key}) \implies \text{Samekey}(\ell, \ell'). & (4) \end{aligned}$$

where $\text{Samekey}(\ell, \ell') \iff [\ell' \neq \ell \wedge \ell'.\text{key} = \ell.\text{key} \neq \perp]$ and $\text{Samekcid}(\ell, \ell')$ is defined analogously.

We then define the Match security game G_{Match} in the sense of Definition 2.2 where π is an AKE protocol, the state, and the setupE algorithm are as in Section 2.1, the query set $Q = \{\text{Send}, \text{Reveal}, \text{Corrupt}\}$ and the behaviour χ are as in Section 2.3, the setupG algorithm sets the LST of each session to (honest, honest, fresh) and the predicate $P = \text{Match}$. The advantage of an adversary \mathcal{A} against the game G_{Match} with identity sets \mathcal{I}, \mathcal{S} is written as

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{Match}}}(1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{Match}}}(1^\lambda) = 0 \right].$$

Definition 2.4 (Match security). *An AKE protocol π is Match-secure for identity sets \mathcal{I}, \mathcal{S} if, for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{Match}}}(1^\lambda) = \text{negl}(\lambda)$.*

Comparison to previous definitions. Considered in the context of the BR model, our definition of *Match* security refines that of [BFWW11] in two ways. We first incorporate the conditions of the *KCIDbind* predicate of [FGSW16] as conditions (2) and (4). As this work builds a unified model of AKE protocols with both authentication and key confirmation, it is reasonable to add this predicate to the definition of *Match* since it concerns notions of correctness and soundness, as the previous definition already did for keys.

Secondly, we remove the requirement that partnered sessions should agree on each other’s identities. This condition implied that *Match*-secure AKE protocols already provided some form of authentication, albeit very weak. This mixed an authentication with design and soundness and we therefore remove it. We present separate definitions for authentication in the next section.

In the context of the CK-style models, as summarised in [Cre11], the usual first requirement of security is for matching sessions of uncorrupted identities to derive equal keys. We note that our notion of *Match* security would capture and extend this requirement if a *StateReveal* query were added to its game. We note that we do not restrict to uncorrupted identities, as we do not consider that the adversary takes control of sessions ℓ managed by the game; thus our *Match* predicate is only evaluated for honestly-behaving sessions.

3 Key authentication and confirmation

We present here our new predicate-based definitions for key authentication notions. We define three distinct flavours of authentication: *implicit*, *confirmation* and *explicit*. For each, we first discuss the intuitive understanding that motivates our definition and then give a formal statement. We then show that our definitions are consistent and that a protocol that combines implicit key authentication and key confirmation also provides explicit key authentication.

3.1 Implicit authentication

We take *implicit* to mean: “should there be a session ℓ' that possesses the same key as session ℓ , then the owner of session ℓ' *must* be the identity designated as the peer of session ℓ .” This does not guarantee the secrecy of the key, nor whether such a session ℓ' exists. Equivalently, this means that any session whose owner is *not* designated by the peer of ℓ is incapable of deriving the same session-key. (Recall that the term “session” refers to sessions executed by the model and that this does not forbid the adversary from deriving the key itself.)

This informal notion of implicit authentication raises the question whether to only consider sessions which interact with an honest peer, or to also allow those with a corrupted peer. The impact of this distinction was first observed by Diffie et al. [DvOW92]. In their design of the station-to-station (STS) protocol they aimed to prevent an attack in which one can make an honest B believe it is sharing a key with a malicious E , whereas the actual other honest key holder A intends to communicate with B . This has later become known as an *unknown key-share* (UKS) attack [BWM99] which, ironically, was shown to apply to the STS protocol in the same work [BWM99].

For our formal definition, the question is then either to restrict the adversary’s valid targets to the sessions that were executed with an honest peer, or to allow all sessions, even those that accepted with a corrupted peer, as valid targets. The first choice would comply with the idea stated in [BWM99]: “By definition, the provision of implicit key authentication is only considered in the case where B engages in the protocol with an honest entity (which E isn’t).” The second choice, would instead lead to a definition where UKS attack scenarios even with dishonest peers are accounted for. Such a scenario would include a corrupt server causing a client to exchange a key with another, unintended, server. It is clear that this yields a stronger security guarantee and we therefore choose the second formulation in our definitions.

We stress that our model also captures *key-compromise impersonation* resistance [BJM97]—where the adversary knows the long-term key of a party A and tries to impersonate another party to A —since our formalization also allows the owner of target sessions to be corrupted.

Definition 3.1 (Implicit key authentication). *The *iKeyAuth* predicate evaluates to 1 if and only if*

$$\forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \forall \ell' \in \text{LSID}, (\text{Samekey}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}),$$

where $\ell.\text{accept}$ holds true if and only if $\ell.\text{accept} = \text{true}$. We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides implicit key authentication if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{iKeyAuth}}}(1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{iKeyAuth}}}(1^\lambda) = 0 \right] = \text{negl}(\lambda),$$

where G_{iKeyAuth} is the same as G_{Match} with $P = \text{iKeyAuth}$.

Note that the predicate only applies to sessions that expect authentication, which is captured by the condition that $\ell.\text{pid} \in \mathcal{S}$. This models the fact that one can only provide authentication to keys if one possesses authenticating information. An artefact of this is that protocols without authentication (i.e. with $\mathcal{S} = \emptyset$) strictly speaking provide implicit key authentication as there is no authenticating party which the adversary can attack. Mathematically, this corresponds to a quantification over the empty set.

On the other end of the spectrum, we see that the case where $\mathcal{S} = \mathcal{I}$ rejoins mutual authentication where, upon completion, session ℓ has authenticated to session ℓ' and vice-versa. Indeed, if $\ell.\text{id} \in \mathcal{S}$ and also $\ell.\text{pid} \in \mathcal{S}$, then it is expected to authenticate itself to its intended peer in the same way that it expects to receive authentication. Upon both sessions completing, the predicate therefore induces a symmetry in the authentication guarantees.

As explained in the introduction, our basic definition considers the strongest adversarial model. We remark that we could relax the above requirement and only consider sessions ℓ with an honest peer, i.e. with $\ell.\delta_{\text{peer}} = \text{honest}$. Indeed, this notion sometimes appears in the literature: it still provides guarantees for parties who engage in sessions of the protocol with an honest peer as intended partner. Clearly, the notion neglects executions in which the intended peer is dishonest in which case one could be vulnerable to certain UKS attacks.

3.2 Key confirmation

Intuitively, this second notion is “the guarantee that another session possesses the same key.” While this does not provide authentication in the sense of binding an identity to a key, we define it here because its existential guarantee is a link between implicit and explicit authentication.

Here, we note that key confirmation only makes sense for honest peers because an adversary impersonating an honest party can always compute the key and provide confirmation to the target session. To prevent this trivial attack, we introduce a freshness condition on the peer.

Definition 3.2 (Authentication freshness). For any $\ell \in \text{LSID}$, the predicate $\text{aFresh}(\ell)$ evaluates to true if and only if

$$\ell.\delta_{\text{peer}} = \text{honest}.$$

We note that this freshness notion does not prohibit `Reveal` queries; this is because key authentication properties are expected to hold upon *derivation* of the key, and knowledge of the key should not benefit the adversary in breaking these. Furthermore, we show that the `Reveal` query is not useful for the adversary to wrongfully provide confirmation. Either another session derives the same key and the adversary reveals it, but then another session with the same key does exist, and therefore confirmation holds even though it does with the adversary’s intervention; or the adversary reveals the target session itself. However, the second option is not possible, since, in our model, setting the key to a non-trivial value is synonymous with accepting and terminating. Therefore the adversary cannot submit a `Reveal` query before the session has already accepted, at which point the adversary has already won if the session does not share a key with any other. Hence our freshness predicate does not need to eliminate trivial attacks using the `Reveal` query.

Fischlin et al. [FGSW16] introduced the distinction between *full* and *almost-full* key confirmation which captures the differences in guarantees that the last sender and last receiver in an AKE session can expect. We present here their predicate-based definitions and refer to [FGSW16] for a discussion. The former one says if an “authentication fresh” session with full key confirmation accepts, then there must be at least one other session also holding the key.

Definition 3.3 (Full key confirmation). *The fKeyConf predicate evaluates to 1 if and only if*

$$\forall \ell \in \text{LSID}, (\text{aFresh}(\ell) \wedge \ell.\text{kconf} = \text{full} \wedge \ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \exists \ell' \in \text{LSID} :: \text{Samekey}(\ell', \ell).$$

We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides full key confirmation if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}} (1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}} (1^\lambda) = 0 \right] = \text{negl}(\lambda),$$

where G_{fKeyConf} is defined similarly to G_{iKeyAuth} .

We see that the session's expected level of key confirmation, which is set when the session is first activated, is captured with the condition that $\ell.\text{kconf} = \text{full}$ — for a given protocol π , a session can decide which key confirmation to expect if it is activated as an initiator or a responder. Also, the session ℓ in question is excluded from the existence condition by the **Samekey** predicate and therefore a session cannot confirm its own key. We note that **fKeyConf** is only tested against sessions that expect authentication, with $\ell.\text{pid} \in \mathcal{S}$, as is discussed in [FGSW16, Section III.D]. Whilst this condition is not strictly required in the predicate for our result in Section 3.4, we adopt it here to align ourselves on the stand-alone definition of key confirmation.

As pointed out by Fischlin et al. [FGSW16], almost-full key confirmation is delicate to define. We adopt their notion saying that if such a fresh session accepts, then there must be a another session holding the same key-confirmation identifier and, moreover, if that other session has already derived a key, then it is the same one as the original session.

Definition 3.4 (Almost-full key confirmation). *The afKeyConf predicate is defined as*

$$\forall \ell \in \text{LSID}, (\text{aFresh}(\ell) \wedge \ell.\text{kconf} = \text{almost} \wedge \ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \exists \ell' \in \text{LSID} :: \left(\text{Samekcid}(\ell', \ell) \wedge [\ell'.\text{key} \neq \perp \implies \text{Samekey}(\ell', \ell)] \right).$$

We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides almost-full key confirmation if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{afKeyConf}}} (1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{afKeyConf}}} (1^\lambda) = 0 \right] = \text{negl}(\lambda)$$

where $G_{\text{afKeyConf}}$ is defined similarly to G_{iKeyAuth} .

3.3 Explicit authentication

This third notion is much stronger than the first. Indeed, we take *explicit* to mean that authentication is obtained *at termination*, and therefore it does not rely on the potential use of the key at a later time. In other words, session ℓ , upon accepting, knows that there is another session ℓ' which has the same key and whose identity is bound to it. Intuitively, it is a combination of implicit authentication and key confirmation, and indeed it was informally defined as such in Menezes, van Oorschot and Vanstone's *Handbook of Applied Cryptography* [MvOV97]; this appears in the logic of the predicates below.

Similarly to key confirmation, the existence of a session which has already derived the same key cannot always be guaranteed due to the asymmetry of the final message. We therefore define the two analogous notions of *full* and *almost-full explicit key authentication*.

As before, we study the requirements of a freshness predicate. As in the case of implicit authentication we do not stipulate that the peer is honest for the target session when it comes to the condition that any partner holding the same key is correctly identified ($\forall \ell' \in \text{LSID}, \text{Samekey}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}$). This again provides safety against all possible UKS attacks. Only for the “liveness” condition (i.e. that there exists a party with the same key) do we require that the intended peer is honest ($\text{aFresh}(\ell) \implies \exists \ell' \in \text{LSID} :: \text{Samekey}(\ell', \ell)$), otherwise the session may have communicated with an impersonating adversary which could trivially compute the key. Similarly to key confirmation, the **Reveal** query would not enable the adversary to conduct trivial attacks; this implies that the **aFresh** predicate is also the correct one here.

In summary, full explicit key authentication demands that for any fresh accepting session, any other session deriving the same key has the correct identity and there exists at least one other session holding the same key, if the peer is honest.

Definition 3.5 (Full explicit key authentication). *The fexKeyAuth predicate evaluates to 1 if and only if*

$$\begin{aligned} \forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{kconf} = \text{full} \wedge \ell.\text{accept}) \implies \\ (\forall \ell' \in \text{LSID}, \text{Samekey}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}) \\ \wedge (\text{aFresh}(\ell) \implies \exists \ell' \in \text{LSID} :: \text{Samekey}(\ell', \ell)). \end{aligned}$$

We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides full explicit key authentication if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{fexKeyAuth}}} (1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{fexKeyAuth}}} (1^\lambda) = 0 \right] = \text{negl}(\lambda)$$

where $G_{\text{fexKeyAuth}}$ is defined similarly to G_{iKeyAuth} .

We see that a session's expectation of both authentication and confirmation appears as $\ell.\text{pid} \in \mathcal{S}$ and $\ell.\text{kconf} = \text{full}$ in the predicate.

Definition 3.6 (Almost-full explicit key authentication). *The predicate afexKeyAuth evaluates to 1 if and only if*

$$\begin{aligned} \forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{kconf} = \text{almost} \wedge \ell.\text{accept}) \implies \\ (\forall \ell' \in \text{LSID}, (\text{Samekey}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid})) \\ \wedge \left(\text{aFresh}(\ell) \implies \exists \ell' \in \text{LSID} :: [\text{Samekcid}(\ell', \ell) \wedge (\ell'.\text{key} \neq \perp \implies \right. \\ \left. \text{Samekey}(\ell', \ell)) \right]). \end{aligned}$$

We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides almost-full explicit key authentication if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{afexKeyAuth}}} (1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{afexKeyAuth}}} (1^\lambda) = 0 \right] = \text{negl}(\lambda)$$

where $G_{\text{afexKeyAuth}}$ is defined similarly to G_{iKeyAuth} .

3.4 Equivalence results

We formally prove the coherence of the authentication definitions presented above. Namely, we show that a protocol which satisfies both implicit key authentication and key confirmation also satisfies explicit key authentication, and we show that the converse holds.

Theorem 3.1. *Let π be an AKE protocol. Then it holds for π that*

$$\text{iKeyAuth} \wedge \text{fKeyConf} \iff \text{fexKeyAuth}, \quad (5)$$

$$\text{iKeyAuth} \wedge \text{afKeyConf} \iff \text{afexKeyAuth}. \quad (6)$$

Proof. We first focus on equation (5) and show that $\text{iKeyAuth} \wedge \text{fKeyConf} \implies \text{fexKeyAuth}$; we proceed by proving the contrapositive. Let \mathcal{A} be a successful adversary against the fexKeyAuth predicate; i.e. \mathcal{A} reaches an execution state where $\neg \text{fexKeyAuth}$ holds true. This is equivalent to

$$\begin{aligned} \exists \ell^* \in \text{LSID} :: \ell^*.\text{pid} \in \mathcal{S} \wedge \ell^*.\text{kconf} = \text{full} \wedge \ell^*.\text{accept} \\ \wedge \left((\exists \ell' :: \text{Samekey}(\ell', \ell^*) \wedge \ell'.\text{id} \neq \ell^*.\text{pid}) \vee (\text{aFresh}(\ell^*) \wedge \forall \ell'', \neg \text{Samekey}(\ell'', \ell^*)) \right) \quad (7) \end{aligned}$$

Thus if $\neg\text{fexKeyAuth}$ holds true, either the first expression of the OR clause holds, which implies $\neg\text{iKeyAuth}$, or the second one holds and implies $\neg\text{fKeyConf}$. We therefore obtain that

$$\neg\text{fexKeyAuth} \implies \neg\text{iKeyAuth} \vee \neg\text{fKeyConf} \quad (8)$$

which completes the first part of the proof.

We now show that $\text{fexKeyAuth} \implies \text{iKeyAuth} \wedge \text{fKeyConf}$. We first show that $\text{fexKeyAuth} \implies \text{iKeyAuth}$. Let \mathcal{A} be a successful adversary against the iKeyAuth predicate; i.e. \mathcal{A} reaches an execution state where $\neg\text{iKeyAuth}$ holds true which is equivalent to

$$\exists \ell^* :: \ell^*.\text{pid} \in \mathcal{S} \wedge (\ell^*.\text{kconf} = \text{full}) \wedge \ell^*.\text{accept} \wedge \exists \ell' :: \text{Samekey}(\ell^*, \ell') \wedge (\ell'.\text{id} \neq \ell^*.\text{pid}). \quad (9)$$

Note that we include $\ell^*.\text{kconf} = \text{full}$ in $\neg\text{iKeyAuth}$ as we only aim to prove that fexKeyAuth implies iKeyAuth for sessions that expect full explicit key authentication. We now assume, for contradiction, that fexKeyAuth holds; this implies

$$\forall \ell, (\text{Samekey}(\ell, \ell^*) \implies \ell.\text{id} = \ell.\text{pid}) \wedge (\text{aFresh}(\ell^*) \implies \exists \ell'' :: \text{Samekey}(\ell^*, \ell'')) \quad (10)$$

for ℓ^* as in (9). We see that the existence of ℓ' that holds from (9) contradicts the first condition of (10) which shows that $\text{fexKeyAuth} \implies \text{iKeyAuth}$ as expected from the formulation of the predicates.

We now show that $\text{fexKeyAuth} \implies \text{fKeyConf}$. Let \mathcal{A} be a successful adversary against the fKeyConf predicate; i.e. \mathcal{A} reaches an execution state where $\neg\text{fKeyConf}$ holds true. This is equivalent to

$$\exists \ell^* :: \text{aFresh}(\ell^*) \wedge \ell^*.\text{pid} \in \mathcal{S} \wedge (\ell^*.\text{kconf} = \text{full}) \wedge \ell^*.\text{accept} \wedge \forall \ell, \neg\text{Samekey}(\ell^*, \ell).$$

This ℓ^* is now exactly one that satisfies $\neg\text{fexKeyAuth}$ and hence we immediately have that $\text{fexKeyAuth} \implies \text{fKeyConf}$. Since we have that fexKeyAuth implies both iKeyAuth and fKeyConf , combined with (8) this concludes the proof that

$$\text{iKeyAuth} \wedge \text{fKeyConf} \iff \text{fexKeyAuth}.$$

The proof of the same equivalence for almost-full confirmation notions, equation (6), follows from a similar argument. \square

4 Protocol examples

In this section we present established protocols and study which of our notions they achieve. We omit the secrecy proofs and sketch the proof for `Match` security to focus on the other properties.

Our results confirm that a “rule of thumb” for protocol design to achieve implicit key authentication is to include the parties’ identities in the key derivation step, $\text{key} = \text{KDF}(K, (\text{id}_A, \text{id}_B, \dots))$. If the key derivation function is collision-resistant then different identities immediately imply different session keys. Indeed this strategy has already been applied in very early protocols proposals, such as [BPR00], and has even been sometimes used to fix insecure protocols, e.g., [CBH05b]. We note that this method is also used in the TLS 1.3 protocol.

From the analysis of key confirmation in TLS 1.3 of [FGSW16], we see that a good strategy to obtain full or almost-full key confirmation is to send a MAC computed over a known value (such as the transcript) with a key derived from the same material as the final session key.

4.1 Plain Diffie-Hellmann

We begin with the plain Diffie-Hellmann (DH) protocol, presented in Figure 2, in which the parties exchange g^x and g^y to derive a key from g^{xy} and the communication transcript. One may also use the identities in the key derivation. The exchanged elements live in a cyclic group \mathbb{G} of prime order $|\mathbb{G}| = q$ with generator g such that $\langle g \rangle = \mathbb{G}$. We assume that this group is known to all parties. Since this is an unauthenticated protocol, we have $\mathcal{S} = \emptyset$.

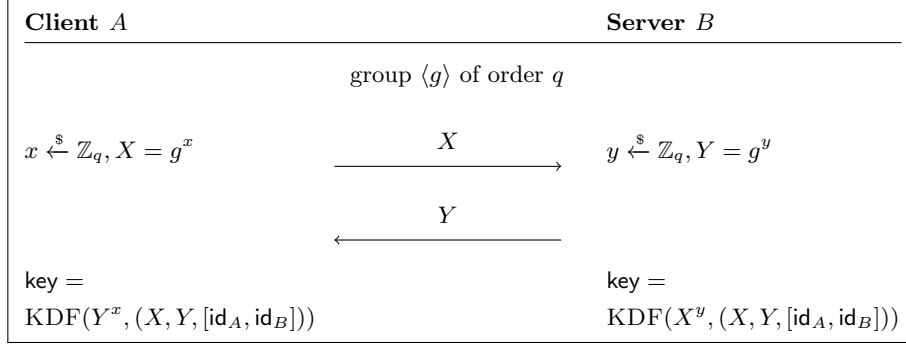


Fig. 2. The plain Diffie-Hellman protocol with identifiers $\text{sid} = \text{kcid} = (\text{id}_A, X, \text{id}_B, Y)$. Note that both parties know the identities id_A and id_B in advance. The notation $[\text{id}_A, \text{id}_B]$ means that the identities are optional.

Match security. The plain DH protocol provides **Match** security. Indeed both the session and key-confirmation identifiers fully determine the key (Properties 1 and 2). Furthermore, since the key-confirmation and session identifiers are identical, equal key-confirmation identifiers imply identical keys (Property 4). Finally, an honest party will contribute a random Diffie-Hellman share, such that the probability of matching any other share, is at most $n_s^2 \cdot \frac{1}{q}$ for a total number of n_s sessions, and thus negligible (Property 3).

Implicit key authentication. The plain DH protocol trivially provides implicit key authentication, as $\mathcal{S} = \emptyset$. We note that setting $\mathcal{S} \neq \emptyset$ would allow an adversary to break implicit key authentication as it could create a mismatch in the expected peer identities for any two sessions.

However, by including the identities in the key derivation function (as shown in Figure 2) this protocol can provide implicit key authentication even in the setting where $\mathcal{S} \neq \emptyset$. Indeed, the adversary has no control over a session’s owner identity $\ell.\text{id}$ which implies that if $\ell.\text{pid} \neq \ell'.\text{id}$ then $\ell.\text{key} = \ell'.\text{key}$ only if there is a collision in the KDF. In the random oracle model, or assuming a collision-resistant KDF, this happens only with negligible probability.

Key confirmation and explicit key authentication. As expected, this protocol does not provide key confirmation since no information about the key is exchanged after it is derived. We can show this formally by taking an adversary which initiates a session with an honest party (either client or server) without initiating a matching partner session. It then creates the message g^x or g^y to complete the exchange with the honest party. Obviously, there is then no other session which holds the same key nor the same key-confirmation identifier.

As the protocol does not provide key confirmation, Theorem 3.1 implies that it cannot provide explicit key authentication either. Intuitively this is clear since no authentication takes place and the adversary can substitute any value.

4.2 HMQV

We next come to one of the most prominent candidates for implicitly authenticated key exchange, the HMQV protocol [Kra05]. The idea here is to run a DH key exchange and to mix Schnorr-type signatures under the parties’ public keys in the key derivation. These signatures are not sent but only used locally, thus “implicitly” authenticating the key.

The protocol works over a group $\langle g \rangle = \mathbb{G}$ and uses a hash function H to compute the Schnorr signature. It is mutually authenticating, i.e. $\mathcal{S} = \mathcal{I}$, for which both parties use a long-term key. We assume that each party holds a certificate cert_i for its public key pk_i , and that the certificate is verified upon receiving it. We also assume that the public key and the owner’s identity can be recovered from the certificate. We set

Client A	group $\langle g \rangle$ of order q	Server B
$(a, A = g^a, \text{cert}_A)$		$(b, B = g^b, \text{cert}_B)$
$x \xleftarrow{\$} \mathbb{Z}_q, X = g^x$		$y \xleftarrow{\$} \mathbb{Z}_q, Y = g^y$
	$\xrightarrow{\text{cert}_A, X}$	
verify cert_B	$\xleftarrow{\text{cert}_B, Y}$	verify cert_A
$d = H(X, \text{cert}_B)$		$d = H(X, \text{cert}_B)$
$e = H(Y, \text{cert}_A)$		$e = H(Y, \text{cert}_A)$
$K_A = (YB^e)^{x+da}$		$K_B = (XA^d)^{y+eb}$
key = KDF(K_A)		key = KDF(K_B)

Fig. 3. HMQV protocol with session and key-confirmation identifiers $\text{sid} = \text{kcid} = (\text{cert}_A, X, \text{cert}_B, Y)$.

$\text{sid} = \text{kcid} = (\text{cert}_A, X, \text{cert}_B, Y)$. Since key derivation in HMQV is also determined by the transcript and the hash function, Match security follows as in the plain DH case.

Implicit key authentication. We provide a proof that the HMQV protocol achieves our strong notion of implicit key authentication and is therefore secure against all possible UKS and KCI attacks. Recall that we need to show that

$$\forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \forall \ell' \neq \ell \in \text{LSID}, (\text{Samekey}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}).$$

Due to the differences in the corruption queries that the adversary is allowed to make, the proof of secrecy for HMQV in [Kra05] is not immediately sufficient to imply our strong notion. Indeed, this proof holds only when the test-session in the secrecy experiment is fresh *in the sense of secrecy freshness* (see Section 5). An attack on implicit key authentication which would require the corruption of the owner before the session took place would therefore not be considered as valid against key secrecy and would not be ruled out by the proof of [Kra05].

Proposition 4.1. *The HMQV protocol provides unconditional mutual implicit key authentication, with*

$$\text{Adv}_{\mathcal{A}, \text{HMQV}, \mathcal{I}, \mathcal{I}}^{\text{G}_{\text{iKeyAuth}}}(1^\lambda) \leq \frac{n_i^2 \cdot n_s \cdot h}{2q} + \text{negl}(\lambda),$$

where h is the number of queries made to H , modelled as a random oracle.

Proof. To break the iKeyAuth predicate, it must be that a session $\ell_A = (A, B, *)$ shares a key with a session $\ell_C = (C, D, *)$ where $C \neq B$. This can happen either if $K_A = K_C$, or if $K_A \neq K_C$ but $\text{KDF}(K_A) = \text{KDF}(K_C)$. The later implies a collision in the KDF and we assume that this happens only with negligible probability. The only freedom that \mathcal{A} then has is to modify the Y value sent to ℓ_A as a response to its first message. Since the value of K_C, x, d and a are already fixed, \mathcal{A} must choose a value of Y such that YB^e , where $e = H(Y, \text{cert}_A)$ is exactly the right value such that $K_A = K_C$. Modelling H as a random oracle ensures that each value of Y yields a new random value of e and therefore that there is a probability of $1/q$ that a given value of Y will yield the correct value of YB^e . Given that there are at $n_i^2 \cdot n_s/2$ pairs of sessions, it holds that the adversary has at most a $n_i^2 \cdot n_s \cdot h/2q$ probability of finding a suitable Y for which the equality holds. \square

Similarly to the plain DH protocol, setting $\text{key} = \text{KDF}(K, (\text{id}_A, \text{id}_B))$ immediately provides implicit key authentication if the KDF is collision resistant.

Key confirmation and explicit key authentication. HMQV does not provide key confirmation in the same way that the plain DH does not. It immediately follows that the protocol does not provide explicit key authentication either.

4.3 TLS 1.3

We give a simplified version of the DH mode of the TLS 1.3 protocol suite in Figure 4 on page 16. Namely, we omit intermediate keys, in particular the handshake key and the encryption of the handshake protocol. We also look at server-only authentication, disallowing client authentication or anonymous executions.




Fig. 4. (Simplified) TLS 1.3 in mode (EC)DH, without handshake encryption and with server-only authentication. The session identifier and key-confirmation identifier are given by $\text{sid} = \text{kcid} = (r, X, S, Y, \text{cert}_B)$. Notation x, \dots, y means all transmitted communication data, ranging from x to y .

Match security and implicit key authentication. TLS 1.3 is **Match**-secure; the argument is identical to the plain DH case and appears in [DFGS15]. Implicit key authentication follows as for the identifier-including HMQV variant, if we assume that the KDF function is collision-resistant: since the server authenticates and cert_B appears in the key derivation, equal keys imply a correct authentication. Similarly to the HMQV protocol, the key secrecy of TLS would not suffice to imply implicit key authentication.

Key confirmation and explicit key authentication. Key confirmation for TLS 1.3 (**draft-10**) was shown in [FGSW16] and our simplified version corresponds to the analysed variant. The idea of the proof is that the parties use the key in the handshake protocol within the MAC. The property holds for full and almost-full key confirmation, such that our version also provides explicit key authentication.

5 Key Secrecy

In this section we define key secrecy in the BR-style via a security game and cast the main implication in terms of a logical formula.

5.1 BR-Style Key Secrecy

We recall the definition of secrecy for AKE protocols from [BFWW11]. Here the adversary is challenged, for a session of its choice, to distinguish between the honestly generated key or a randomly sampled one. To define the BR-secrecy game $G_{\text{BRSec}, \mathcal{D}}$, where \mathcal{D} denotes the distribution of keys, we use the same execution, session and local session states as in Section 2. The model state contains two bits, $b_{\text{test}} \in \{0, 1\}$ (initialised at random) and $b_{\text{guess}} \in \{0, 1, \perp\}$ (initialised to \perp) along with a session identifier $\ell_{\text{test}} \in \text{LSID} \cup \{\perp\}$ (initialised to \perp). The identifier ℓ_{test} is the local session to which the **Test** query is made (see below). The bit b_{test} determines whether the adversary receives the real key from ℓ_{test} , or a random value, in response to the **Test** query—we assume that only one **Test** query is made. The bit b_{guess} stores the adversary’s guess.

There are two additional queries, namely **Test** and **Guess**. The **Test**(ℓ) query sets $\ell_{\text{test}} \leftarrow \ell$ and returns $\text{key} = \ell.\text{key}$ if $b_{\text{test}} = 1$ or $\text{key} \xleftarrow{\$} \mathcal{D}$ otherwise. The query **Guess**(b) sets $b_{\text{guess}} \leftarrow b$. The **Valid** predicate requires that the **Test** query is made only once and to a session which has derived a key, and that only one **Guess** query is made. The adversary may therefore chose to submit queries that will trivially allow him to win the BR-secrecy game.

To catch this, if \mathcal{A} tests a session ℓ with $\ell.\delta_{\text{ownr}} = \text{corrupt}$ or with $\ell.\delta_{\text{peer}} = \text{corrupt}$, the **sFresh** predicate returns **false**. It also does so if the session, or any of its partners, has been revealed. We also take care of sessions which do not expect authentication as \mathcal{A} may impersonate the unauthenticated party and learn the session key. Hence, **sFresh** also returns **false** if the intended partner $\ell.\text{pid}$ belongs to the set $\mathcal{I} \setminus \mathcal{S}$ of unauthenticated parties. However, there is one exception: if there exists an honest session which is partnered to ℓ_{test} , even if it does not belong to the intended partner, then the session took place between two honest session and the key is still expected to remain secret.

Definition 5.1 (Secrecy freshness). *For any $\ell \in \text{LSID}$, the **sFresh**(ℓ) predicate evaluates to true if and only if,*

$$\begin{aligned} & (\ell.\delta_{\text{ownr}} = \ell.\delta_{\text{peer}} = \text{honest}) \wedge (\ell.\delta_{\text{sess}} = \text{fresh}) \\ & \wedge (\forall \ell' \in \text{LSID}, \text{Partner}(\ell', \ell) \implies \ell'.\delta_{\text{sess}} = \text{fresh}) \\ & \wedge [\ell.\text{pid} \in \mathcal{I} \setminus \mathcal{S} \implies (\exists \ell' \in \text{LSID} :: \text{Partner}(\ell, \ell'))] \end{aligned}$$

We could further relax the requirement for sessions with unauthenticated parties by defining almost-partnered sessions and allowing these to be tested. This would be the least strict requirement but we refrain from introducing another identifier and instead keep the simpler definition with partnering.

Definition 5.2 (BR-secrecy). The BRSec predicate is defined differently from the authentication ones due to its distinguishing nature. Instead of returning 0 or 1 to signify whether a certain condition holds, the BRSec predicates evaluates to $\text{MST}.b_{\text{guess}}$ if and only if

$$\text{MST}.l_{\text{test}} \neq \perp \wedge \text{sFresh}(\text{MST}.l_{\text{test}})$$

and evaluates to \perp otherwise. We also denote by $G_{\text{BRSec}, \mathcal{D}}^{b_{\text{test}}}$ the secrecy game with a specific value for b_{test} . We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} is BR-secret w.r.t. output key distribution \mathcal{D} if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}}(1^\lambda) = \left| \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}^0}(1^\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}^1}(1^\lambda) = 1 \right] \right|$$

is a negligible function in λ .

5.2 Key-Match Soundness

We now define the KMSoundness property which captures the essence of secrecy as a predicate over the core key-exchange mechanisms, without Test and Guess queries. It says that for any *authentication* fresh and accepting session ℓ , there does not exist another session ℓ' which holds the same key but is not partnered with ℓ .

Definition 5.3 (Key-Match Soundness). The KMSoundness predicate evaluates to 1 if and only if

$$\forall \ell \in \text{LSID}, (\text{aFresh}(\ell) \wedge \ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \forall \ell' \in \text{LSID} :: (\text{Samekey}(\ell', \ell) \implies \text{Partner}(\ell', \ell)).$$

We then say that the AKE protocol π with identity sets \mathcal{I}, \mathcal{S} provides key-match soundness if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{KMSoundness}}}(1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{KMSoundness}}}(1^\lambda) = 0 \right] = \text{negl}(\lambda)$$

where $G_{\text{KMSoundness}}$ is defined similarly to G_{iKeyAuth} .

The next theorem establishes that an adversary who can contradict Key-Match soundness must break either BR-secrecy or Match-security.

Theorem 5.1. Let π be an AKE protocol with Match security and BR secrecy w.r.t. \mathcal{D} . Then it also provides key-match soundness. More precisely, for any PPT algorithm \mathcal{A} attacking KMSoundness in at most n sessions, it holds that

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{KMSoundness}}}(1^\lambda) \leq n^2 \cdot \text{Adv}_{\mathcal{B}_2, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}}(1^\lambda) + \text{Adv}_{\mathcal{B}_1, \pi, \mathcal{I}, \mathcal{S}}^{G_{\text{Match}}}(1^\lambda) + 2^{-|\text{key}|}$$

for some PPT algorithms $\mathcal{B}_1, \mathcal{B}_2$ and the output length $|\text{key}|$ of keys.

Proof. The first observation is that Match security implies that any partnered session to either ℓ or to ℓ' must hold the same key as the corresponding session. If this would not hold with overwhelming probability, we could build an algorithm \mathcal{B}_1 to refute Match security in a straightforward way. This enables us to assume that there are two sessions ℓ_0 and ℓ'_0 with the above property and which accept first. That is, ℓ_0 and ℓ'_0

- hold identical keys, $\text{Samekey}(\ell_0, \ell'_0)$,
- are not partnered, $\neg \text{Partner}(\ell_0, \ell'_0)$, and
- for neither of the two sessions, in the moment when the session accepts, there is another session which is yet partnered with the session.

Note that ℓ_0 must have accepted by assumption, such that $\ell_0.\text{key} \neq \perp$, and therefore $\text{Samekey}(\ell_0, \ell'_0)$ implies that session ℓ'_0 , too, must have a valid key. In particular it must have accepted (and can both be tested and revealed in an attack on secrecy).

Assume now that there was a successful adversary \mathcal{A} against key-match soundness (with two sessions ℓ_0, ℓ'_0 as above). We show how to break BRSec through an adversary \mathcal{B}_2 with non-negligible probability in this case.

Our adversary \mathcal{B}_2 will try to predict the sessions ℓ_0, ℓ'_0 by picking two session numbers i, j at random from $\{1, 2, \dots, n\}$, where we count sessions according to their initialisation in \mathcal{A} 's simulated attack. Next, \mathcal{B}_2 runs \mathcal{A} 's attack, relaying all inputs and oracle queries and answers between \mathcal{B}_2 's game and \mathcal{A} . Note that \mathcal{B}_2 has the same oracle interfaces as \mathcal{A} , but in addition may call the **Test** and the **Guess** oracle.

Adversary \mathcal{B}_2 diverges from \mathcal{A} with respect to two points: If the i -th session in the attack accepts, then \mathcal{B}_2 immediately asks to **Reveal** the session key key_i . If the j -th session accepts, then \mathcal{B}_2 immediately calls **Test** to get a key value key_j . Adversary \mathcal{B}_2 makes the **Guess**(b_{guess}) query and stops, where the bit b_{guess} is set to 1 if $\text{key}_i = \text{key}_j$, and to 0 otherwise.

Note that up to the point when \mathcal{B}_2 makes the **Test** query, the simulation to \mathcal{A} is perfect. Assume that \mathcal{B}_2 predicts ℓ_0, ℓ'_0 correctly for the **Test** query resp. the **Reveal** query, which happens with probability at least $1/n^2$. Then the **Test** session ℓ_0 does not have a partner yet, is authentication fresh and the authenticating partner is still honest, such that the session is still secrecy fresh. In this case, if $b_{\text{test}} = 1$ the **Test** oracle returns the actual session key, such that the keys match, and we have \mathcal{B}_2 output 1, too. In summary, the probability of this happening is at least $\frac{1}{n^2}$ times the probability that \mathcal{A} succeeds, minus a negligible term for refuting **Match** security. This is non-negligible.

Next consider the case $b_{\text{test}} = 0$ such that the test session ℓ_0 returns a random key. Then the probability that this independent random key matches the other key is $2^{-|\text{key}|}$ and thus negligible, such that \mathcal{B}_2 only returns 1 with this negligible probability. \square

6 Relation with CK-style Security

In all CK-style models (CK, CK_{HMQV} and eCK) the definition of “matching sessions” includes, among others, the requirement that parties agree on each other’s identities. As there is no other mention of matching expected identities, this seems to be the only capture of authentication in such models.

Capturing authentication in CK-style models. Some light is shed on this by Cremers’s argument in [Cre11]. He states that since “the test session-key must be indistinguishable from keys computed by non-matching sessions”, then sessions that compute the same key must be matching sessions. This is analogous to our Theorem 5.1 concerning **KMSoundness** except that the CK-style definition of matching also includes expected identities, and therefore this result implies that CK-style security also guarantees (at least) implicit authentication.

Crucially, this argument however only applies for sessions for which CK-style security holds, that is sessions that remain $\text{sFresh}_{\text{CK}}$, for a suitable definition of this extended secrecy freshness predicate that captures the restriction of the CK-style models as presented in [Cre11]. In contrast, our standalone definition of implicit authentication (Definition 3.1) captures security against a much wider range of attacks due to its absence of freshness requirements on the target session.

Extending authentication freshness. To establish our authentication definitions of Section 3 in a CK-style model, we consider whether the freshness conditions changes upon the addition of a **StateReveal** query. This reveals to the adversary either the entire state of a session or only its ephemeral key, as defined by the protocol, for the CK and eCK models respectively [Cre11].

Following from the discussion of Section 3.1, we allow the adversary to perform **StateReveal** queries when attacking the $\text{iKeyAuth}_{\text{CK}}$ predicate as we wish to capture the widest possible range of attacks. We recall that this implies that \mathcal{A} can both **StateReveal** a session and **Corrupt** its owner during its attack, which is not allowed by the freshness conditions of CK-style models. Following similarly from Sections 3.2 and 3.3,

we have a restriction that the adversary cannot trivially obtain the key when attacking key confirmation and explicit key authentication. Therefore we add the condition that the adversary cannot both corrupt the intended partner and also `StateReveal` a partner session for attacks against CK variants of these predicates.

Separating authentication from secrecy. Let \mathcal{E} denote the event, in a CK-style secrecy experiment, where the adversary succeeds in causing two game-controlled sessions to share a key without matching in the CK-style sense, i.e. without agreeing on each other's identities. Let G_{CK} denote the usual CK-style secrecy game and let also G_{CK}^- be the secrecy game which instead uses our definition of partnering of Section 2.2 without expected identities (i.e. an extension of the secrecy game of Section 5.1). Using a rather informal terminology, we then have

$$\Pr[\mathcal{A} \text{ wins } G_{\text{CK}}] = \Pr[\mathcal{A} \text{ wins } G_{\text{CK}}|\mathcal{E}] \cdot \Pr[\mathcal{E}] + \Pr[\mathcal{A} \text{ wins } G_{\text{CK}}|\neg\mathcal{E}] \cdot \Pr[\neg\mathcal{E}].$$

First, \mathcal{E} corresponds exactly to a break of `iKeyAuthCK` and so we have

$$\Pr[\mathcal{A} \text{ wins } G_{\text{CK}}|\mathcal{E}] \cdot \Pr[\mathcal{E}] \leq \Pr[\neg\text{iKeyAuth}_{\text{CK}}].$$

Second, if \mathcal{A} wins G_{CK} without triggering \mathcal{E} , then his attack can be reproduced in G_{CK}^- and so we have

$$\Pr[\mathcal{A} \text{ wins } G_{\text{CK}}|\neg\mathcal{E}] \cdot \Pr[\neg\mathcal{E}] \leq \Pr[\mathcal{A} \text{ wins } G_{\text{CK}}^-].$$

In conclusion,

$$\Pr[\mathcal{A} \text{ wins } G_{\text{CK}}] \leq \Pr[\neg\text{iKeyAuth}_{\text{CK}}] + \Pr[\mathcal{A} \text{ wins } G_{\text{CK}}^-],$$

which shows that we can represent CK-style key secrecy through our separate notions of authentication and key secrecy when adapted to these models.

7 Key-confirming protocols

In this section we define the class of symmetric *key-confirming protocols*. These are protocols which provide guarantees on the existence of a session with the same key. We note that this guarantee may be secondary to the main functionality of these protocols. For example, we would expect that protocols for authenticated message transmission would belong to this class.

Based on the model of [BFWW11, Section 4], we describe the syntax for such protocols and the security game for key confirmation. We formalise these protocols as $\pi = (\text{kg}, \zeta)$ and write \mathcal{D}_{kg} to denote the output distribution of the randomized key generation algorithm `kg`; we use the same mechanism of local session identifiers. For the game execution state, as there are not long-term keys, `EST` is not defined.

Session state. The session state for key-confirming protocols consists of:

- `crypt` $\in \{0, 1\}^*$ is some protocol-specific private session state.
- `key` $\in \{0, 1\}^* \cup \{\perp\}$ is the symmetric key used by the protocol.
- `kcind` $\in \{\text{true}, \text{false}, \perp\}$ indicates key confirmation achieved. Initially set to \perp , it must be changed to `true` or `false` before termination. The value of `kcind` is always public.

The difference with [BFWW11] is the addition of the *key confirmation indicator* `kcind` which informs the owner if key confirmation has been obtained. We stress that setting `kcind` is done independently of termination. For example, a secure channel protocol could achieve key confirmation with the first messages exchanged but continue running for much longer as the channel is used for communication. We will only focus on guarantees related to the setting of `kcind` and will not make assumptions or requirements on the session's termination.

Local session state. As in [BFWW11], the local session state consists of:

- $\delta_{\text{key}} \in \{\text{fresh}, \text{revealed}\}$ denotes whether the key is known to the adversary.
- `lst` $\in \{0, 1\}^*$ is any other local session state required to model the protocol's other security requirements.

Setup. The `setupE` algorithm only initialises `crypt`, `key` and `kcind` to \perp for each $\ell \in \text{LSID}$. The `setupG` algorithm also only initialises $\delta_{\text{key}} \leftarrow \text{fresh}$ for every session as our security game for key confirmation does not require any model-wide state.

Queries. As in [BFWW11], our model allows \mathcal{A} to initialise sessions with three different queries. The first, `InitS`(ℓ), initialises a session with an honestly generated key, $\ell.\text{key} \leftarrow \text{kg}(1^\lambda)$, which remains outside of the adversary’s view. The second, `InitP`(ℓ_1, ℓ_2), initialises a session with the same key as another. Here, the game sets $\ell_2.\text{key} \leftarrow \ell_1.\text{key}$ and $\ell_2.\delta_{\text{key}} \leftarrow \ell_1.\delta_{\text{key}}$. The third query, `InitK`(ℓ, κ), allows \mathcal{A} to set his own key for a given session. It sets $\ell.\text{key} \leftarrow \kappa$ and immediately sets $\ell.\delta_{\text{key}} \leftarrow \text{revealed}$ to note that \mathcal{A} already knows the key. As before, `Send`(ℓ, m) and `Reveal`(ℓ) allow \mathcal{A} to control the network and view honestly generated keys.

The **Valid** predicate verifies that `Send` and `Reveal` queries are only made to sessions which have been initialised and that initialisation queries are only made to sessions for which the key has not yet been set. For the `InitP` query, it also verifies that ℓ_1 has already been initialised.

Key confirmation guarantee. We note that here there no longer exists a distinction between full and almost-full key confirmation since keys are set upon initialisation. This notion says that for any session which has set the key confirmation identifier to true, there is another session which uses the same key.

Definition 7.1 (Key confirmation guarantee).

The `symKeyConf` predicate evaluates to 1 if and only if

$$\forall \ell \in \text{LSID}, (\ell.\delta_{\text{key}} = \text{fresh} \wedge \ell.\text{kcind} = \text{true}) \implies \exists \ell' \in \text{LSID} :: \text{Samekey}(\ell', \ell),$$

where `Samekey` is defined as for AKE protocols, including the condition that $\ell' \neq \ell$ ⁸. The game $G_{\text{symKeyConf}}$ is then defined with state, `setupE`, `setupG` and behaviour as above, with query set $Q = \{\text{Send}, \text{InitS}, \text{InitP}, \text{InitK}, \text{Reveal}\}$ and winning predicate $P = \text{symKeyConf}$.

The protocol π provides (secure) key confirmation, or is a key-confirming protocol, if, for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}, \pi, \mathcal{I}}^{G_{\text{symKeyConf}}}(1^\lambda) = \Pr \left[\text{Exp}_{\mathcal{A}, \pi, \mathcal{I}}^{G_{\text{symKeyConf}}}(1^\lambda) = 0 \right] = \text{negl}(\lambda).$$

We note that a symmetric protocol π which always sets `kcind` = false trivially achieves secure key confirmation. This is similar to an AKE protocol formally achieving implicit key authentication by setting $\mathcal{S} = \emptyset$.

Protocol example. We present an example of key-confirming protocols. Let $\mathcal{M} = (\text{kg}, \text{MAC}, \text{Vf})$ be an unforgeable message authentication code (MAC) which requires that, for any PPT adversary \mathcal{A} ,

$$\Pr \left[\begin{array}{l} \text{key} \leftarrow \text{kg}(\lambda), \\ \text{Vf}(\text{key}, m, t) = 1; (m, t) \leftarrow \mathcal{A}^{\text{MAC}(\text{key}, \cdot)}(\lambda), \\ m \notin \mathcal{Q} \end{array} \right] \leq \text{negl}(\lambda),$$

where $\text{MAC}(\text{key}, \cdot)$ denotes access to a tagging oracle and \mathcal{Q} denotes the messages queried by \mathcal{A} for tagging. From such a MAC, we construct the protocol π_{kconf} as follows. If a session is activated with $m = \text{init}$, it sends $t = \text{MAC}(\text{key}, 1)$. When it receives a second message m^* , it verifies that it is a tag for the message “2” by checking if $\text{Vf}(\text{key}, 2, m^*) = 1$. If this holds, then it sets `kcind` \leftarrow true, otherwise it sets `kcind` \leftarrow false. If a session is instead activated as a receiver, then it plays the counterpart and checks that it correctly receives a tag for the message “1” and, if so, sets `kcind` \leftarrow true and replies with a tag for “2”. We present the formal description of π_{kconf} in Figure 5.

To show that the protocol π_{kconf} is key-confirming, we use the concept of single-session reducible games presented in [BFWW11]. Without specifying the formal details, we see that the independence of sessions in the game $G_{\text{symKeyConf}}^{\pi_{\text{kconf}}}$, apart from their potential partner, implies that this game is session restricted.

⁸ We note that $\ell.\delta_{\text{key}} = \text{fresh}$ is analogous to a freshness condition.

π_{kconf} (with secret key)	
1 :	Initialise $\text{kconf} \leftarrow \perp$ and role $\rho \leftarrow \perp$
2 :	while $\text{kconf} = \perp$ do
3 :	Receive m^*
4 :	if $m^* = \text{init}$ then
5 :	Set $\rho \leftarrow \text{init}$
6 :	Send $t = \text{MAC}(\text{key}, 1)$
7 :	elseif $\rho = \text{init}$ then
8 :	if $\text{Vf}(\text{key}, 2, m^*) = 1$ then
9 :	Set $\text{kconf} \leftarrow \text{true}$
10 :	else set $\text{kconf} \leftarrow \text{false}$
11 :	else
12 :	if $\text{Vf}(\text{key}, 1, m^*) = 1$ then
13 :	Set $\text{kconf} = \text{true}$
14 :	Send $t = \text{MAC}(\text{key}, 2)$
15 :	else set $\text{kconf} \leftarrow \text{false}$

Fig. 5. A simple key confirmation protocol that expects a tag on the message “1” or “2” depending on the role (initiator or responder) played by the session.

Theorem 2 of [BFWW11, Appendix B] then gives us that $G_{\text{symKeyConf}}^{\pi_{\text{kconf}}}$ is single session reducible and therefore that the key-confirmation property of π_{kconf} depends only on the security of a single session.

We then reduce the security of one session to the unforgeability of the MAC \mathcal{M} . The reduction sets up a session and waits for the adversary to submit a message. If it submits $m^* = \text{init}$, then the reduction uses the oracle $\text{MAC}(\text{key}, \cdot)$ to respond with a correct tag. If it submits any other message, then the reduction submits $(1, m^*)$ or $(2, m^*)$ as its forgery, depending on its role. If the reduction has queried the MAC oracle on “1” in response to an init message, then the adversary must create a tag for the message “2” to make π_{kconf} accept and therefore there is no risk of the reduction outputting a message which it has already queried. This shows that the reduction creates a forgery exactly when the adversary is capable of winning $G_{\text{symKeyConf}}^{\pi_{\text{kconf}}}$ and thus π_{kconf} provides secure key confirmation if \mathcal{M} is an unforgeable MAC.

While this protocol is a simple example that has no objective beyond providing secure key confirmation, it nonetheless provides justification for expecting useful constructions, such as authenticated encryption schemes of secure channel protocols, to provide the same guarantees.

8 Key confirming protocols give explicit authentication

In this section we define the composition of an AKE protocol with a key-confirming protocol, similarly to [BFWW11, Section 5]. As a significant difference, we consider the composed protocol *as an AKE protocol*, and not as a symmetric protocol, and we then prove that composing an AKE protocol with implicit key authentication and key secrecy together with a secure key-confirming protocol yields an *AKE protocol* with *explicit* key authentication.

Syntax of composed protocols. The composed protocol first runs the AKE protocol and then, once this accepts, it runs the symmetric protocol, initialised with the key derived in the first step, until it obtains key confirmation.

We first recall that a key-confirming protocol $\pi = (\text{kg}_\pi, \zeta_\pi)$ must set `kcind` before terminating and that a session can then continue its execution; this is not the case when forming an AKE protocol by composing with π since it must terminate upon acceptance of the key, as in Section 2. We therefore define the protocol $\bar{\pi}$ to be π with the algorithm $\zeta_{\bar{\pi}}$ the same as ζ_π but halted after `kcind` is set. Thus the key derived by $\zeta_{\bar{\pi}}$ is only accepted as the final key for the composed protocol once $\zeta_{\bar{\pi}}$ has terminated with `kcind` set to `true`. Now, given an AKE protocol $\text{ke} = (\text{kg}_{\text{ke}}, \zeta_{\text{ke}})$, we write $\text{ke}; \bar{\pi} = (\text{kg}_{\text{ke}; \bar{\pi}}, \zeta_{\text{ke}; \bar{\pi}})$ for the composition.

As we consider $\text{ke}; \bar{\pi}$ as an AKE protocol, it uses the same long-term key generation as ke and therefore $\text{kg}_{\text{ke}; \bar{\pi}} = \text{kg}_{\text{ke}}$. The algorithm $\zeta_{\text{ke}; \bar{\pi}}$ first runs ζ_{ke} . If it rejects, then $\zeta_{\text{ke}; \bar{\pi}}$ rejects the entire session; otherwise it runs $\zeta_{\bar{\pi}}$ with the derived key and accepts or rejects depending on `kcind`. Formally, given an incoming message m , the algorithm $\zeta_{\text{ke}; \bar{\pi}}$ first examines the value of `acceptke` from the ke protocol. If `acceptke` = \perp , it passes m to ζ_{ke} ; if `acceptke` = `true`, it passes m to $\zeta_{\bar{\pi}}$. We note that this value of `acceptke` from ke is different from the value of `acceptke; \bar{\pi}}` for $\text{ke}; \bar{\pi}$ which is only set once both ke and $\bar{\pi}$ have terminated.

Syntax of composed games. Our composed game $G_{\text{Pred}}^{\text{ke}; \bar{\pi}}$ enables the adversary to interact with simultaneous sessions of the composed protocol. We stress that the adversary’s goal here is to “break” the authentication property `Pred` of $\text{ke}; \bar{\pi}$ seen as an AKE protocol. We build our composed game from the elements of the games for the protocols ke and π and use indices to distinguish them.

Game state. The composed execution state $\text{EST}_{\text{ke}; \bar{\pi}}$ is set to be EST_{ke} as key-confirming protocols do not have one. The session state $\text{SST}_{\text{ke}; \bar{\pi}}$ of a session ℓ is made up of the same elements as for key exchange protocols but constructed from the composing session states in the following way:

- The long-term keying information $(\text{pk}_i, \text{sk}_i)$ and pk_j is the same as in SST_{ke} .
- The protocol-specific private session state is defined as the concatenation of both protocols’ states: $\text{crypt}_{\text{ke}; \bar{\pi}} = \text{crypt}_{\text{ke}} \parallel \text{crypt}_{\bar{\pi}}$.
- The `acceptke; \bar{\pi}}` indicator is set to `true` once `kcind\bar{\pi}}` is set to `true`. Note that this only happens if `acceptke` is also already set to `true` as $\bar{\pi}$ does not begin until ke has accepted.
- The session identifier `sidke; \bar{\pi}` is set to `sidke` *only* when `acceptke; \bar{\pi}}` is set to `true`.
- The `keyke; \bar{\pi}` is set to `keyke` *only* when `acceptke; \bar{\pi}}` is set to `true`. Before then, the value of `keyke` is kept internally and passed on to $\bar{\pi}$, so that `key\bar{\pi}}` \leftarrow `keyke`, when `acceptke` \leftarrow `true`.
- As $\bar{\pi}$ always provides full key confirmation, we have that `kconfke; \bar{\pi}}` = `true` for all sessions.

The local session state $\text{LST}_{\text{ke}; \bar{\pi}}$ remains the same as for AKE protocols and the model state remains undefined as not required for authentication.

Setup, queries and Valid predicate. These are the same as described in Section 2 with the addition that the **Valid** predicate makes use of the **Valid_{\bar{\pi}}** predicate for `Send` and `Reveal` queries to sessions that are executing $\bar{\pi}$.

Winning predicate. We can use any of the predicates defined in Sections 2 and 3.

Composition result. We show that a implicitly authenticated key exchange protocol composed with a key-confirming protocol produces an explicitly authenticated key exchange protocol. Our choice of public session identifiers means we do not require a session matching algorithm as in [BFWW11, Section 3].

Theorem 8.1. *Let ke be a Match-secure key exchange protocol which provides implicit key authentication and BR-secrecy w.r.t. key distribution \mathcal{D} . Let π be a symmetric-key protocol with key generation distribution \mathcal{D} which provides secure key confirmation. Then $\text{ke}; \bar{\pi}$ is a key exchange protocol with provides explicit key authentication.*

Proof. We make use of Theorem 3.1 to separate our work into two steps. First we prove in Lemma 8.1 that $\text{ke}; \bar{\pi}$ provides implicit key authentication under the assumption that ke does. Then we prove in Lemma 8.2 that $\text{ke}; \bar{\pi}$ provides full key confirmation under the assumption that ke is BR-secret and that $\bar{\pi}$ provides secure key confirmation. As the two property hold separately, Theorem 3.1 immediately gives us that $\text{ke}; \bar{\pi}$ provides explicit key authentication. \square

Lemma 8.1. *Let ke be a key exchange protocol and let π be a symmetric-key protocol. For any PPT adversary \mathcal{A} , it holds that*

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{iKeyAuth}}^{\text{ke}; \bar{\pi}}}(1^\lambda) = \text{Adv}_{\mathcal{B}, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{iKeyAuth}}^{\text{ke}}}(1^\lambda)$$

for some PPT algorithm \mathcal{B} .

Proof. Let \mathcal{A} be an adversary against $\text{ke}; \bar{\pi}$ in the $G_{\text{iKeyAuth}}^{\text{ke}; \bar{\pi}}$ game, which we refer to $G^{\text{ke}; \bar{\pi}}$ when the context is clear. We build an adversary \mathcal{B} against ke in the $G_{\text{iKeyAuth}}^{\text{ke}}$ game, which we similarly refer to G^{ke} .

\mathcal{B} sets up $G^{\text{ke}; \bar{\pi}}$ for \mathcal{A} as described above using elements from G^{ke} . It then responds to \mathcal{A} 's queries in the following way. (We use the notation $\ell_{\text{ke}; \bar{\pi}}$ to denote session identifiers used by \mathcal{A} in $G^{\text{ke}; \bar{\pi}}$ and the notation ℓ_{ke} to denote the corresponding identifiers used by \mathcal{B} in G^{ke} .)

- When \mathcal{A} submits $\text{Send}(\ell_{\text{ke}; \bar{\pi}}, m)$, \mathcal{B} checks the value of $\ell_{\text{ke}; \bar{\pi}}.\text{accept}$. If it is either true or false, \mathcal{B} responds \perp to \mathcal{A} as the sessions has either already accepted or rejected. If it is still \perp , \mathcal{B} examines the value of $\ell_{\text{ke}}.\text{accept}$.
 - If $\ell_{\text{ke}}.\text{accept} = \perp$, \mathcal{B} submits $\text{Send}(\ell_{\text{ke}}, m)$ to G^{ke} and responds to \mathcal{A} with m' returned by G^{ke} . If $\ell_{\text{ke}}.\text{sid}$ is set at that step, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{sid} \leftarrow \ell_{\text{ke}}.\text{sid}$. If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{false}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{false}$ and makes this known to \mathcal{A} . If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{true}$, \mathcal{B} submits $\text{Reveal}(\ell_{\text{ke}})$ to G^{ke} to obtain $\ell_{\text{ke}}.\text{key}$.
 - If $\ell_{\text{ke}}.\text{accept} = \text{true}$ and $\ell_{\bar{\pi}}.\text{kcind} = \perp$, \mathcal{B} has obtained $\ell_{\text{ke}}.\text{key}$ so it can respond to \mathcal{A} according to $\bar{\pi}$ by computing the response internally. If $\ell_{\bar{\pi}}.\text{kcind} \leftarrow \text{false}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{false}$. If $\ell_{\bar{\pi}}.\text{kcind} \leftarrow \text{true}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{true}$ and $\ell_{\text{ke}; \bar{\pi}}.\text{key} \leftarrow \ell_{\text{key}}.\text{key}$.
 - The case of $\ell_{\text{ke}}.\text{accept} = \text{false}$ is never examined as it would already hold that $\ell_{\text{ke}; \bar{\pi}}.\text{accept} = \text{false}$.
- When \mathcal{A} submits $\text{Reveal}(\ell_{\text{ke}; \bar{\pi}})$, \mathcal{B} checks the value of $\ell_{\text{ke}; \bar{\pi}}.\text{accept}$. If it is either \perp or false, \mathcal{B} responds \perp to \mathcal{A} as his Reveal query is invalid. If it is true, then $\ell_{\text{ke}; \bar{\pi}}.\text{key}$ was set when $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{true}$ so \mathcal{B} responds with $\ell_{\text{ke}; \bar{\pi}}.\text{key}$ to \mathcal{A} and sets $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{sess}} \leftarrow \text{revealed}$.
- When \mathcal{A} submits $\text{Corrupt}(i)$, \mathcal{B} submits $\text{Corrupt}(i)$ to G^{ke} and receives sk_i which it returns to \mathcal{A} . At that moment, G^{ke} will mark the values of $\ell_{\text{ke}}.\delta_{\text{ownr}}$ and $\ell_{\text{ke}}.\delta_{\text{peer}}$ as **corrupt** for relevant ℓ_{ke} as described in Section 2. However, \mathcal{B} will not update the corresponding sessions in $G^{\text{ke}; \bar{\pi}}$ in the same way as this would leak information to \mathcal{A} about the internal stage of the sessions. Instead, \mathcal{B} marks the values as **corrupt** for the sessions $\ell_{\text{ke}; \bar{\pi}}$ which have not completed the entire composed protocol, even if they have already completed the key exchange protocol and would not be marked as **corrupt** in G^{ke} .

We now argue that if \mathcal{A} is able to reach an execution state in $G^{\text{ke}; \bar{\pi}}$ for which the iKeyAuth predicate evaluates to 0, then \mathcal{B} , by behaving as described above, reaches a state in G^{ke} for which the iKeyAuth predicate also evaluates to 0. This means that the composed protocol preserves implicit key authentication. If \mathcal{A} reaches such a state, then we have that

$$\begin{aligned} \exists \ell_{\text{ke}; \bar{\pi}} \in \text{LSID}_{\text{ke}; \bar{\pi}} :: (\ell_{\text{ke}; \bar{\pi}}.\text{pid} \in \mathcal{S} \wedge \ell_{\text{ke}; \bar{\pi}}.\text{accept}) \\ \wedge (\exists \ell'_{\text{ke}; \bar{\pi}} \in \text{LSID}_{\text{ke}; \bar{\pi}} :: \text{Samekey}(\ell'_{\text{ke}; \bar{\pi}}, \ell_{\text{ke}; \bar{\pi}}) \wedge \ell'_{\text{ke}; \bar{\pi}}.\text{id} \neq \ell_{\text{ke}; \bar{\pi}}.\text{pid}) . \end{aligned}$$

We show that this also holds for the corresponding sessions ℓ_{ke} and ℓ'_{ke} in G^{ke} . We first have that $\ell_{\text{ke}}.\text{pid} \in \mathcal{S}$ as all the sessions and the set \mathcal{S} match one-to-one between the two games. We then have that $\ell_{\text{ke}}.\text{accept} = \text{true}$ as $\ell_{\text{ke}; \bar{\pi}}.\text{accept}$ is set to true only if the ke session accepts, and \mathcal{B} relays \mathcal{A} 's Send queries exactly which causes ℓ_{ke} to accept in G^{ke} .

As per the definition of the composed protocol $\text{ke}; \bar{\pi}$, the final key is fixed as soon as the ke part completes, therefore it holds that if two sessions accept with the same key in $G^{\text{ke}; \bar{\pi}}$, then they have derived that same key in the first part. As \mathcal{B} relays \mathcal{A} 's queries exactly, we have that $\text{Samekey}(\ell'_{\text{ke}}, \ell_{\text{ke}})$ holds in G^{ke} for the sessions corresponding to $\ell'_{\text{ke}; \bar{\pi}}$ and $\ell_{\text{ke}; \bar{\pi}}$. Furthermore, $\ell'_{\text{ke}}.\text{id} \neq \ell_{\text{ke}}.\text{pid}$ also holds as these values are the same as the ones for the sessions in $G^{\text{ke}; \bar{\pi}}$. This shows that the following holds for the corresponding sessions:

$$\begin{aligned} \exists \ell_{\text{ke}} \in \text{LSID}_{\text{ke}} :: (\ell_{\text{ke}}.\text{pid} \in \mathcal{S} \wedge \ell_{\text{ke}}.\text{accept}) \\ \wedge (\exists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: \text{Samekey}(\ell'_{\text{ke}}, \ell_{\text{ke}}) \wedge \ell'_{\text{ke}}.\text{id} \neq \ell_{\text{ke}}.\text{pid}) , \end{aligned}$$

which implies that \mathcal{B} is a successful adversary for the game $G_{\text{iKeyAuth}}^{\text{ke}}$ exactly when \mathcal{A} is successful for the game $G_{\text{iKeyAuth}}^{\text{ke};\bar{\pi}}$. \square

Lemma 8.2. *Let ke be a Match-secure key exchange protocol with output key distribution \mathcal{D} . Let π be a symmetric-key protocol with key generation distribution \mathcal{D} . Let $n = n_i^2 \cdot n_s$. For any PPT adversary \mathcal{A} , it holds that*

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}}}(1^\lambda) \leq n \cdot \text{Adv}_{\mathcal{B}_1, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}}(1^\lambda) + \text{Adv}_{\mathcal{B}_2, \pi, \mathcal{I}}^{G_{\text{SymKeyConf}}}(1^\lambda),$$

for some PPT algorithms \mathcal{B}_1 and \mathcal{B}_2 .

Proof. We use a strategy similar to the proof of Theorem 1 in [BFWW11], namely we first replace all the keys derived by the ke part of the composed protocol by randomly sampled keys from the correct distribution, using BR-secrecy to show that the final game is indistinguishable from the first. Then we show, similarly to Lemma 8.1, that if an adversary manages to break the key confirmation property of the composed protocol, then a reduction can break the key confirmation property of the symmetric protocol π .

To replace all the keys used, we proceed with a hybrid argument. Let the game $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}$ be the game G_{fKeyConf} played against protocol $\text{ke}; \bar{\pi}$, where the first Σ sessions to accept a new key, i.e. where a partner session has not already accepted a key, have their keys from ke replaced by a random value from \mathcal{D} for the π part, where $\mathcal{D} = \mathcal{D}_{\text{kg}}$ is the output distribution of the key generation algorithm for π . We remove the mention of fKeyConf when the context is clear. The original game G_{fKeyConf} for \mathcal{A} is therefore $G^{\text{ke}; \bar{\pi}, 0, \mathcal{D}}$ where only honestly computed keys are used for π .

The game $G^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}$ runs just as $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}}$ does with the following modifications. It maintains a counter σ to keep track of the number of new keys that are accepted (not counting those which the adversary might already know by corrupting one of the parties); this is set to 0 initially. The behaviour of $G^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}$ is then the same as $G^{\text{ke}; \bar{\pi}}$ with the following differences to the $\text{Send}(\ell_{\text{ke}; \bar{\pi}}, m)$ query:

- If $\sigma \geq \Sigma$, behave as in $G^{\text{ke}; \bar{\pi}}$, otherwise:
- If ℓ_{ke} has accepted already, simulate the π part honestly with $\ell_{\bar{\pi}}.\text{key}$;
- Compute the response and the state update according to the ke algorithm;
- If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{true}$:
 - If there exists an $\ell'_{\text{ke}; \bar{\pi}} \in \text{LSID}$ such that $\text{Partner}(\ell_{\text{ke}; \bar{\pi}}, \ell'_{\text{ke}; \bar{\pi}}) = \text{true}$ and $\ell'_{\text{ke}}.\text{accept} = \text{true}$, then set $\ell_{\bar{\pi}}.\text{key} \leftarrow \ell'_{\bar{\pi}}.\text{key}$;
 - If there does not exist such an $\ell'_{\text{ke}; \bar{\pi}}$ that is partnered and whose ke part has already accepted, but either $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{ownr}} = \text{corrupt}$ or $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{peer}} = \text{corrupt}$ then set $\ell_{\bar{\pi}}.\text{key} \leftarrow \ell_{\text{ke}}.\text{key}$;
 - If both identities are still honest, and no partner session exists or has already accepted a ke key, then set $\ell_{\bar{\pi}}.\text{key} \xleftarrow{\$} \mathcal{D}$ and update $\sigma \leftarrow \sigma + 1$.

With this new behaviour, we have that the first Σ new keys that are unknown to the adversary at the time of their acceptance are replaced with keys sampled from \mathcal{D} for the π part of the protocol.

Lemma 8.3 now allows us to change the game $G^{\text{ke}; \bar{\pi}, 0, \mathcal{D}}$ into the game $G^{\text{ke}; \bar{\pi}, n, \mathcal{D}}$ for $n = n_i^2 \cdot n_s$ where the indistinguishability of the two games is guaranteed by the BR-secrecy of the ke protocol. This yields

$$\left| \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, 0, \mathcal{D}}}(1^\lambda) - \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, n, \mathcal{D}}}(1^\lambda) \right| \leq n \cdot \text{Adv}_{\mathcal{B}_1, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}}^{\mathcal{D}}}(1^\lambda),$$

for a first reduction \mathcal{B}_1 . In Lemma 8.4, we then show that key confirmation of the composed protocol follows from key confirmation of the symmetric-key protocol:

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, n, \mathcal{D}}}(1^\lambda) = \text{Adv}_{\mathcal{B}_2, \pi, \mathcal{I}}^{G_{\text{SymKeyConf}}}(1^\lambda)$$

for a second reduction \mathcal{B}_2 . This allows us to conclude that

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}}}(1^\lambda) \leq n \cdot \text{Adv}_{\mathcal{B}_1, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}}(1^\lambda) + \text{Adv}_{\mathcal{B}_2, \pi, \mathcal{I}}^{G_{\text{SymKeyConf}}}(1^\lambda).$$

\square

Lemma 8.3. *Let ke be a Match-secure key exchange protocol with output key distribution \mathcal{D} . Let π be a symmetric-key protocol with key generation distribution \mathcal{D} . For $\Sigma = 1, \dots, n_i^2 \cdot n_s$ and for any PPT adversary \mathcal{A} , we have*

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma-1, \mathcal{D}}} (1^\lambda) \leq \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}} (1^\lambda) + \text{Adv}_{\mathcal{B}, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}} (1^\lambda),$$

for some PPT algorithm $\mathcal{B} = \mathcal{B}(\Sigma)$.

Proof. Given an adversary \mathcal{A} against the game $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma-1, \mathcal{D}}$, we construct an algorithm \mathcal{B} against the game $G_{\text{BRSec}, \mathcal{D}}$. The reduction \mathcal{B} sets up the game for \mathcal{A} as described at the beginning of this section and keeps track of the internal variable of each of the stages of the protocol. It also initialises $\sigma \leftarrow 0$.

As \mathcal{A} runs, \mathcal{B} responds to a $\text{Send}(\ell_{\text{ke}; \bar{\pi}}, m)$ query as follows. (We recall that $\ell_{\text{ke}; \bar{\pi}}$ refers here to the variables of $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}$ simulated by \mathcal{B} to \mathcal{A} , ℓ_{ke} refers here to the variables of $G_{\text{BRSec}, \mathcal{D}}$ played by \mathcal{B} and that $\ell_{\bar{\pi}}$ refers to the variables for the execution of $\bar{\pi}$ simulated by \mathcal{B} .)

- If $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \in \{\text{true}, \text{false}\}$, \mathcal{B} returns \perp to \mathcal{A} ; otherwise:
- If $\ell_{\text{ke}}.\text{accept} = \perp$, \mathcal{B} submits $\text{Send}(\ell_{\text{ke}}, m)$ to $G_{\text{BRSec}, \mathcal{D}}$ and receives an updated state for ℓ_{ke} and a response m' . If $\ell_{\text{ke}}.\text{sid}$ is set, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{sid} \leftarrow \ell_{\text{ke}}.\text{sid}$ for \mathcal{A} . If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{false}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{false}$ and informs \mathcal{A} . If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{true}$, the following takes place:
 - If $\sigma = \Sigma$, $\nexists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: \text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}})$ and $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{ownr}} \neq \text{corrupt}$ and $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{peer}} \neq \text{corrupt}$, then
 - * Submit $\text{Test}(\ell)$ to $G_{\text{BRSec}, \mathcal{D}}$ and receive key_{ke} .
 - * Set $\ell_{\bar{\pi}}.\text{key} \leftarrow \text{key}_{\text{ke}}$.
 - * Update $\sigma \leftarrow \sigma + 1$.
 - Else, if $\sigma \leq \Sigma$ then
 - * If there does not exist $\ell'_{\text{ke}} \in \text{LSID}_{\text{ke}}$ such that $\text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}})$ and $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{ownr}} \neq \text{corrupt} \neq \ell_{\text{ke}; \bar{\pi}}.\delta_{\text{peer}}$, then sample a random key $\text{key}_\pi \xleftarrow{\$} \mathcal{D}$ and set $\ell_{\bar{\pi}}.\text{key} \leftarrow \text{key}_\pi$. Update $\sigma \leftarrow \sigma + 1$.
 - * Else, if $\nexists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: \text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}})$ and either $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{ownr}} = \text{corrupt}$ or $\ell_{\text{ke}; \bar{\pi}}.\delta_{\text{peer}} = \text{corrupt}$, then submit the query $\text{Reveal}(\ell_{\text{ke}})$ to $G_{\text{BRSec}, \mathcal{D}}$ and receive key_{ke} . Then set $\ell_{\bar{\pi}}.\text{key} \leftarrow \text{key}_{\text{ke}}$.
 - * Else, there exists an $\ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: \text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}})$ for which $\ell'_{\bar{\pi}}.\text{key}$ has already been set. Then set $\ell_{\bar{\pi}}.\text{key} \leftarrow \ell'_{\bar{\pi}}.\text{key}$.
 - Else $\sigma > \Sigma$ so perform the following:
 - * If $\exists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: \text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}})$ then set $\ell_{\bar{\pi}}.\text{key} \leftarrow \ell'_{\bar{\pi}}.\text{key}$.
 - * Else submit the query $\text{Reveal}(\ell_{\text{ke}})$ to $G_{\text{BRSec}, \mathcal{D}}$, receive key_{ke} and set $\ell_{\bar{\pi}}.\text{key} \leftarrow \text{key}_{\text{ke}}$. Update $\sigma \leftarrow \sigma + 1$.

If \mathcal{A} submits a $\text{Reveal}(\ell_{\text{ke}; \bar{\pi}})$ query, $\ell_{\text{ke}; \bar{\pi}}$ must have accepted for it to be valid. Therefore \mathcal{B} has already manually set the internal key $\ell_{\bar{\pi}}.\text{key}$ and it can return it to \mathcal{A} consistently.

If \mathcal{A} submits a $\text{Corrupt}(i)$ query, \mathcal{B} marks all relevant sessions $\ell_{\text{ke}; \bar{\pi}} \in \text{LSID}_{\text{ke}; \bar{\pi}}$ as **corrupt** if they are still running and then submits $\text{Corrupt}(i)$ to $G_{\text{BRSec}, \mathcal{D}}$ to receive sk_i and return it to \mathcal{A} .

In the processing of a Send query, when $\sigma > \Sigma$ and there is an existing partner session, we initialise the key directly from the partner session's. As we assume that ke is Match-secure, these two partner sessions will derive the same key with overwhelming probability.

We note that if the Test query returns the real key, then \mathcal{B} will perfectly simulate $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma-1, \mathcal{D}}$ to \mathcal{A} , but if it returns a random key from \mathcal{D} , then \mathcal{B} will perfectly simulate $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}$. When \mathcal{A} terminates and was successful, \mathcal{B} submits $\text{Guess}(1)$ to $G_{\text{BRSec}, \mathcal{D}}$; it submits $\text{Guess}(0)$ otherwise. The advantage of \mathcal{B} in $G_{\text{BRSec}, \mathcal{D}}$ therefore corresponds to the difference in the success probability of \mathcal{A} as we have

$$\Pr \left[\text{Exp}_{\mathcal{B}, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}^0} (1^\lambda) = 1 \right] = \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}} (1^\lambda)$$

and

$$\Pr \left[\text{Exp}_{\mathcal{B}, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}^1} (1^\lambda) = 1 \right] = \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma-1, \mathcal{D}}} (1^\lambda)$$

which gives

$$\text{Adv}_{\mathcal{B}, \text{ke}, \mathcal{I}, \mathcal{S}}^{G_{\text{BRSec}, \mathcal{D}}}(1^\lambda) = \left| \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma-1, \mathcal{D}}}(1^\lambda) - \text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, \Sigma, \mathcal{D}}}(1^\lambda) \right|$$

and yields the desired result. \square

Lemma 8.4. *Let ke be a Match-secure key exchange protocol with output key distribution \mathcal{D} and π be a symmetric-key protocol with key generation distribution \mathcal{D} . Let $n = n_i^2 \cdot n_s$. For any PPT adversary \mathcal{A} , it holds that*

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, n, \mathcal{D}}}(1^\lambda) = \text{Adv}_{\mathcal{B}, \pi, \mathcal{I}}^{G_{\text{SymKeyConf}}}(1^\lambda),$$

for some PPT algorithm \mathcal{B} .

Proof. Similarly to the proof of Lemma 8.1, we build a reduction \mathcal{B} against π in $G_{\text{SymKeyConf}}$ which uses an adversary against $\text{ke}; \bar{\pi}$ in $G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, n, \mathcal{D}}$, which we refer to as $G^{\text{ke}; \bar{\pi}, n}$ in this proof for simplicity.

The algorithm \mathcal{B} sets up $G^{\text{ke}; \bar{\pi}, n}$ for \mathcal{A} by simulating all the elements relevant to the ke stage of the composed protocol. It then responds to \mathcal{A} 's queries as follows (we once again use $\ell_{\text{ke}; \bar{\pi}}$ to refer to identifiers used by \mathcal{A} , ℓ_{ke} for corresponding identifiers simulated internally by \mathcal{B} and ℓ_π for those used by \mathcal{B} in $G_{\text{SymKeyConf}}$.

- When \mathcal{A} submits $\text{Send}(\ell_{\text{ke}; \bar{\pi}}, m)$: if $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \in \{\text{true}, \text{false}\}$, \mathcal{B} returns \perp to \mathcal{A} . Otherwise:
 - If $\ell_{\text{ke}}.\text{accept} = \perp$, \mathcal{B} simulates the execution of ke. If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{false}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{false}$. If $\ell_{\text{ke}}.\text{accept} \leftarrow \text{true}$, \mathcal{B} leaves $\ell_{\text{ke}}.\text{key} = \perp$ and then:
 - * If $\ell_{\text{ke}}.\delta_{\text{peer}} = \text{honest} \wedge \nexists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: (\text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}}) \wedge \ell'_{\text{ke}}.\text{accept} = \text{true})$, then \mathcal{B} submits $\text{InitS}(\ell_\pi)$ to $G_{\text{SymKeyConf}}$.
 - * If $\ell_{\text{ke}}.\delta_{\text{peer}} = \text{corrupt} \wedge \nexists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: (\text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}}) \wedge \ell'_{\text{ke}}.\text{accept} = \text{true})$, then \mathcal{B} submits $\text{InitS}(\ell_\pi)$ and then $\text{Reveal}(\ell_\pi)$ to $G_{\text{SymKeyConf}}$ to generate and obtain key_π which it saves by setting $\ell_{\text{ke}}.\text{key} \leftarrow \text{key}_\pi$.
 - * If $\exists \ell'_{\text{ke}} \in \text{LSID}_{\text{ke}} :: (\text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}}) \wedge \ell'_{\text{ke}}.\text{accept} = \text{true})$, then \mathcal{B} submits $\text{InitP}(\ell'_\pi, \ell_\pi)$ to $G_{\text{SymKeyConf}}$ and sets $\ell_{\text{ke}}.\text{key} \leftarrow \ell'_{\text{ke}}.\text{key}$.
 - If $\ell_{\text{ke}}.\text{accept} = \text{true}$, \mathcal{B} submits $\text{Send}(\ell_\pi, m)$ to $G_{\text{SymKeyConf}}$ and returns the reply to \mathcal{A} . If $\ell_\pi.\text{kind} \leftarrow \text{false}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{false}$. If $\ell_\pi.\text{kind} \leftarrow \text{true}$, \mathcal{B} sets $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \leftarrow \text{true}$ and sets $\ell_{\text{ke}; \bar{\pi}}.\text{key} \leftarrow \ell_{\text{ke}}.\text{key}$.
- When \mathcal{A} submits $\text{Reveal}(\ell_{\text{ke}; \bar{\pi}})$: if $\ell_{\text{ke}; \bar{\pi}}.\text{accept} \in \{\perp, \text{false}\}$, \mathcal{B} returns \perp to \mathcal{A} . Otherwise:
 - If $\ell_{\text{ke}; \bar{\pi}}.\text{key} \neq \perp$, \mathcal{B} returns $\ell_{\text{ke}; \bar{\pi}}.\text{key}$ to \mathcal{A} .
 - If $\ell_{\text{ke}; \bar{\pi}}.\text{key} = \perp$, \mathcal{B} submits $\text{Reveal}(\ell_\pi)$ to $G_{\text{SymKeyConf}}$ to obtain key_π , sets $\ell_{\text{ke}; \bar{\pi}}.\text{key} \leftarrow \text{key}_\pi$ and sets $\ell'_{\text{ke}; \bar{\pi}}.\text{key} \leftarrow \text{key}_\pi$ for any $\ell'_{\text{ke}} \in \text{LSID}_{\text{ke}}$ such that $\text{Partner}(\ell_{\text{ke}}, \ell'_{\text{ke}}) \wedge \ell'_{\text{ke}}.\text{accept} = \text{true}$.
- When \mathcal{A} submits $\text{Corrupt}(i)$: \mathcal{B} marks all relevant sessions $\ell_{\text{ke}; \bar{\pi}} \in \text{LSID}_{\text{ke}; \bar{\pi}}$ as **corrupt** (either δ_{ownr} or δ_{peer}) and returns sk_i to \mathcal{A} .

By processing each query as above, the algorithm \mathcal{B} ensures that the first session that accepts the ke stage within a potential partnership pair is mapped to a new session in $G_{\text{SymKeyConf}}$ by an InitS query. If the peer of that session was already corrupt, then \mathcal{B} submits a Reveal query so that this session is flagged as revealed in the game for π . If a session is the second to accept within a partnership pair at the ke stage, then \mathcal{B} uses an InitP query to initialise it with the same key as its partner and to give it the same value for δ_{key} within $G_{\text{SymKeyConf}}$. As we assume that ke is Match-secure, these two partner sessions will derive the same key with overwhelming probability. This, together with \mathcal{B} 's handling of \mathcal{A} 's Reveal and Corrupt queries ensures that every session for which \mathcal{A} could trivially obtain the session key is immediately marked as revealed in $G_{\text{SymKeyConf}}$.

Furthermore, since the key derived by \mathcal{B} 's internal simulation of the ke stage is never used by the π stage, but instead replaced with a randomly generated key using an InitS query, \mathcal{B} provides a perfect simulation of $G^{\text{ke}; \bar{\pi}, n}$ to \mathcal{A} . Therefore, as \mathcal{B} relays \mathcal{A} 's Send queries exactly, we see that if \mathcal{A} wins against the fKeyConf predicate in $G^{\text{ke}; \bar{\pi}, n}$ then \mathcal{B} will also reach a state that wins against $G_{\text{SymKeyConf}}$. We therefore have

$$\text{Adv}_{\mathcal{A}, \text{ke}; \bar{\pi}, \mathcal{I}, \mathcal{S}}^{G_{\text{fKeyConf}}^{\text{ke}; \bar{\pi}, n, \mathcal{D}}}(1^\lambda) = \text{Adv}_{\mathcal{B}, \pi, \mathcal{I}}^{G_{\text{SymKeyConf}}}(1^\lambda).$$

\square

9 Entity authentication

This second form of authentication does not involve the key in its security guarantees. Our *entity authentication* definitions are instead based on the `Partner` predicate, instead of `Samekey`, and say that if two sessions terminate with the same `sid`, then they should agree on each other’s identities. This corresponds to the intuitive notion of entity authentication where sessions obtain guarantees upon terminating and deriving an identifier. We first present our definitions and then demonstrate that for `Match-secure` and `BR-secret` protocols, entity and key authentication are equivalent.

This equivalence also shows that, similarly to including the identities in the key to ensure implicit key authentication, including the identities in the session identifiers ensures implicit entity authentication. It also suggests that involving the `sids` in a MAC is a good method for ensuring entity confirmation.

9.1 Implicit and full explicit entity authentication and confirmation

To adapt our definitions to entity authentication, we replace `Samekey` by `Partner`. Below we present only the predicates for implicit entity authentication, entity confirmation and full explicit entity authentication; the full case requires the definition of the session state variable `econf` $\in \{\text{full}, \text{almost}, \text{no}, \perp\}$, indicating, analogously to `kconf`, which form of authentication a session expects. The almost-full case is more involved because it also requires entity confirmation identifiers, similarly to the `kcid`; for the sake of brevity, we do not include it here. The full definitions can be derived from the following predicates.

Implicit entity authentication: The `iEntAuth` predicate is defined as

$$\forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept}) \implies \forall \ell' \in \text{LSID}, (\text{Partner}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}).$$

Full entity confirmation: The `fexEntConf` predicate is defined as

$$\forall \ell \in \text{LSID}, (\text{aFresh}(\ell) \wedge \ell.\text{econf} = \text{full} \wedge \ell.\text{pid} \in \mathcal{S}) \wedge \ell.\text{accept} \implies \exists \ell' \in \text{LSID} :: \text{Partner}(\ell', \ell).$$

Full explicit entity authentication: The `fexEntAuth` predicate is defined as

$$\begin{aligned} \forall \ell \in \text{LSID}, (\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{econf} = \text{full} \wedge \ell.\text{accept}) \implies \\ (\forall \ell' \in \text{LSID}, \text{Partner}(\ell', \ell) \implies \ell'.\text{id} = \ell.\text{pid}) \\ \wedge (\text{aFresh}(\ell) \implies \exists \ell' \in \text{LSID} :: \text{Partner}(\ell', \ell)). \end{aligned}$$

We note that implicit entity authentication is a very weak notion as the session terminates neither with an explicit guarantee nor with a secret element, such as a key, that it may use later to obtain a stronger guarantee. Separating the two different properties that constitute explicit authentication may be helpful to understand and guide protocol design. For example, with this separation in mind, a signature over the transcript sent at the end of an execution can be seen as providing entity confirmation, and therefore boosting implicit entity authentication to explicit authentication. A similar argument could show that password authentication over a secure channel can serve a similar purpose.

9.2 Key and entity authentication relationships

We now present the necessary conditions for key and entity authentication notions to be equivalent to one another. We use the “secrecy and match-security predicate” `KMSoundness` from Definition 5.3 to state the relationships in terms of predicates. Recall that this predicate holds (with overwhelming probability) for a `Match-secure` and `BR-secret` protocol. The results are intuitively compelling. Under the assumption that partnered sessions derive equal keys and equal keys can only be derived in partnered sessions, authentication guarantees obtained via keys are equivalent to those obtained directly via session identifiers.

Proposition 9.1. *Let π be an AKE protocol. Then it holds for π that*

$$\text{iKeyAuth} \wedge \text{Match} \implies \text{iEntAuth}, \tag{11}$$

$$\text{fexKeyAuth} \wedge \text{Match} \wedge \text{KMSoundness} \implies \text{fexEntAuth}, \tag{12}$$

Proof. The first implication (11) can be seen as follows. Both predicates are identical, except that entity authentication uses the `Partner` predicate instead of `Samekey`. Hence, a mismatch—in the sense that `iKeyAuth` holds but `iEntAuth` does not—can only occur if:

$$\exists \ell \in \text{LSID} :: [((\ell.\text{pid} \in \mathcal{S}) \wedge \ell.\text{accept}) \wedge \exists \ell' \in \text{LSID} :: (\neg \text{Samekey}(\ell', \ell) \wedge \text{Partner}(\ell', \ell))].$$

If the sessions ℓ, ℓ' would have the same key, and then they would also satisfy the identity requirement $\ell'.\text{id} = \ell.\text{pid}$, because of the `iKeyAuth` property.

Note that the partnering predicate stipulates that the session identifiers of ℓ and ℓ' are equal (and different from \perp). According to our specification of key exchange protocols this, in turn, implies that both sessions must have accepted and, if so, that they have set the keys to values different from \perp . But now we would get an immediate contradiction to Property (1) of `Match` security:

$$\exists \ell, \ell' \in \text{LSID} :: \text{Partner}(\ell, \ell') \wedge (\ell.\text{key} \neq \perp \neq \ell'.\text{key}) \wedge \neg \text{Samekey}(\ell, \ell').$$

The second implication (12) follows similarly, there are two possibilities for a mismatch (`fexKeyAuth` holds, but `fexEntAuth` does not). Either the first property in the implication in `fexEntAuth`, which also appears in the implicit definition, is false, in which case we get the same contradiction as before. Or the second property in the implication (`aFresh` $\implies \exists \ell' :: \text{Partner}(\ell', \ell)$) is false, although it holds in `fexKeyAuth` for the `Samekey` case. This means that

$$\begin{aligned} \exists \ell \in \text{LSID} :: & \left(\text{aFresh}(\ell) \wedge \ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept} \right. \\ & \left. \wedge [\exists \ell' \in \text{LSID} :: \text{Samekey}(\ell', \ell)] \wedge [\forall \ell'' \in \text{LSID} :: \neg \text{Partner}(\ell'', \ell)] \right). \end{aligned}$$

Note that this means that there will be a session ℓ' which has the same key as ℓ , and since ℓ has accepted it must be a valid key $\ell.\text{key} \neq \perp$, but such that no other session is partnered with ℓ . This, however, contradicts the `KMSoundness` predicate. \square

Proposition 9.2. *Let π be an AKE protocol. Then it holds for π that*

$$\text{iKeyAuth} \iff \text{iEntAuth} \wedge \text{Match} \wedge \text{KMSoundness}, \quad (13)$$

$$\text{fexKeyAuth} \iff \text{fexEntAuth} \wedge \text{Match} \wedge \text{KMSoundness} \quad (14)$$

Proof. We start with the first implication (13). Similar to the previous proposition one can show that any mismatch in the predicates implies

$$\exists \ell \in \text{LSID} :: [\ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept} \wedge \exists \ell' \in \text{LSID} :: (\text{Samekey}(\ell', \ell) \wedge \neg \text{Partner}(\ell', \ell))].$$

This, once more, would contradict the `KMSoundness` predicate.

The second implication (14) derives a contradiction as in the implicit case, or we can analogously to the other direction conclude that

$$\begin{aligned} \exists \ell \in \text{LSID} :: & \left(\text{aFresh}(\ell) \wedge \ell.\text{pid} \in \mathcal{S} \wedge \ell.\text{accept} \right. \\ & \left. \wedge [\exists \ell' \in \text{LSID} :: \text{Partner}(\ell', \ell)] \wedge [\forall \ell'' \in \text{LSID} :: \neg \text{Samekey}(\ell'', \ell)] \right). \end{aligned}$$

This, of course, would contradict Property (1) of the `Match` predicate, since we would have partnered sessions ℓ, ℓ' which do not hold the same (valid) key $\ell.\text{key} \neq \perp$. Here we use the fact that partnering implies non-trivial session identifiers, and if session ℓ' has set the identifier, it has accepted and set a key $\ell'.\text{key} \neq \perp$, too. \square

Finally, we show that implicit entity authentication together with (full) explicit entity confirmation is equivalent to full explicit entity authentication.

Proposition 9.3. *Let π be an AKE protocol. Then it holds for π that*

$$\text{iEntAuth} \wedge \text{fexEntConf} \iff \text{fexEntAuth}$$

The proof follows as in the case of key authentication.

References

- [BFWW11] Christina Brzuska, Marc Fischlin, Bogdan Warinschi, and Stephen C. Williams. Composability of Bellare-Rogaway key exchange protocols. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 11*, pages 51–62. ACM Press, October 2011.
- [BJM97] Simon Blake-Wilson, Don Johnson, and Alfred Menezes. Key agreement protocols and their security analysis. In *Cryptography and Coding, 6th IMA International Conference, 1997, Proceedings*, volume 1355 of *LNCS*, pages 30–45. Springer, 1997.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 139–155. Springer, Heidelberg, May 2000.
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, August 1994.
- [BR95] Mihir Bellare and Phillip Rogaway. Provably secure session key distribution: The three party case. In *27th ACM STOC*, pages 57–66. ACM Press, May / June 1995.
- [BSWW13] Christina Brzuska, Nigel P. Smart, Bogdan Warinschi, and Gaven J. Watson. An analysis of the EMV channel establishment protocol. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 373–386. ACM Press, November 2013.
- [BWM99] Simon Blake-Wilson and Alfred Menezes. Unknown key-share attacks on the station-to-station (STS) protocol. In Hideki Imai and Yuliang Zheng, editors, *PKC'99*, volume 1560 of *LNCS*, pages 154–170. Springer, Heidelberg, March 1999.
- [CBH05a] Kim-Kwang Raymond Choo, Colin Boyd, and Yvonne Hitchcock. Examining indistinguishability-based proof models for key establishment protocols. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 585–604. Springer, Heidelberg, December 2005.
- [CBH05b] Kim-Kwang Raymond Choo, Colin Boyd, and Yvonne Hitchcock. On session key construction in provably-secure key establishment protocols. In *Progress in Cryptology - Mycrypt 2005, First International Conference on Cryptology in Malaysia, Kuala Lumpur, Malaysia, September 28-30, 2005, Proceedings*, volume 3715 of *Lecture Notes in Computer Science*, pages 116–131. Springer, 2005.
- [CF12] Cas J. F. Cremers and Michele Feltz. Beyond eCK: Perfect forward secrecy under actor compromise and ephemeral-key reveal. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *ESORICS 2012*, volume 7459 of *LNCS*, pages 734–751. Springer, Heidelberg, September 2012.
- [CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001.
- [Cre11] Cas Cremers. Examining indistinguishability-based security models for key exchange protocols: The case of CK, CK-HMQV, and eCK. ASIACCS '11, pages 80–91, New York, NY, USA, 2011. ACM.
- [DFGS15] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 15*, pages 1197–1210. ACM Press, October 2015.
- [DvOW92] Whitfield Diffie, Paul C. van Oorschot, and Michael J. Wiener. Authentication and authenticated key exchanges. *Des. Codes Cryptography*, 2(2):107–125, 1992.
- [FGSW16] Marc Fischlin, Felix Günther, Benedikt Schmidt, and Bogdan Warinschi. Key confirmation in key exchange: A formal treatment and implications for TLS 1.3. In *2016 IEEE Symposium on Security and Privacy*, pages 452–469. IEEE Computer Society Press, May 2016.
- [Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 546–566. Springer, Heidelberg, August 2005.
- [LLM07] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Heidelberg, November 2007.
- [LS17] Yong Li and Sven Schäge. No-match attacks and robust partnering definitions: Defining trivial attacks for security protocols is not trivial. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 1343–1360. ACM Press, October / November 2017.
- [MTI86] Tsutomu Matsumoto, Youichi Takashima, and Hideki Imai. On seeking smart public-key-distribution systems. *Transactions of the Institute of Electronics and Communication Engineers of Japan. Section E*, E69(2):99–106, 2 1986.

- [MU08] Alfred Menezes and Berkant Ustaoglu. Comparing the pre- and post-specified peer models for key agreement. In Yi Mu, Willy Susilo, and Jennifer Seberry, editors, *ACISP 08*, volume 5107 of *LNCS*, pages 53–68. Springer, Heidelberg, July 2008.
- [MvOV97] Alfred Menezes, Paul van Oorschot, and Scott Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1997.
- [RS09] Phillip Rogaway and Till Stegers. Authentication without elision: Partially specified protocols, associated data, and cryptographic models described by code. In *CSF*, pages 26–39. IEEE Computer Society, 2009.
- [Ust09] Berkant Ustaoglu. Comparing sessionstatereveal and ephemeralkeyreveal for Diffie-Hellman protocols. In Josef Pieprzyk and Fangguo Zhang, editors, *ProvSec 2009*, volume 5848 of *LNCS*, pages 183–197. Springer, Heidelberg, November 2009.
- [Yan13] Zheng Yang. Modelling simultaneous mutual authentication for authenticated key exchange. In *Foundations and Practice of Security - 6th International Symposium, FPS 2013, La Rochelle, France, October 21-22, 2013, Revised Selected Papers*, volume 8352 of *Lecture Notes in Computer Science*, pages 46–62. Springer, 2013.