

ASICS: Authenticated Key Exchange Security Incorporating Certification Systems

Colin Boyd¹, Cas Cremers², Michèle Feltz², Kenneth G. Paterson³, Bertram Poettering³, and Douglas Stebila¹

¹ Queensland University of Technology, Brisbane, Australia

² Institute of Information Security, ETH Zurich, Switzerland

³ Royal Holloway, University of London, Egham, United Kingdom
{c.boyd, stebila}@qut.edu.au, {cas.cremers, feltzm}@inf.ethz.ch,
{kenny.paterson, bertram.poettering}@rhul.ac.uk

Abstract. Most security models for authenticated key exchange (AKE) do not explicitly model the associated *certification system*, which includes the certification authority (CA) and its behaviour. However, there are several well-known and realistic attacks on AKE protocols which exploit various forms of malicious key registration and which therefore lie outside the scope of these models. We provide the first systematic analysis of *AKE security incorporating certification systems* (ASICS). We define a family of security models that, in addition to allowing different sets of standard AKE adversary queries, also permit the adversary to register arbitrary bitstrings as keys. For this model family we prove generic results that enable the design and verification of protocols that achieve security even if some keys have been produced maliciously. Our approach is applicable to a wide range of models and protocols; as a concrete illustration of its power, we apply it to the CMQV protocol in the natural strengthening of the eCK model to the ASICS setting.

Keywords: authenticated key exchange (AKE), unknown key share (UKS) attacks, certification authority (CA), invalid public keys, PKI

1 Introduction

After public key encryption and digital signatures, authenticated key establishment (AKE) is perhaps the most important public key primitive. From a real-world perspective, AKE protocols relying on public key techniques are widely deployed in systems that are used every day by billions of users, including systems such as TLS, IPsec, SSH, and various single sign-on systems. From a theoretical perspective, formal, cryptographically sound modelling for AKE protocols began in the symmetric setting with the seminal work of Bellare and Rogaway [5], and was later extended to the public key setting [7]. Since then, there has been a large body of work in this tradition, and many additions and modifications have been proposed. The most prominent current models in this tradition [4, 13, 29, 41] strengthen or add to the required security properties, cover different protocol classes, and strengthen adversary powers.

Despite intensive study over two decades, important elements of AKE protocols have not been sufficiently modelled, preventing our deeper understanding of this important primitive and limiting its applicability to real-world protocols. Specifically, the public key infrastructure (PKI) needed to support the authenticity of public keys in AKE, and the interactions between the certification authority (CA), honest parties, and the adversary, are rarely modelled. Rather, with exceptions as noted below, in typical AKE models and proofs it is assumed that all public keys are honestly generated and authentically distributed at the start of the security game, and that there is a single key per party; certificates are excluded from the model. The adversary can corrupt parties, learning all their secrets, but has limited ability to register malicious keys. Roughly speaking, this modelling approach corresponds to an ideal CA, who zealously generates perfect key pairs and securely distributes them to the correct parties.

However, CAs in the real world simply do not operate in such rigorous ways. They have differing strengths of procedures for checking claimed identities⁴, so malicious parties might in some cases get arbitrary public keys certified against identifiers of their choice. The most egregious examples involve CAs who, either willingly, under coercion, or as a result of security compromises, have issued certificates for keys and identifiers that they should not have.⁵ CAs following best-practices may require that a user

⁴ For example, issuance of Extended Validation (EV) certificates requires stronger identity-checking requirements than non-EV certificates, see <https://www.cabforum.org/certificates.html> for more details.

⁵ In June and July 2011, Dutch CA DigiNotar was hacked [20], with the intruder taking control of all 8 of the CA's signing servers; at least 531 rogue certificates were then issued. In August 2011, TURKTRUST CA [18] issued

requesting a certificate submit a certificate signing request (CSR, as per [43]) to the CA. This involves the user self-signing the data that is to be certified. Other approaches to providing *proofs of possession* are standardised in [1, 2, 38]. However, even these basic tests of private key ownership are not mentioned in industry guidelines issued by the CA/Browser Forum [11, 12]. Furthermore, these procedures all fall short of the proofs of knowledge [37] required to match what is assumed in typical AKE models. Thus, an attacker may be able to register another party’s public key under his own identifier, or register a malformed key which then interacts with properly generated keys in an unfortunate way.

Critically, there are realistic attacks on AKE protocols which cannot be captured by AKE security models that omit CA and PKI aspects:

- Kaliski’s *unknown key share* (UKS) attack [25] on early versions of MQV exploits the ability of the adversary to dynamically register a public key (which is valid and for which the adversary *does* know the secret key).
- The UKS attack on KEA described by Lauter and Mityagin [30, p. 380] exploits the fact that the adversary is able to re-register some party’s static public key as his own public key.
- Blake-Wilson and Menezes [9] introduced the *duplicate-signature key selection* (DSKS) attack on signature schemes: after observing a user’s signature σ on a message m , the adversary E is able to compute a signature key pair (sk_E, vk_E) (or sometimes just a verification key vk_E) such that σ is also E ’s signature on the message m . Now, for example, if the Station-to-Station (STS) protocol is implemented using a signature scheme that is vulnerable to DSKS attacks, and the adversary can register arbitrary public keys with the CA, then the protocol is vulnerable to an online UKS attack [9].
- In Lim and Lee small subgroup attacks [31], the adversary extracts partial (or even complete) information about a party’s long-term secret key. Some of these attacks require registering invalid public keys with the CA before engaging in protocol runs with honest participants. Of particular note are the Lim–Lee-style attacks of Menezes and Ustaoglu [34] on the HMQV protocol [27].

We claim that to date there has been no *systematic* treatment in the literature of the behaviour of CAs with respect to public keys and identifiers chosen by the adversary and, relatedly, the treatment of public keys and ephemeral values by protocol participants (especially with respect to group membership testing). Our paper sets out to rectify this situation, providing a comprehensive and self-contained treatment of these features, as well as establishing generic results to make protocols resilient against such attacks.

Contributions. Our paper has three main contributions.

First, we present in Section 2 a framework for reasoning about the security of AKE protocols with respect to various CA key registration procedures. This framework allows us to capture several attacks based on adversarial key registration, including UKS attacks, small-subgroup attacks, attacks that occur when the CA does not check if public keys are registered twice, and attacks that occur when multiple public keys can be registered per identifier.

Second, we provide in Section 3 a generic approach to achieve strong security guarantees against adversaries that can register arbitrary public keys for certain types of protocols. In particular, we show how to transform Diffie–Hellman type AKE protocols that are secure in a model where only honest key registration is allowed into protocols that are secure even when adversaries can register arbitrary valid or invalid public keys. In such cases, security is still guaranteed for all sessions (that were considered *clean* or *fresh* in the base model) except those in which the peer’s public key is valid but registered by the adversary.

Third, we demonstrate in Section 4 how our methodology can be used to establish strong security guarantees, even when the adversary can register arbitrary public keys, for concrete protocols such as CMQV, NAXOS, and UP, using CMQV as a running example. We provide in Section 5 recommendations for the design of protocols that are secure in our models.

Related work. The original computational model for key exchange of Bellare and Rogaway [5] has a long-lived key generator, which is used to initialise all parties’ keys at the start of the game. This is a standard part of most computational models today. However, in common with several later models [13, 24, 28], the adversary cannot influence long-term keys: only honestly generated keys are considered. Starting with the 1995 model of Bellare and Rogaway [6] it was recognised that the adversary may be able to choose long-term keys for certain parties, whether public keys or symmetric keys. It is possible to identify three different methods that have been used to model such an adversary capability.

special certificates with wildcard signing capabilities, allowing impersonation of any domain in the Internet. This was discovered, by coincidence, only 18 months later.

1. The adversary can replace long-term keys by providing them as an input to a `corrupt` query. This was the method used originally by Bellare and Rogaway [6] and was subsequently used in the public key setting by others [8, 35].
2. The adversary is allowed to generate arbitrary keys for corrupted parties at any time during the protocol run [27].
3. An additional query is added specifically to set up a user with a new key chosen by the adversary [15, 22, 44]. This query is typically called `establishparty` and takes as input the user name and its long-term public key.

These methods allow the models to capture the Kaliski attack [25], which requires the adversary to register a new public key after certain protocol messages have been obtained. However, none of these currently used methods has the generality of our model and, in particular, all of them omit the following realistic features:

- registration of multiple public keys per user;
- flexible checking by certification authorities via a verification procedure;
- adversarial choice of public keys per session.

Special mention should also be made of the model of Shoup [41]. Unlike most popular AKE models today, it uses a simulatability definition of security comparing ideal and real world views. Security is defined to mean that for any real world adversary there is an ideal world adversary (benign by definition) such that the transcripts of the two are computationally indistinguishable. Real-world adversaries have the ability to assign users to public key certificates. Shoup’s model has not been widely used and the examples in [41] are not fully worked through. Furthermore, the model cannot represent an adversary who obtains only ephemeral secret keys without knowing the long-term key of the same user and therefore cannot capture security properties common in more modern models.

Cash et al. [14] and Hofheinz et al. [21] have considered the security of non-interactive key exchange (NIKE) in settings where the adversary can register arbitrary public keys, analogously to our ASICS setting for interactive key exchange. It is an interesting open problem to examine how the security models and constructions for NIKE in [14, 21] can be built upon to achieve security in the ASICS setting.

Boldyreva et al. [10] model certification and PKI for public key encryption and signatures. They point to legitimate concerns regarding what certification authorities can be expected to check in practice. It is interesting to note that while they use key exchange as a prime motivating example for modelling certification and there has been some follow-up work [19], key exchange is not explored in their models. Nevertheless, [10, 19] are in the spirit of our work: like ours, their models allow (amongst other things) corrupt key registration. However, they do not allow the adversary to reveal long-term or ephemeral keys, which is fundamental to modern AKE security models.

Critically, all of the approaches mentioned above have only been used to establish results for a handful of specific protocols. In contrast, we establish *generic results* that facilitate the design and verification of AKE protocols, and that can be applied to a large class of protocols.

2 ASICS model family

In this section we define a parameterized AKE security model that allows for explicit modelling of the certification of public keys. Prominent AKE security frameworks can be instantiated in this family of models, as well as extensions that allow dynamic adversarial registration of arbitrary bitstrings as public keys.

Generally speaking, from a user’s point of view, participation in key exchange encompasses three consecutive phases: First, users set up their individual key pairs; more precisely, each user invokes a randomized algorithm `KeyGen` that outputs a fresh secret-key/public-key pair (sk, pk) . Second, users contact a certification authority (CA) to get their keys certified: each user provides the CA with its identifier \hat{P} and its public key pk , and obtains a certificate C that binds the identifier to the key. After completing these setup steps, in the third phase, users can engage in interactive sessions with other users to establish shared keys. To do so, they usually require knowledge of their own key pair (sk, pk) , their identifier \hat{P} , and the corresponding certificate C . In addition to that, protocols may require a priori knowledge of (a subset of) the peer’s public key pk' , peer’s identifier \hat{Q} , and peer’s certificate C' . As we will see, our execution model is general enough to cover all these settings.

To ease notation, we assume that public key pk and identifier \hat{P} can be readily derived from any certificate C ; we use notation $C.\text{pk} = \text{pk}$ and $C.\text{id} = \hat{P}$ correspondingly. This assumption holds for all practical PKIs we are aware of, and in particular for X.509 certificates (as used in SSL/TLS); the latter

carry public key and identifier in a canonically formatted data structure. In potential other cases, where keys and identifiers cannot be extracted from certificates but instead are auxiliary input to the verification routine, certificates can still be expected to be valid for only a single key/identifier pair, i.e., the mapping from C to compatible pk and \hat{P} is again unambiguous. Proposed notations $C.\text{pk}$ and $C.\text{id}$ are hence meaningful also in such cases.

In this paper, we are particularly interested in the interaction between users and the CA that occurs in the certification process. We consider different degrees of rigour in the checks of consistency and ownership of public keys pk presented to the CA. On the one hand, CAs could be pedantic with such verifications (e.g., require a proof of knowledge⁶ of the secret key corresponding to pk); on the other hand, CAs could also just accept any given bitstring pk as valid and issue a certificate on it. The ability to precisely assess the security of key establishment in the face of different CA behaviours is a key contribution of our new model family.

Definition 1. An ASICS protocol Π consists of a set of domain parameters, a key generation algorithm KeyGen , a public key verification procedure VP , and the protocol description π that describes how key exchange protocol messages are generated and responded to as well as how the session key is derived.

We denote by VP the specific *verification procedure* on public keys and identifiers that a considered CA deploys. As different checks on pk and \hat{P} might require different levels of interaction between the registering user and the CA, we model it as a procedure, as opposed to a function. We require that VP is efficient and has binary output. Furthermore, we require that the CA issues the requested certificate only if VP outputs value 1; all certification requests where VP outputs value 0 are rejected. Note that, for simplicity, we only consider non-interactive verification procedures (i.e., two-message registration protocols) between the user and the CA. A more general treatment covering interactive verification procedures as well would introduce additional complexities to our framework.

Specific key exchange protocols might be insecure for one (liberal) instantiation of VP , and be secure for another (stricter) one. Note that CAs that do not perform any check on pk and \hat{P} are modelled by a verification procedure VP that always outputs 1. A verification procedure that performs few checks may output 1 for at least all $\text{pk} \in \mathbf{PK}$, where \mathbf{PK} denotes the set of possible public keys output by KeyGen . Precisely, if the inputs of algorithm KeyGen are security parameter 1^k and randomness $r \in_R \{0, 1\}^k$, then we define

$$\mathbf{PK} = \{\text{pk} \mid \text{there exists } r \in \{0, 1\}^k \text{ such that } \text{KeyGen}(1^k; r) = (\cdot, \text{pk})\} \quad .$$

A verification procedure with high assurance may require a zero-knowledge argument that the requester knows the secret key corresponding to the public key, and even that the key was generated verifiable at random. Note that we allow VP to keep an internal state between invocations; our model hence covers possible implementations of CAs that reject certification requests with public keys that have already been registered (e.g., for a different identifier).

2.1 Security model

At a high level, our model stipulates users that generate one or more keys, obtain certificates for these keys from a CA, and use keys and certificates to run (potentially concurrent) sessions of the key agreement protocol. Similar to other security models [5, 13], the adversary controls all communication in these sessions, corrupts users at will to obtain their secret keys, and arbitrarily reveals established session keys. Innovative is the adversary’s additional ability to steer the registration process with the CA: it can obtain from the CA valid certificates for public keys and identifiers of its choosing (as long as VP evaluates to 1), and provides users with such certificates.

To keep our model simple and comprehensible, we abstract away any forgeability issues of certificates and assume the following ideal functionality: no certificate will be considered valid unless it has been issued by the CA. We model this by letting the challenger keep a list \mathcal{C} of all CA-issued certificates and by equipping users with a *certificate verification oracle* \mathcal{O}_{CV} that checks membership in that list; concretely, we assume that $\mathcal{O}_{CV}(C) = 1 \Leftrightarrow C \in \mathcal{C}$. Of course, in concrete implementations, this oracle is replaced by an explicit local verification routine; for instance, if certification is implemented via a signature scheme, this will include its verification procedure (besides further checking of validity, chain-validity, revocation state, etc.).

⁶ Although interactive or non-interactive zero-knowledge proof systems seem to yield ideal solutions in this context, standards like X.509, OpenPGP, and PKCS#10 content themselves with the purely heuristic (and inferior) approach of demanding a so-called *proof-of-possession* (PoP), mostly implemented via self-signed certificates or self-signed certificate requests (cf. [1, 38, 39]). The security of such constructions seems to be difficult to formally assess [37].

Sessions and session state. Users, once they have created their keys and obtained corresponding certificates, can execute protocol sessions. Within a user, each such session is uniquely identified by a pair $s = (C, i)$, where C denotes the certificate used by the user (by himself) in that session, and i is a counter. The user maintains session-specific variables as indicated in Table 1. Some session variables are fixed upon session creation, whereas others can be assigned or updated during protocol execution. Some, such as `pcert`, `status`, and `key`, are considered to be outputs of the key establishment and might be used in higher-level protocols or applications. A session s has *accepted* if $s_{\text{status}} = \text{accepted}$.

<code>acert</code>	certificate of the actor (the user running this session)
<code>pcert</code>	certificate of this session's peer
<code>role</code>	taken role; either \mathcal{I} (initiator) or \mathcal{R} (responder)
<code>sent</code>	concatenation of all messages sent in this session
<code>rcvd</code>	concatenation of all messages received in this session
<code>status</code>	session status; either active , accepted , or rejected
<code>key</code>	key in $\{0, 1\}^k$ established in this session
<code>rand</code>	randomness used in this session
<code>data</code>	any additional protocol-specific data

Table 1. Elements of session state

$\mathcal{Q}_N = \{\text{kgen}, \text{hregister}, \text{create}, \text{send}\}$	(Normal protocol behaviour)
$\mathcal{Q}_S = \{\text{corrupt}, \text{randomness}, \text{session-key}\}$	(corruption of Secrets)
$\mathcal{Q}_R = \{\text{pkregister}, \text{npkregister}\}$	(adversarial key Registration)

Table 2. Overview of query sets. Additionally, there is a `test-session` query.

Adversarial queries. The adversary interacts with users by issuing queries. The adversary can direct users to establish long-term key pairs and certificates (`kgen`, `hregister`), to initiate protocol sessions (`create`), and to respond to protocol messages (`send`). The adversary may be able to learn long-term keys (`corrupt`), session-specific randomness (`randomness`), or session keys (`session-key`) from users. The adversary can also maliciously obtain certificates from the CA (`pkregister`, `npkregister`).

The queries in set $\mathcal{Q}_N = \{\text{kgen}, \text{hregister}, \text{create}, \text{send}\}$, defined as follows, model normal operation of the protocol; they are required in any security model. Initially, the auxiliary variables \mathcal{HK} , \mathcal{C} , \mathcal{C}_h , \mathcal{C}_{pk} , and \mathcal{C}_{npk} are set to \emptyset .

- `kgen()` By running algorithm `KeyGen`, a fresh key pair (sk, pk) is generated. Public key pk is returned to the adversary; secret key sk is stored for processing potential later queries corresponding to pk . The public key is added to the set of honestly generated keys: $\mathcal{HK} \leftarrow \mathcal{HK} \cup \{\text{pk}\}$.
- `hregister(pk, \hat{P})` The query requires that $\text{pk} \in \mathcal{HK}$ and that VP outputs 1 on input pk^7 and \hat{P} ; otherwise, it returns \perp . The public key pk is registered at the CA for the identifier \hat{P} . The resulting certificate C is added to the global set of certificates and to the set of honestly generated certificates: $\mathcal{C} \leftarrow \mathcal{C} \cup \{C\}$ and $\mathcal{C}_h \leftarrow \mathcal{C}_h \cup \{C\}$. The query returns C .
- `create($s = (C, i), r, [C']$)` The query requires that $C \in \mathcal{C}_h$, that a session with counter i for certificate C does not already exist, and that $r \in \{\mathcal{I}, \mathcal{R}\}$; otherwise, it returns \perp . A new session s is created for the user with public key $C.\text{pk}$ and identifier $C.\text{id}$. Session variables are initialized as

$$(s_{\text{acert}}, s_{\text{pcert}}, s_{\text{role}}, s_{\text{sent}}, s_{\text{rcvd}}, s_{\text{status}}, s_{\text{key}}) \leftarrow (C, \perp, r, \epsilon, \epsilon, \text{active}, \perp) .$$

If the optional certificate C' is provided, we set $s_{\text{pcert}} \leftarrow C'$. In addition, a string in $\{0, 1\}^k$ is sampled uniformly at random and assigned to s_{rand} ; we assume that all randomness required during the execution of session s is deterministically derived from s_{rand} . The user also runs the initialization procedure for the key exchange protocol, which may further initialize its own (internal) state variable s_{data} and optionally generate a message M . If M was generated, set $s_{\text{sent}} \leftarrow M$, and return M . Otherwise, return \perp .

⁷ Reasonable implementations of `VP` output 1 on all keys $\text{pk} \in \mathcal{HK}$, because $\mathcal{HK} \subseteq \mathbf{PK}$.

- **send**(s, M) The query requires that session s exists and that $s_{\text{status}} = \text{active}$; otherwise, it returns \perp . The user continues the protocol execution for this session with incoming message M , which may optionally generate a response message M' . Next, s_{rcvd} is set to $(s_{\text{rcvd}} \parallel M)$ and, if M' is output, s_{sent} is set to $(s_{\text{sent}} \parallel M')$. The protocol execution may (re-)assign values to s_{status} and s_{key} , and to the session's internal state variable s_{data} . Also, if the value s_{pcert} was not provided to the **create** query, then protocol execution may assign a value to s_{pcert} . (For example, in TLS the client does not learn the server's certificate until it receives the **ServerCertificate** message.) If M' was generated, return M' ; otherwise return \perp .

Remark 1 (Multiple CAs). Our model stipulates a PKI with a single CA. This could be considered a limitation for the analysis of many practically relevant scenarios. For example, the PKI currently in place to secure web traffic in the Internet consists of a large number of independently acting CAs. However, our single-CA setting can readily emulate a multi-CA setting: If $\text{VP}_1, \dots, \text{VP}_n$ denote the verification procedures of n independent CAs, then composed verification procedure $\text{VP}_{1\dots n}$ that outputs 1 if there exists at least one i , $1 \leq i \leq n$, such that VP_i outputs 1, effectively folds the n CAs into a single one, adequately modeling the aforementioned multi-CA setting.

Remark 2 (One key for many identifiers). A frequently observed setting in the web PKI is that users hold one public key that they bind to a set of identifiers (with the corresponding number of certificates). This happens, for instance, if globally acting companies register domain names in several different countries, but process HTTPS requests on a single centralized facility. Observe that our model covers such cases of multiple registration of a key: **hregister** may be queried many times for the same **pk**. Moreover, through being identified by pairs (C, i) (as opposed to (pk, i)), sessions associated with one **pk** are aware of the identifier $C.\text{id}$ in use.

The queries in set $\mathcal{Q}_S = \{\text{corrupt}, \text{randomness}, \text{session-key}\}$ model the corruption of a user's secrets. Similar queries are found in other standard AKE models [5, 13].

- **corrupt**(**pk**) The query requires $\text{pk} \in \mathcal{HK}$; otherwise, it returns \perp . This query returns the secret key **sk** corresponding to public key **pk**.
- **randomness**(s) The query requires that session s exist; otherwise, it returns \perp . The query returns the randomness s_{rand} .
- **session-key**(s) The query requires that session s exist and that $s_{\text{status}} = \text{accepted}$; otherwise, it returns \perp . The query returns the session key s_{key} .

Remark 3 (Semantics of randomness query). The **randomness** query gives the adversary access to the randomness s_{rand} that is used in a particular session. Formally, one may think of the session as an instance of a probabilistic Turing machine: the query gives the adversary the values that are read from the random tape. Practically, this models a randomness generator whose values are revealed after they are produced, for example, by a side-channel attack. This is similar to the **ephemeral key reveal** query from [29].

The **hregister** query introduced above only allows registration of keys $\text{pk} \in \mathcal{HK}$, i.e., keys held by honest users. In contrast, the adversary can obtain certificates on *arbitrary* (valid) public keys using the following **pkregister** query. Going even further, the **npkregister** query allows registration of objects that are not even public keys (always assuming that **VP** outputs 1 on the candidate object). These queries will allow modelling Kaliski's attack on MQV [25] and small subgroup attacks [31], amongst others. We emphasize that the queries in set $\mathcal{Q}_R = \{\text{pkregister}, \text{npkregister}\}$ have no counterparts in standard definitions of key exchange security.

- **pkregister**(pk, \hat{P}) The query requires that $\text{pk} \in \mathbf{PK}$ and that **VP** outputs 1 on input pk and \hat{P} ; otherwise, it returns \perp . The public key pk is registered at the CA for identifier \hat{P} . The resulting certificate C is added to the global set of certificates and to the set of certificates generated through **pkregister** query: $\mathcal{C} \leftarrow \mathcal{C} \cup \{C\}$ and $\mathcal{C}_{\text{pk}} \leftarrow \mathcal{C}_{\text{pk}} \cup \{C\}$. The query returns C .
- **npkregister**(pk, \hat{P}) The query requires that $\text{pk} \notin \mathbf{PK}$ and that **VP** outputs 1 on input pk and \hat{P} ; otherwise, it returns \perp . The public key pk is registered at the CA for the identifier \hat{P} . The resulting certificate C is added to the global set of certificates and to the set of certificates generated through **npkregister** query: $\mathcal{C} \leftarrow \mathcal{C} \cup \{C\}$ and $\mathcal{C}_{\text{npk}} \leftarrow \mathcal{C}_{\text{npk}} \cup \{C\}$. The query returns C .

Remark 4 (Efficiency of pkregister and npkregister processing). Observe that **pkregister** and **npkregister** queries require a membership test for set \mathbf{PK} . In some settings, such tests might correspond to computa-

tionally hard problems like quadratic residuosity or DDH⁸. Although we stress that the named queries and the ensuing security experiments are well-defined nevertheless, we discuss two possible problems that might arise when analyzing key exchange protocols where **PK**-membership is difficult to assess: (1) If the queries, by whatever means, correctly decide membership in **PK**, then the adversary implicitly gets access to a decision oracle for that set; this has to be taken into account in corresponding security reductions, e.g., when selecting appropriate hardness assumptions. (2) If an analyzed ASICS protocol Π is used as a component in a higher-level construction Π' , then any security argument that reduces the security of Π' to the security of Π will inevitably have to specify how `pkregister` and `npkregister` oracles are to be simulated. As mentioned before, in some settings this seems to be infeasible.

2.2 Security experiment

Using the above queries, we define a parameterized family of AKE security models. As is common in BR-style AKE models, we must restrict query usage so that the adversary cannot trivially win the security experiment. The conditions under which queries are disallowed are expressed by a *freshness* condition, which typically uses a *matching* condition to formalize intended partner sessions.

Definition 2 (Matching, freshness, ASICS model). *Let Π be an ASICS protocol. A matching condition M for Π is an efficient binary relation on the set of sessions of Π . Let Q be a set of queries such that $\mathcal{Q}_N \subseteq Q \subseteq \mathcal{Q}_N \cup \mathcal{Q}_S \cup \mathcal{Q}_R$. A freshness condition F for (Π, Q) is an efficient predicate (usually depending on a matching condition M) that takes a session of Π and a sequence of queries (including arguments and results) of a security experiment over queries in Q . We call $X = (M, Q, F)$ an ASICS model for Π .*

Definition 3 gives two possible matching conditions. We will later give examples of freshness conditions, in Example 1 on the following page and in Section 4.1.

We refer to [16] for more details on the intricacies of matching definitions in AKE protocols. Two issues are important here. First, there is a strong connection between the information used in a matching definition and the information used to compute the session key. Second, some protocols like the two-message versions of MQV and HMQV allow sessions to compute the same key even if they perform the same role, whereas other protocols such as NAXOS require the sessions that compute the same key to perform different roles. In the remainder of the paper we will use one of the definitions below, depending on the type of protocol.

Definition 3 (M1-matching, M2-matching). *Let s and s' denote two sessions of an ASICS protocol. We say that session s' M1-matches (or is M1-matching) session s if $s_{\text{status}} = s'_{\text{status}} = \text{accepted}$ and*

$$\begin{aligned} & (s_{\text{acert.pk}}, s_{\text{acert.id}}, s_{\text{pcert.pk}}, s_{\text{pcert.id}}, s_{\text{sent}}, s_{\text{rcvd}}) \\ & = (s'_{\text{pcert.pk}}, s'_{\text{pcert.id}}, s'_{\text{acert.pk}}, s'_{\text{acert.id}}, s'_{\text{rcvd}}, s'_{\text{sent}}) \end{aligned}$$

Similarly, we say that session s' M2-matches (or is M2-matching) session s if s' M1-matches session s and $s_{\text{role}} \neq s'_{\text{role}}$.

Thus, we will use M1 for protocols such as the two-message versions of MQV and HMQV, and we will use M2 for protocols such as NAXOS [16].

Remark 5 (Comparison to matching definition if only one key per identifier is allowed). Note that when users are allowed to have multiple public keys, matching sessions (that is, intended communication partners) of some AKE protocol π do not always compute the same session key if the definition of “matching” does not require agreement on the public keys used in the respective sessions.

The goal of the adversary is to distinguish the session key of a fresh session from a completely random string. This is modelled through an additional query:

- `test-session`(s) This query requires that session s exists and that $s_{\text{status}} = \text{accepted}$; otherwise, it returns \perp . A bit b is chosen at random. If $b = 1$, then s_{key} is returned. If $b = 0$, a random element of $\{0, 1\}^k$ is returned.

⁸ For instance, consider an RSA modulus N and the (cyclic) group QR_N of quadratic residues modulo N . As the CDH problem in QR_N is provably as hard as factoring N [32, 40], it is conceivable to instantiate a DL-based key agreement protocol (like UM, HMQV, etc.) with that group. This would result in $\mathbf{PK} = QR_N$, and `pkregister` and `npkregister` queries would have to perform membership tests for this set. However, the latter is assumed to be a hard problem, by the quadratic residuosity assumption [23].

Definition 4 (ASICS_X experiment). Let Π be an ASICS protocol and $X = (M, Q, F)$ be an ASICS model. We define experiment ASICS_X, between an adversary E and a challenger who implements all users and the CA, as follows:

1. The experiment is initialized with domain parameters for security parameter k .
2. The adversary E can perform any sequence of queries from Q .
3. At some point in the experiment, E issues a test-session query for a session s that has accepted and satisfies F at the time the query is issued.
4. The adversary may continue with queries from Q , under the condition that the test session must continue to satisfy F .
5. Finally, E outputs a bit b' as E 's guess for b .

Definition 5 (ASICS_X advantage). The adversary E wins the security experiment if it correctly outputs the bit b chosen in the test-session query. The ASICS_X-advantage of E is defined as $\text{Adv}_{\Pi, E}^{\text{ASICS}_X}(k) = |2\Pr(b = b') - 1|$.

Definition 6 (ASICS security). Let Π be an ASICS protocol and $X = (M, Q, F)$ be an ASICS model. Π is said to be secure in ASICS model X if, for all PPT adversaries E , it holds that

1. if two users successfully accept in M -matching sessions, then they both compute the same session key, and
2. E has no more than a negligible advantage in winning the ASICS_X experiment; that is, there exists a negligible function negl in the security parameter k such that $\text{Adv}_{\Pi, E}^{\text{ASICS}_X}(k) \leq \text{negl}(k)$.

Remark 6 (Implicit authentication). Note that the ASICS security definition, like eCK-style security definitions, only provides implicit peer authentication, meaning that the key could only be known by the peer, not explicit authentication that the peer actually was active in the session.

Remark 7 (Comparison to other models, including BR, CK and eCK). By choosing the appropriate query subset, freshness, and matching definitions, our parameterised model can be instantiated to produce many existing AKE models. The vast majority of existing models assume honestly-generated key pairs that are securely distributed to users. They therefore do not have queries to model `pkregister` and `npkregister`. Our model family includes the models eCK [29] and eCK^w [17], which capture the leakage of session specific ephemeral data, as well as the eCK-PFS model [17] incorporating perfect forward secrecy. The Bellare-Rogaway model [5] can be captured in our approach by including timestamps as part of the messages, to accurately model the matching definition from [5]. Unlike the CK model [13], we do not assume unique session identifiers to be available to the protocol up-front. We do not consider this a drawback as real-world protocols may not have such identifiers either. Instead, our family includes the basic CK_{HMQV} model [27] in which the HMQV protocol was analysed. For the CK and CK_{HMQV} models we have chosen to interpret the `session-state` as the generated randomness, in the same fashion as the models have been used by their authors [13, 27].

Remark 8 (Freshness will depend on the set of allowed queries and on the matching definition). Note that the exact definition of freshness F used in a security model X will likely depend heavily on the set of queries Q available to the adversary in that model, as well as the specific matching definition. For example, when Q includes `pkregister` and/or `npkregister` queries, then we will require an additional restriction on the registration status of the specific public key of the peer used in the test session.

Example 1. Let us consider the following ASICS model as a concrete example. Let $X = (M1, Q, F)$ be the ASICS model given by $Q = Q_N \cup \{\text{session-key}\} \cup Q_R$ and F defined as follows. Given a sequence of queries and a session s , F holds if:

- no `session-key(s)` query has been issued, and
- for all sessions s' such that s' M1-matches s , no query `session-key(s')` has been issued, and
- no query `pkregister(spcert.pk, spcert.id)` has been issued.

The model X is an extension of a BR-like model with a CA that allows registration of arbitrary keys. If a protocol is secure in X , then it is secure even if the adversary can register arbitrary bitstrings as public keys, as long as the specific peer key used in the test session is not an adversary-generated valid public key.

2.3 Capturing attacks

We illustrate how several attacks exploiting the adversary's ability to register valid or invalid public keys can be captured in ASICS models.

Kaliski's online UKS attack against MQV [25]. As noted in Example 1, Kaliski's attack against MQV can be captured in an ASICS model where the adversary can register a specific valid public key with his own identifier via a `pkregister` query. As the adversary knows the secret key corresponding to the registered public key, the attack cannot be prevented by VP requiring a proof-of-possession of the secret key.

UKS attack against KEA based on public-key re-registration [30, p. 380]. Suppose that public key `pk` has been honestly registered at the CA for some user with identifier \hat{P} via the query `hregister(pk, \hat{P})`. In this UKS attack on the KEA protocol, the adversary re-registers the public key `pk` under his own identifier $\hat{L} \neq \hat{P}$ by issuing the query `pkregister(pk, \hat{L})`. The attack is prevented if VP checks for uniqueness of the public key and outputs 0 when the public key was certified before (as observed in [30, p. 381]). Note that the UKS attack can also be prevented by making the session key derivation depend on users' identifiers.

UKS attack against KEA+ based on impersonation attack. Lauter and Mityagin [30] produced the KEA+ protocol from the KEA protocol and Protocol 4 in [7] by incorporating the identifiers of the user and its peer in the session key computation to prevent UKS attacks; however, a similar but previously unreported UKS attack still works on the KEA+ protocol. This UKS attack involves a type of impersonation attack [42, p. 3]: it requires the adversary to successfully impersonate a user to the CA who then issues a certificate containing the user's identifier, but the adversary's valid public key. We stress that the attack does not arise when only one public key per identifier can be registered.

The UKS attack against KEA+ is captured in the ASICS model $X = (\text{M2}, Q, F)$ given by $Q = \mathcal{Q}_N \cup \{\text{session-key}\} \cup \{\text{pkregister}\}$ and F defined as follows. We say that a session s satisfies F if it holds that

- no `session-key(s)` query has been issued, and
- for all sessions s' such that s' M2-matches s , no query `session-key(s')` has been issued, and
- no query `pkregister($s_{\text{pcert.pk}}, s_{\text{pcert.id}}$)` has been issued.

We next show that KEA+ is insecure in ASICS model X . The adversary creates an initiator session via the query `create($s = (C, i), \mathcal{I}, C'$)`, where the public keys $C.\text{pk} = A = g^a$ and $C'.\text{pk} = B = g^b$ were honestly registered via the query `hregister`. Let $C.\text{id} = \hat{A}$. Now the adversary registers a second public key $A' = A^r$, for some r chosen uniformly at random from \mathbb{Z}_p , with identifier \hat{A} , by issuing the query `pkregister(A', \hat{A})` returning certificate C'' . He then creates a responder session via the query `create($s' = (C', j), \mathcal{R}, C''$)` and activates session s' by sending the message $X = g^x$ sent by session s to session s' . The adversary intercepts the message $Y = g^y$ sent by session s' , modifies it by setting $Y' = Y^r$, and then sends message Y' to session s . Session s accepts the key $s_{\text{key}} = H(Y'^a, B^x, \hat{A}, \hat{B})$ as the session key, while session s' accepts as its key $s'_{\text{key}} = H(A'^y, X^b, \hat{A}, \hat{B})$. The completed session s is chosen as the test session. Now a `session-key` query to session s' reveals the session key of the test session. The sessions s and s' are not M2-matching (since $s_{\text{rcvd}} \neq s'_{\text{sent}}$), but compute the same session key. This leads to our UKS attack. Adding either the long-term public keys of both actor and peer or the exchanged messages (in addition to the identifiers) in the session key computation would prevent the attack, since then the session keys computed in both sessions would be different with overwhelming probability (assuming that H is modelled as a random oracle). Note that if KCI attacks are permitted in the model, the UKS attack cannot be prevented by VP requiring a proof-of-possession of the secret key corresponding to the public key A' .

Online UKS attack on STS-MAC based on duplicate-signature key selection (DSKS) [9]. Suppose that the signature scheme employed in the STS-MAC protocol is vulnerable to DSKS attacks. Note that resistance to DSKS attacks is not covered by standard security notions for signatures such as EUF-CMA or SEUF-CMA security, which concern only single keys. The UKS attack on STS-MAC [9, p. 160] exploits the ability of the adversary to register a valid public key `pk` under his own identifier during the run of the protocol. More precisely, the adversary first intercepts a user's message containing a signature σ on message m . He then issues a query `pkregister(pk, \hat{L})` such that σ is also a valid signature on m under `pk`. The query associates `pk` with the adversary's identifier \hat{L} . Since the adversary knows the secret key corresponding to `pk`, he obtains a certificate from the CA even if VP requires a proof-of-possession. Countermeasures to such UKS attacks via modification of the protocol are available [9].

Lim-Lee style attack against HMQV with special domain parameters, without validation of ephemeral public keys [34]. Suppose that the underlying group G is a subgroup of the group of units $U(R)$ of a ring R containing an element $T \neq 0$ such that $T^2 = 0$. Then, the HMQV protocol is vulnerable to the attack described in [34, p. 137] with respect to an ASICS model in which the adversary is only given access to standard protocol execution queries in the set \mathcal{Q}_N . However, as mentioned in [34, p. 134], this attack on HMQV can only be launched in certain exotic groups that have never been considered for practical

use. The attack can be eliminated by restricting the group used for analyzing HMQV in some way or by adding group membership tests on ephemeral public keys to the protocol.

Lim-Lee style attack against HMQV with DSA domain parameters, without validation of ephemeral public keys [33]. Let $G = \langle g \rangle$ denote a q -order subgroup of \mathbb{Z}_p^* , where q and p are prime and $(p-1)/q$ is smooth. The attack on two-pass HMQV [33, p. 5] can be captured in an ASICS model where the adversary is given access to the queries in the set $Q = \mathcal{Q}_N \cup (\mathcal{Q}_S \setminus \{\text{corrupt}\}) \cup (\mathcal{Q}_R \setminus \{\text{pkregister}\})$. In particular, the adversary can register invalid public keys via the `npkregister` query. This attack can be prevented by countermeasures such as requiring VP to include a group membership test on the public key submitted for certification, or by including group membership tests on both ephemeral and long-term public keys during protocol execution. Small-subgroup attacks may also exist in other settings, for instance in groups over elliptic curves.

We revisit some of the previously mentioned attacks in Section 4.2 to illustrate that the modular approach that we develop in Section 3 cannot be applied to the corresponding protocols.

3 Achieving ASICS security

We provide a modular approach to obtain provable ASICS security for certain types of protocols. We first show in Theorem 1 how a result from Kudla and Paterson [28, Theorem 2] can be adapted to incorporate adversarial registration of valid public keys. Then, in Theorem 2, we indicate how to transform protocols to achieve security in the presence of adversaries that can register arbitrary invalid public keys. We start by defining an adapted version of *strong partnering* [28].

Definition 7 (Strong partnering). *Let Π be an ASICS protocol, and let $X = (M, Q, F)$ be an ASICS model. We say that Π has strong partnering in the ASICS_X experiment if no PPT adversary, when attacking Π in the ASICS_X experiment, can establish two sessions s and s' of protocol Π holding the same session key without being M -matching, with more than negligible probability in the security parameter k .*

Given an ASICS model $X = (M, Q, F)$, we denote by $\text{cNR-}X$ (using terminology from [28], where *cNR* stands for “computational No-Reveals”, referring to the absence of session-key reveals) the reduced *computational* ASICS_X experiment which is similar to the ASICS_X experiment except that the adversary (a) is not allowed to issue `session-key` and `test-session` queries, (b) must pick a session that has accepted and satisfies F at the end of its execution, and (c) output the session key for this session. See Kudla and Paterson [28] for a more detailed description of reduced games.

Definition 8 (cNR- X security). *Let Π be an ASICS protocol and $X = (M, Q, F)$ be an ASICS model. Π is said to be cNR- X -secure if, for all PPT adversaries E , it holds that*

1. *if two users successfully accept in M -matching sessions, then they both compute the same session key, and*
2. *E has no more than a negligible advantage in winning the cNR- X experiment; that is, there exists a negligible function negl in the security parameter k such that $\text{Adv}_{\Pi, E}^{\text{cNR-}X}(k) \leq \text{negl}(k)$, where $\text{Adv}_{\Pi, E}^{\text{cNR-}X}(k)$ is defined as the probability that E outputs (s, s_{key}) for a session s that has accepted and satisfies F .*

Our first theorem deals with the security of DH-type ASICS protocols, which are a generalization of DH-type AKE protocols of Cremers and Feltz [17] to include certificates and to explicitly identify session strings. This class of protocols includes the most prominent modern two-message AKE protocols.

Definition 9 (DH-type ASICS protocol). *A DH-type ASICS protocol is an ASICS protocol of the following form, specified by functions $f_{\mathcal{I}}, f_{\mathcal{R}}, F_{\mathcal{I}}, F_{\mathcal{R}}, H$:*

- *Domain parameters (G, g, q) , where $G = \langle g \rangle$ is a group of prime order q generated by g .*
- *KeyGen(): Choose $a \in_R [0, q-1]$. Set $A \leftarrow g^a$. Return secret key $\text{sk} = a$ and public key $\text{pk} = A$.*
- *VP(x, \hat{P}) = 1 for all x and all \hat{P} (i.e., the CAs do not perform any checks).*
- *When the initiator is activated to create a new session s with a `create($s = (C, i), r = \mathcal{I}, C'$)` query, where the secret key corresponding to $C.\text{pk}$ is a , the initiator computes an outgoing ephemeral public key $X \leftarrow g^{f_{\mathcal{I}}(s_{\text{rand}}, a, C, C')}$ and returns X as an outgoing message.*
- *The responder is activated to create a new session s' with a `create($s = (C', j), r = \mathcal{R}, C$)` query, where the secret key corresponding to $C'.\text{pk}$ is b .*
- *When the responder is activated in a session s' with a `send($s', M = X$)`, the responder computes an outgoing ephemeral public key $Y \leftarrow g^{f_{\mathcal{R}}(s'_{\text{rand}}, b, C', C)}$ and returns Y as an outgoing message. The responder computes a session string (an intermediate value to which we give a special name) $ss' \leftarrow F_{\mathcal{R}}(f_{\mathcal{R}}(s'_{\text{rand}}, b, C', C), b, X, C', C)$, a session key $s'_{\text{key}} \leftarrow H(ss')$, and accepts: $s'_{\text{status}} \leftarrow \text{accepted}$.*

- When the initiator is activated in a session s with a $\text{send}(s, M = Y)$, the initiator computes a session string $ss \leftarrow F_{\mathcal{I}}(f_{\mathcal{I}}(s_{\text{rand}}, a, C, C'), a, Y, C, C')$, a session key $s_{\text{key}} \leftarrow H(ss)$, and accepts: $s_{\text{status}} \leftarrow \text{accepted}$.

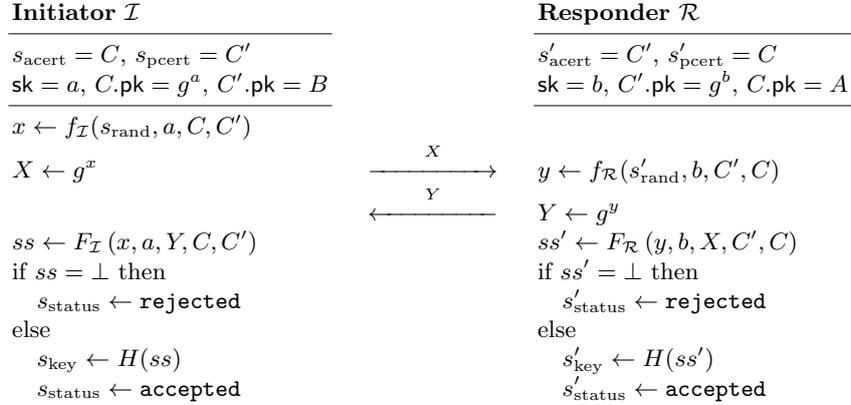


Fig. 1. Messages for generic DH-type ASICS protocol

Theorem 1. Let $X = (M, Q, F)$ be an ASICS model with $\mathcal{Q}_N \subseteq Q \subseteq \mathcal{Q}_N \cup \mathcal{Q}_S$. Let $Y = (M, Q', F')$ be the ASICS model where $Q' = Q \cup \{\text{pkregister}\}$ and F' is defined as follows. A session s is said to satisfy F' if it satisfies F and no $\text{pkregister}(s_{\text{pcert}}.\text{pk}, s_{\text{pcert}}.\text{id})$ query has been issued. Let Π be a DH-type ASICS protocol. Suppose that

- Π has strong partnering in the ASICS_Y experiment,
 - cNR-X security of the related protocol π (defined in the same way as Π except that the session key generated in π is the session string of Π (i.e., $s_{\text{key}}^\pi = ss^\Pi$)) is probabilistic polynomial-time reducible to the hardness of the computational problem of some relation ϕ ,
 - the session string decisional problem in the ASICS_Y experiment for Π is polynomial-time reducible to the decisional problem of ϕ , and
 - there is a polynomial-time algorithm that decides whether an arbitrary bitstring is an element of G ,
- then the security of Π in ASICS model Y is probabilistic polynomial-time reducible to the hardness of the gap problem of ϕ , if H is modelled as a random oracle.

In the cNR-X experiment of Theorem 1 the queries session-key and pkregister are not allowed, whereas in ASICS_Y both queries are allowed. Theorem 1 states that for any DH-type protocol Π , under certain conditions, it holds that security of the related protocol π in a reduced model (in which public keys can only be honestly registered) implies security of Π in the stronger non-reduced model that additionally captures adversarial registration of valid public keys.

The following theorem, which is applicable to a wider range of protocols than Theorem 1 (e.g., to three-message protocols such as UM [35] or HMQV-C [26]), allows us to achieve security against adversaries that can obtain certificates from the CA for invalid public keys by transforming the protocol to include a group membership test on the peer's public key. In contrast to Theorem 1, no additional requirement is imposed on the freshness condition of model Y .

Theorem 2. Let $X = (M, Q, F)$ be an ASICS model with $\mathcal{Q}_N \subseteq Q \subseteq \mathcal{Q}_N \cup \mathcal{Q}_S \cup (\mathcal{Q}_R \setminus \{\text{npkregister}\})$.

Let Π be an ASICS protocol where the domain parameters (G, g, q) , the key generation algorithm KeyGen and the verification procedure VP are as in Definition 9.

Let $f(\Pi)$ denote the ASICS protocol derived from Π by adding the following protocol step for each role of the protocol. Upon creation with (or, via send , receipt of) the certificate C' to be used for the peer of session s , the user running session s checks whether the public key $C'.\text{pk}$ belongs to the group G before continuing the execution of the protocol. In case the check fails, the protocol execution is aborted and s_{status} is set to rejected .

Suppose that protocol Π is secure in ASICS model X and that there is a polynomial-time algorithm that decides whether an arbitrary bitstring is an element of G . Then the transformed protocol $f(\Pi)$ is secure in ASICS model $Y = (M, Q \cup \{\text{npkregister}\}, F)$.

Combining both theorems, we obtain the following result.

Corollary 1. *Let Π be a DH-type ASICS protocol. Let $X = (M, Q, F)$ and $Y = (M, Q', F')$ be defined as in Theorem 1, and let the conditions of Theorem 1 hold with respect to protocol Π . Let $f(\Pi)$ denote the protocol derived from Π as specified in Theorem 2. Then the transformed protocol $f(\Pi)$ is secure in ASICS model $Z = (M, Q'', F')$, where $Q'' = Q' \cup \{\text{npkregister}\}$, if H is modelled as a random oracle.*

Applying Corollary 1 to a concrete DH-type ASICS protocol that satisfies all the preconditions, we obtain a protocol that is secure in an ASICS model in which (a) sessions (including the test session) may use a certificate for the peer that resulted from an `npkregister` query, and (b) the certificate of the test session's peer was not the result of a `pkregister` query.

4 Applications

To illustrate the power of our generic approach, we examine in Section 4.1 how to apply our technique to Ustaoglu's CMQV protocol [44]. CMQV is a modern DH-type protocol that is comparable in efficiency to HMQV, but enjoys a simpler security proof in the eCK model.

Our results allow us to analyse CMQV in a model that does not include `session-key`, `pkregister`, and `npkregister` queries, which simplifies the overall proof. We verify that CMQV meets the preconditions of Corollary 1, and conclude that a variant of CMQV with group membership test on the peer's public key is ASICS-secure in an eCK-like model. Similarly, our generic approach can be applied to other DH-type candidates such as NAXOS [29] and UP [45]. We discuss in Section 4.2 where the preconditions of Corollary 1 are violated for protocols from Section 2.3 such as MQV or HMQV.

4.1 CMQV

CMQV [44] was originally proven secure in the eCK model, where there is only one public key per identifier. In the ASICS setting, there is no such unique mapping between user identifiers and public keys. Hence, to be able to prove CMQV secure in the ASICS model, we need to include the public keys of the users in the session string to ensure that they have the same view of these keys when deriving the session key.

CMQV as a DH-type ASICS protocol. Two-pass CMQV can be stated as a DH-type ASICS protocol, by instantiating Definition 9 with the following functions. Let $\mathcal{H}_1 : \{0, 1\}^k \times \mathbb{Z}_q^* \rightarrow \mathbb{Z}_q^*$, $\mathcal{H}_2 : \{0, 1\}^* \rightarrow \mathbb{Z}_q$, and $H : \{0, 1\}^* \rightarrow \{0, 1\}^k$ be hash functions. We define $f_{\mathcal{I}}, f_{\mathcal{R}}, F_{\mathcal{I}}, F_{\mathcal{R}}$ as:

$$\begin{aligned} f_{\mathcal{I}}(r, a, C, C') &= \mathcal{H}_1(r, a) \\ F_{\mathcal{I}}(x, a, Y, C, C') &= \begin{cases} \perp & , \text{ if } Y \notin G \setminus \{1\} \\ ((YB^e)^{x+da} \parallel g^x \parallel Y \parallel C.\text{id} \parallel A \parallel C'.\text{id} \parallel B) & , \text{ if } Y \in G \setminus \{1\} \end{cases} \\ f_{\mathcal{R}}(r, b, C', C) &= \mathcal{H}_1(r, b) \\ F_{\mathcal{R}}(y, b, X, C', C) &= \begin{cases} \perp & , \text{ if } X \notin G \setminus \{1\} \\ ((XA^d)^{y+eb} \parallel X \parallel g^y \parallel C.\text{id} \parallel A \parallel C'.\text{id} \parallel B) & , \text{ if } X \in G \setminus \{1\}, \end{cases} \end{aligned}$$

where $d = \mathcal{H}_2(X \parallel C.\text{id} \parallel C'.\text{id})$, $e = \mathcal{H}_2(Y \parallel C.\text{id} \parallel C'.\text{id})$, $A = C.\text{pk}$, $B = C'.\text{pk}$; \parallel denotes tagged concatenation to avoid ambiguity with variable-length strings.

We now show, using Corollary 1, that the resulting DH-type CMQV protocol is a secure ASICS protocol in an ASICS model with leakage queries corresponding to the eCK model.

ASICS model for eCK-like leakage. Define the ASICS model $\text{eCK} = (M2, Q, F)$ for eCK-like leakage [29] as follows. Let $Q = Q_N \cup Q_S$. Let F be the condition that a session s satisfies F if, for all sessions s' such that s' M2-matches s , none of the following conditions hold:

- a `session-key(s)` query has been issued;
- if s' exists:
 - a `session-key(s')` query has been issued;
 - both `corrupt(sacert.pk)` and `randomness(s)` queries have been issued;
 - both `corrupt(s'acert.pk)` and `randomness(s')` queries have been issued;
- if s' does not exist:
 - both `corrupt(sacert.pk)` and `randomness(s)` queries have been issued;

- a $\text{corrupt}(s_{\text{pcert}}.\text{pk})$ query has been issued.

Theorem 3. *Let $f(\text{CMQV})$ be the DH-type ASICS protocol derived from the CMQV protocol defined above, as specified in Theorem 2. If $\mathcal{H}_1, \mathcal{H}_2$ and H are modelled as random oracles, G is a group where the gap Diffie–Hellman assumption holds and membership in G is decidable in polynomial time, then $f(\text{CMQV})$ is secure in ASICS model $Z = (\text{M2}, \mathcal{Q}_N \cup \mathcal{Q}_S \cup \mathcal{Q}_R, F')$, where a session s is said to satisfy F' if it satisfies the freshness condition F from the eCK model and no $\text{pkregister}(s_{\text{pcert}}.\text{pk}, s_{\text{pcert}}.\text{id})$ query has been issued.*

Proof (Sketch). We can readily show that CMQV satisfies the preconditions of Corollary 1 under the above formulation of the eCK model as an ASICS model:

1. *Strong partnering.* It is straightforward to see that CMQV has strong partnering in the $\text{ASICS}_{\text{eCK}'}$ game (where eCK' is derived from eCK as described in Theorem 1): since the session key in CMQV is computed via a random oracle, the probability that two sessions derive the same session key without using the same session string input to the random oracle is negligible.
2. *cNR-eCK-security of the session string variant of CMQV.* This can be shown by an adaptation of Ustaoglu’s original proof of CMQV. In large part, the main proof can be followed. However, a few simplifications can be made because the simulation need not answer session-key queries (so preventing key replication attacks and simulating sessions where the public key is a challenge value are easier).
3. *Hardness of the session string decision problem.* It can be easily seen that this is polynomial-time reducible to the decisional problem for Diffie–Hellman triples (U, V, W) by noting that the first component of the CMQV session string σ is equal to $g^{(y+eb)(x+da)} = g^{xy}g^{ady}g^{bex}g^{abde}$; the DDH values (U, V) can be injected into either (X, Y) , (A, Y) , (B, X) , or (A, B) , with W inserted into the corresponding part of σ , yielding a polynomial-time reduction.

Detailed proofs of each of the above claims can be found in Appendix B.

4.2 Discussion

In Section 4.1 we provided evidence that our modular approach can indeed be successfully applied to reason about the security of prominent key exchange protocols. Here we inspect which preconditions of Corollary 1 are violated for some commonly analyzed protocols.

Let $X_1 = (\text{M1}, Q, F)$ be the ASICS model given by $Q = \mathcal{Q}_N \cup \{\text{session-key}\}$ and F defined as follows. We say that a session s satisfies F if it holds that neither a $\text{session-key}(s)$ query nor a $\text{session-key}(s')$ query for any session s' that is M1-matching s have been issued. Let $Y_1 = (\text{M1}, Q', F')$ be derived from X_1 as specified in Theorem 1. The UKS attacks described in Section 2.3 against the DH-type protocols MQV and KEA caused two sessions that were not M1-matching to compute the same session key. Thus, the first precondition of Corollary 1 is violated for these protocols as they do not have strong partnering in the ASICS_{Y_1} experiment. The two-pass HMQV protocol without ephemeral public key validation is insecure in ASICS model X_1 when special domain parameters are used (see Section 2.3) due to the lack of ephemeral public key validation. Hence, it is easy to see that the related protocol, that we call HMQV', is not cNR- X_1 -secure, which violates the second precondition of Corollary 1. Even though it is unclear whether there is an attack against two-pass HMQV with DSA domain parameters without ephemeral public key validation, the third precondition of Corollary 1 is violated for this version of the HMQV protocol as the lack of ephemeral public key validation hinders us from using a DDH oracle to solve the session string decisional problem for the protocol since there is no guarantee that all inputs to the DDH oracle belong to the group G . The latter observations on HMQV lead us to the conclusion that validation of ephemeral public keys is necessary to be able to apply Corollary 1 to the HMQV protocol, even when the adversary is not given access to the randomness query.

Let $X_2 = (\text{M2}, Q, F)$ be the ASICS model given by $Q = \mathcal{Q}_N \cup \{\text{session-key}\}$ and F defined as follows. We say that a session s satisfies F if it holds that neither a $\text{session-key}(s)$ query nor a $\text{session-key}(s')$ query for any session s' that is M2-matching s have been issued. Let $Y_2 = (\text{M2}, Q', F')$ be derived from X_2 as specified in Theorem 1. Note that model Y_2 corresponds to the model X defined in Section 2.3 capturing the UKS attack against KEA+. The latter UKS attack described against the DH-type protocol KEA+ caused two sessions that were not M2-matching to compute the same session key. Thus, the first precondition of Corollary 1 is violated for KEA+ as this protocol does not have strong partnering in the ASICS_{Y_2} experiment.

5 Lessons learned and recommendations

As we started our systematic investigation we assumed that certification authorities would need to perform some minimal checks on public keys to obtain secure KE protocols. Perhaps surprisingly, nearly all of the effort can be shifted to the protocols; and modern protocols often perform sufficient checks. In particular, our results provide formal foundations for some of the protocol-specific observations of Menezes and Ustaoglu [34]: checking that short-term as well as long-term public keys are in the key space (i.e., in group G for DH-type protocols) is not superfluous.

Based on these observations, and given M public keys, N users may need to perform on the order of $M \times N$ such checks in total, even when caching the results. Reasoning purely about the overall amount of computation time used, one could consider moving the burden to the CAs. If the CAs only create certificates after a successful check, the CAs would only perform on the order of M checks in total. Depending on the deployment scenario, this might be a preferable alternative.

Similarly, CAs do not necessarily need to check uniqueness of public keys. As long as the key derivation involves the identifiers in an appropriate way, UKS attacks such as the one on KEA can be prevented. Even if public keys are associated with multiple identifiers, secrecy of the corresponding private key is sufficient to enable ASICS security for the honest user.

In general, our results further justify using as much information as possible in the key derivation function (KDF). This helps with establishing formal security proofs and it is also a prudent engineering principle. In particular, we recommend that in settings where users may have multiple long-term public keys, the input to the KDF should not only include the identifiers and the message transcript, but also the specific public keys used in the session.

We hope our work can serve as a foundation for the development of a range of protocols specifically designed to incorporate certification systems, offering different tradeoffs between efficiency and trust assumptions of the involved parties.

References

1. Adams, C., Farrell, S., Kause, T., Mononen, T.: Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP). RFC 4210 (Proposed Standard) (Sep 2005), <http://www.ietf.org/rfc/rfc4210.txt>, updated by RFC 6712
2. Barker, E., Barker, W., Burr, W., Polk, W., Smid, M.: Recommendation for key management — Part 1: General. NIST Special Publication (March 2007), http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf
3. Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: Juels, A., Wright, R.N., Vimercati, S. (eds.) ACM CCS 06. pp. 390–399. ACM Press (2006)
4. Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated key exchange secure against dictionary attacks. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 139–155. Springer (2000)
5. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO'93. LNCS, vol. 773, pp. 232–249. Springer (1994)
6. Bellare, M., Rogaway, P.: Provably secure session key distribution: The three party case. In: 27th ACM STOC. pp. 57–66. ACM Press (1995)
7. Blake-Wilson, S., Johnson, D., Menezes, A.: Key agreement protocols and their security analysis. In: Darnell, M. (ed.) 6th IMA International Conference on Cryptography and Coding. LNCS, vol. 1355, pp. 30–45. Springer (1997)
8. Blake-Wilson, S., Menezes, A.: Entity authentication and authenticated key transport protocols employing asymmetric techniques. In: Christianson, B., et al. (eds.) Security Protocols Workshop. Lecture Notes in Computer Science, vol. 1361, pp. 137–158 (1997)
9. Blake-Wilson, S., Menezes, A.: Unknown key-share attacks on the station-to-station (STS) protocol. In: Imai, H., Zheng, Y. (eds.) PKC'99. LNCS, vol. 1560, pp. 154–170. Springer (1999)
10. Boldyreva, A., Fischlin, M., Palacio, A., Warinschi, B.: A closer look at PKI: Security and efficiency. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 458–475. Springer (2007)
11. CA/Browser Forum: Baseline requirements for the issuance and management of publicly-trusted certificates, version 1.1 (2011), https://cabforum.org/Baseline_Requirements_V1_1.pdf
12. CA/Browser Forum: Guidelines for the issuance and management of extended validation certificates, version 1.4 (2012), https://cabforum.org/Guidelines_v1_4.pdf
13. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer (2001)
14. Cash, D., Kiltz, E., Shoup, V.: The twin Diffie-Hellman problem and applications. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer (2008)

15. Chatterjee, S., Menezes, A., Ustaoglu, B.: Combined security analysis of the one- and three-pass Unified Model key agreement protocols. In: Gong, G., Gupta, K.C. (eds.) INDOCRYPT 2010. LNCS, vol. 6498, pp. 49–68. Springer (2010)
16. Cremers, C.: Examining indistinguishability-based security models for key exchange protocols: the case of CK, CK-HMQV, and eCK. In: Cheung, B.S.N., Hui, L.C.K., Sandhu, R.S., Wong, D.S. (eds.) ASIACCS 11. pp. 80–91. ACM Press (2011)
17. Cremers, C.J.F., Feltz, M.: Beyond eCK: Perfect forward secrecy under actor compromise and ephemeral-key reveal. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 734–751. Springer (2012)
18. Ducklin, P.: The TURKTRUST SSL certificate fiasco — what really happened, and what happens next? (January 2013), <http://nakedsecurity.sophos.com/2013/01/08/the-turktrust-ssl-certificate-fiasco-what-happened-and-what-happens-next/>
19. Farshim, P., Warinschi, B.: Certified encryption revisited. In: Preneel, B. (ed.) AFRICACRYPT 09. LNCS, vol. 5580, pp. 179–197. Springer (2009)
20. FOX IT: Black Tulip: Report of the investigation into the DigiNotar Certificate Authority breach (2012), <http://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/rapporten/2012/08/13/black-tulip-update/black-tulip-update.pdf>
21. Freire, E.S.V., Hofheinz, D., Kiltz, E., Paterson, K.G.: Non-interactive key exchange. In: Kurosawa, K., Hanaoka, G. (eds.) Public Key Cryptography. Lecture Notes in Computer Science, vol. 7778, pp. 254–271. Springer (2013)
22. Goldberg, I., Stebila, D., Ustaoglu, B.: Anonymity and one-way authentication in key exchange protocols. *Designs, Codes and Cryptography* 67(2), 245–269 (May 2013)
23. Goldwasser, S., Micali, S.: Probabilistic encryption. *Journal of Computer and System Sciences* 28(2), 270–299 (1984)
24. Jeong, I.R., Katz, J., Lee, D.H.: One-round protocols for two-party authenticated key exchange. In: Jakobsson, M., Yung, M., Zhou, J. (eds.) ACNS 04. LNCS, vol. 3089, pp. 220–232. Springer (2004)
25. Kaliski, B.S.: An unknown key-share attack on the MQV key agreement protocol. In: *ACM Transactions on Information and System Security (TISSEC)*. vol. 4, pp. 275–288. ACM New York, NY, USA (August 2001)
26. Krawczyk, H.: HMQV: A high-performance secure Diffie-Hellman protocol. *Cryptology ePrint Archive*, Report 2005/176 (2005), <http://eprint.iacr.org/>
27. Krawczyk, H.: HMQV: A high-performance secure Diffie-Hellman protocol. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 546–566. Springer (2005)
28. Kudla, C., Paterson, K.G.: Modular security proofs for key agreement protocols. In: Roy, B.K. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 549–565. Springer (2005)
29. LaMacchia, B.A., Lauter, K., Mityagin, A.: Stronger security of authenticated key exchange. In: Susilo, W., Liu, J.K., Mu, Y. (eds.) ProvSec 2007. LNCS, vol. 4784, pp. 1–16. Springer (2007)
30. Lauter, K., Mityagin, A.: Security analysis of KEA authenticated key exchange protocol. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 378–394. Springer (2006)
31. Lim, C.H., Lee, P.J.: A key recovery attack on discrete log-based schemes using a prime order subgroup. In: Kaliski Jr., B.S. (ed.) CRYPTO’97. LNCS, vol. 1294, pp. 249–263. Springer (1997)
32. McCurley, K.S.: A key distribution system equivalent to factoring. *Journal of Cryptology* 1(2), 95–105 (1988)
33. Menezes, A.: Another look at HMQV. *Cryptology ePrint Archive*, Report 2005/205 (2005), <http://eprint.iacr.org/>
34. Menezes, A., Ustaoglu, B.: On the importance of public-key validation in the MQV and HMQV key agreement protocols. In: Barua, R., Lange, T. (eds.) INDOCRYPT 2006. LNCS, vol. 4329, pp. 133–147. Springer (2006)
35. Menezes, A., Ustaoglu, B.: Security arguments for the UM key agreement protocol in the NIST SP 800-56A standard. In: Abe, M., Gligor, V. (eds.) ASIACCS 08. pp. 261–270. ACM Press (2008)
36. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *Journal of Cryptology* 13(3), 361–396 (2000)
37. Ristenpart, T., Yilek, S.: The power of proofs-of-possession: Securing multiparty signatures against rogue-key attacks. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 228–245. Springer (2007)
38. Schaad, J.: Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF). RFC 4211 (Proposed Standard) (Sep 2005), <http://www.ietf.org/rfc/rfc4211.txt>
39. Shaw, D.: The Camellia Cipher in OpenPGP. RFC 5581 (Informational) (Jun 2009), <http://www.ietf.org/rfc/rfc5581.txt>
40. Shmueli, Z.: Composite Diffie-Hellman public-key generating systems are hard to break. Technical Report No. 356, Computer Science Department, Technion-Israel Institute of Technology (1985)
41. Shoup, V.: On formal methods for secure key exchange (version 4) (November 1999), revision of IBM Research Report RZ 3120 (April 1999) <http://www.shoup.net/papers/skey.pdf>
42. Turner, P., Polk, W., Barker, E.: ITL Bulletin for July 2012: Preparing for and responding to certification authority compromise and fraudulent certificate issuance (2012), http://csrc.nist.gov/publications/nistbul/july-2012_itl-bulletin.pdf (Accessed 12/03/2013)
43. Turner, S.: The application/pkcs10 Media Type. RFC 5967 (Informational) (Aug 2010), <http://www.ietf.org/rfc/rfc5967.txt>

44. Ustaoglu, B.: Obtaining a secure and efficient key agreement protocol from (H)MQV and NAXOS. *Designs, Codes and Cryptography* 46(3), 329–342 (2008)
45. Ustaoglu, B.: Comparing SessionStateReveal and EphemeralKeyReveal for Diffie-Hellman protocols. In: Pieprzyk, J., Zhang, F. (eds.) *ProvSec 2009*. LNCS, vol. 5848, pp. 183–197. Springer (2009)

A Proofs of Theorem 1 and Theorem 2

Proof of Theorem 1. The proof structure is similar to the proof of Theorem 2 in [28]. We denote by Λ the session key space associated to protocol Π . Since the cNR- X security of protocol π is probabilistic polynomial-time reducible to the hardness of the computational problem of some relation ϕ , there exists an algorithm A that on input of a problem instance of the computational problem of ϕ and interacting with an adversary E which has non-negligible probability η of winning the cNR- X game for π in time τ is able to solve the computational problem of ϕ with non-negligible probability $h(\eta)$ and in time $v(\tau)$, for some polynomial functions h and v .

By assumption, the session string decisional problem in the ASICS $_Y$ experiment for Π is polynomial-time reducible to the decisional problem of ϕ . Hence there is an algorithm W which solves the session string decisional problem for Π in polynomial-time τ'' given access to a decisional oracle for ϕ .

Let D be an adversary winning the ASICS $_Y$ experiment against protocol Π with non-negligible probability η' in time τ' . Let K denote the event that D does not query H with the session string ss^* of the test session s^* . Since Π has strong partnering in the ASICS $_Y$ experiment, it holds that, with overwhelming probability, if two sessions compute the same session key, then they must be M -matching. Thus, if event K occurs, then D can only win the experiment with negligible probability $u(k) + 1/|\Lambda|$, where $u(k)$ denotes the probability that D issues a session-key query to a session s that is not M -matching s^* and $s_{\text{key}} = s_{\text{key}}^*$.

We next define an algorithm B which solves the gap problem of ϕ with non-negligible probability $h'(\eta')$ and in time $v'(\tau')$, for some polynomial functions h' and v' , using adversary D as a subroutine. B will also run algorithm A on the problem instance of the computational problem of ϕ , and an algorithm that decides, in polynomial-time τ''' , whether an arbitrary bitstring pk submitted for certification is an element of G . We now define B 's responses to D 's queries for the pre-specified peer setting; the post-specified peer case proceeds similarly. Algorithm B maintains sets of certificates \mathcal{C}_h and \mathcal{C}_{pk} as well as lists H -List and G -List, all of which are initially empty.

1. $q \in Q \cap \{\text{kgen, randomness, corrupt}\}$: B forwards the query to A and passes A 's response back to D .
2. $\text{hregister}(\text{pk}, \hat{P})$: B forwards the query to A and passes A 's response back to D . In case A returns a certificate C , B adds C to the set \mathcal{C}_h , i.e., $\mathcal{C}_h \leftarrow \mathcal{C}_h \cup \{C\}$.
3. $\text{pkregister}(\text{pk}, \hat{P})$: B checks whether $\text{pk} \in G$ and VP outputs 1 on input pk and \hat{P} . If all the checks succeed, then B adds a certificate C to the set \mathcal{C}_{pk} , i.e., $\mathcal{C}_{\text{pk}} \leftarrow \mathcal{C}_{\text{pk}} \cup \{C\}$, and returns C . Else, B returns \perp .
4. $\text{create}(s = (C, i), r, C')$: B checks whether $C \in \mathcal{C}_h$, a session s with counter i has not yet been created, $r \in \{\mathcal{I}, \mathcal{R}\}$, and $C' \in \mathcal{C}_h \cup \mathcal{C}_{\text{pk}}$. If one of the checks fails, then B returns \perp . Else if $C' \in \mathcal{C}_{\text{pk}}$, then B answers D 's query by simulating the protocol execution itself. Else, B forwards the query to A and passes A 's response (if any) to D .
5. $\text{send}(s, M)$: If session s does not exist or if $s_{\text{status}} \neq \text{active}$, then B returns \perp . Else if $s_{\text{pcert}} \in \mathcal{C}_{\text{pk}}$, then B responds to the query by simulating the protocol execution itself. Else B forwards the query to A and passes A 's response (if any) to D .
6. H query: To answer D 's queries to the random oracle for H , B stores entries of the form (x_i, λ_i) with $\lambda_i \in \Lambda$ in the H -List. When D makes a query x to the random oracle for H , B determines the return value for D as follows:
 - If there exists an entry (x_i, λ_i) in the H -List with $x_i = x$, then return λ_i .
 - Else if there exists an entry $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, s_{\text{role}}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda_i)$ in the G -List, for some session s that has accepted and $\lambda_i \in \Lambda$, such that x is the session string of session s (i.e., $x = ss$) using algorithm F , then store the entry (x, λ_i) in the H -List and return λ_i .
 - Else, B chooses $\lambda \in_R \Lambda$, stores the entry (x, λ) in the H -List and return λ .
7. $\text{session-key}(s)$: To answer D 's session-key queries, B stores entries of the form $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, s_{\text{role}}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda_i)$ with $\lambda_i \in \Lambda$ in the G -List. When D makes a session-key(s) query to an initiator session s that has accepted, B determines the return value for D as follows:
 - If there exists an entry $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, \mathcal{I}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda_i)$ in the G -List, for some $\lambda_i \in \Lambda$, then return λ_i .

- Else if there exists an entry $(s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, \mathcal{R}, s_{\text{rcvd}}, s_{\text{sent}}, \lambda_i)$ in the G -List, then B stores the entry $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, \mathcal{I}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda_i)$ in the G -List and returns λ_i .
- Else if there exists an entry of the form (x_i, λ_i) in the H -List, where $x_i = ss$ using algorithm F , then B stores the entry $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, \mathcal{I}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda_i)$ in the G -List and returns λ_i .
- Else, B chooses $\lambda \in_R \Lambda$, stores the entry $(s_{\text{acert}}.\text{id}, s_{\text{acert}}.\text{pk}, s_{\text{pcert}}.\text{id}, s_{\text{pcert}}.\text{pk}, \mathcal{I}, s_{\text{sent}}, s_{\text{rcvd}}, \lambda)$ in the G -List, and returns λ .

A session-key query to a responder session that has accepted is answered similarly.

8. $\text{test-session}(s^*)$: B selects $\mu \in_R \Lambda$ and returns μ to D .
9. D outputs a guess: B aborts with failure.

B can detect the complementary event K^c by checking which of the entries (x_i, λ_i) in the H -List has $x_i = ss^*$ using algorithm W . B then passes x_i to A . Since the test session s^* must be fresh, no $\text{pkregister}(s_{\text{pcert}}^*.\text{pk}, s_{\text{pcert}}^*.\text{id})$ occurred in the ASICS_Y experiment and hence the certificate s_{pcert}^* has been output through an $\text{hregister}(s_{\text{pcert}}^*.\text{pk}, s_{\text{pcert}}^*.\text{id})$ query. A solves the computational problem of ϕ with non-negligible probability $h(\eta)$, where $\eta = \eta'(1 - u(k) - 1/|A|)$. B is successful by outputting A 's solution to the instance of the computational problem of ϕ and solves the gap problem of ϕ with non-negligible probability $h(\eta)$ and in time $v(\tau)$, where $\tau = \tau' + \tau''n_H(n_{\text{session-key}} + 1) + \tau'''n_{\text{pkregister}}$ with $n_H, n_{\text{session-key}}$ and $n_{\text{pkregister}}$ denoting the number of H , session-key and pkregister queries issued by D , respectively. \square

Remark 9. We cannot show that Theorem 1 holds for more complex protocols Π such as UM or HMVQ-C in arbitrary ASICS base models as the simulation of non-test sessions s of Π with s_{pcert} being the result of a pkregister query cannot be performed in the appropriate way without the knowledge of long-term secret keys and without violating the freshness condition.

Proof of Theorem 2. Let Π be an ASICS protocol secure in model X . It is straightforward to verify the first condition of Definition 6, that is, that M -matching sessions of protocol $f(\Pi)$ compute the same session key. This follows from the fact that M -matching sessions of protocol Π compute the same key as protocol Π is secure in ASICS model X . We next verify that the second condition of Definition 6 holds.

Claim. If there is a PPT adversary E succeeding in the ASICS_Y experiment against protocol $f(\Pi)$ with non-negligible advantage in time τ' , then we can construct a PPT adversary E' succeeding in the ASICS_X experiment against protocol Π with non-negligible advantage in time $v(\tau)$ (for some polynomial function v) using adversary E as a subroutine. Let L be an algorithm that decides, in polynomial-time τ'' , whether an arbitrary bitstring pk submitted for certification is an element of G .

Proof. Fix a PPT adversary E succeeding in the ASICS_Y experiment against protocol $f(\Pi)$ with non-negligible advantage. We define an algorithm E' which succeeds in the ASICS_X experiment against protocol Π with non-negligible advantage using E as a subroutine. Algorithm E' maintains sets of certificates $\mathcal{C}_h, \mathcal{C}_{\text{pk}}$ and \mathcal{C}_{npk} , all of which are initially empty, and answers E 's queries in the pre-specified peer setting as follows.

1. $q \in Q \cap \{\text{kgen}, \text{randomness}, \text{corrupt}, \text{session-key}\}$: E' issues the same query and returns the answer to E .
2. $\text{hregister}(\text{pk}, \hat{P})$: When E issues an $\text{hregister}(\text{pk}, \hat{P})$ query, E' issues the same query and returns the answer to E . In case a certificate C is returned, E' adds C to the set \mathcal{C}_h , i.e., $\mathcal{C}_h \leftarrow \mathcal{C}_h \cup \{C\}$.
3. $\text{pkregister}(\text{pk}, \hat{P})$: When E issues a $\text{pkregister}(\text{pk}, \hat{P})$ query, E' issues the same query and returns the answer to E . In case a certificate C is returned, E' adds C to the set \mathcal{C}_{pk} , i.e., $\mathcal{C}_{\text{pk}} \leftarrow \mathcal{C}_{\text{pk}} \cup \{C\}$.
4. $\text{npkregister}(\text{pk}, \hat{P})$: E' checks whether $\text{pk} \notin G$ (using algorithm L) and VP outputs 1 on input pk and \hat{P} . If the checks succeed (i.e., $\text{pk} \notin G$ and $\text{VP}(\text{pk}) = 1$), then E' returns a certificate C to E and adds C to the set \mathcal{C}_{npk} , i.e., $\mathcal{C}_{\text{npk}} \leftarrow \mathcal{C}_{\text{npk}} \cup \{C\}$. Otherwise, E' returns \perp .
5. $\text{create}(s = (C, i), r, C')$: E' checks whether $C \in \mathcal{C}_h$, a session s with counter i has not yet been created, $r \in \{\mathcal{I}, \mathcal{R}\}$, and $C' \in \mathcal{C}_h \cup \mathcal{C}_{\text{pk}} \cup \mathcal{C}_{\text{npk}}$. If one of the checks fails, then E' returns \perp . Else if $C' \in \mathcal{C}_h \cup \mathcal{C}_{\text{pk}}$, then E' issues the same query and returns the answer (if any) to E . Else, E' rejects the session creation and sets s_{status} to **rejected**.
6. $\text{send}(s, M)$: If session s does not exist or if $s_{\text{status}} \neq \text{active}$, then E' returns \perp . Else E' issues the same query and returns the response (if any) to E .
7. $\text{test-session}(s)$: When E issues a $\text{test-session}(s)$ query to a session s that has accepted, E' issues the same test-session query and returns the answer to E .
8. At the end of E 's execution, that is, after it has output its guess b' , E' outputs b' as well.

It follows that $\text{Adv}_{f(\Pi), E'}^{\text{ASICS}_Y}(k) \leq \text{Adv}_{\Pi, E'}^{\text{ASICS}_X}(k)$, and adversary E' runs in time $v(\tau)$ with $\tau = \tau' + \tau''n_{\text{npkregister}}$, for some polynomial function v , where $n_{\text{npkregister}}$ denotes the number of npkregister queries made by E .

Since Π is secure in ASICS model X , $\text{Adv}_{f(\Pi), E}^{\text{ASICS}_Y}(k)$ is bounded above by a negligible function in the security parameter k . \square

B Analysis of CMQV

Let $\text{eCK}' = (\text{M2}, Q', F')$ be the ASICS model where $Q' = Q \cup \{\text{pkregister}\}$ and F' is defined as F with the additional requirement that no $\text{pkregister}(s_{\text{pcert}}.\text{pk}, s_{\text{pcert}}.\text{id})$ query has been issued.

Lemma 1. *Let eCK and eCK' be as above. CMQV has strong partnering in the $\text{ASICS}_{\text{eCK}'}$ experiment under the assumption that H is a random oracle.*

Proof. Suppose otherwise. Namely, suppose there exists two sessions s and s' of CMQV that hold the same session key but are not M2-matching. Since the session key in CMQV is derived by applying a random oracle, except with negligible probability, the input to the random oracle in both sessions must be the same. Since they are not M2 matching, either $s_{\text{acert}}.\text{id} \neq s'_{\text{pcert}}.\text{id}$, or $s_{\text{acert}}.\text{pk} \neq s'_{\text{pcert}}.\text{pk}$, or $s_{\text{pcert}}.\text{id} \neq s'_{\text{acert}}.\text{id}$, or $s_{\text{pcert}}.\text{pk} \neq s'_{\text{acert}}.\text{pk}$, or $s_{\text{sent}} \neq s'_{\text{rcvd}}$, or $s_{\text{rcvd}} \neq s'_{\text{sent}}$, or $s_{\text{role}} \neq s'_{\text{role}}$.

First suppose $s_{\text{role}} \neq s'_{\text{role}}$. Then either the public keys, identifiers, or transcripts of the two sessions do not correspond. But these are all inputs to the random oracle, so except with negligible probability the outputs of the random oracle will be different, contradicting that the two sessions hold the same session key.

Now suppose $s_{\text{role}} = s'_{\text{role}}$. Except with negligible probability, two distinct honest sessions will have $s_{\text{rand}} \neq s'_{\text{rand}}$, and hence $s_{\text{sent}} \neq s'_{\text{sent}}$. But since both s and s' think of themselves as the initiator, they will each put their own sent ephemeral public key in the second component of the call to H , and these values are different, so except with negligible probability the outputs of the random oracle will be different, contradicting that the two sessions hold the same key. \square

Let (G, g, q) be as in Definition 9. Let $\phi \subseteq (G \times G) \times G$ be the *Diffie-Hellman relation* on $G = \langle g \rangle$. In particular, (g^a, g^b) is related under ϕ to g^c if and only if $ab \equiv c \pmod{q}$.

Lemma 2. *The cNR-eCK security of the variant of CMQV in which the session string is output as the session key is polynomial-time reducible to the computational problem of the Diffie-Hellman relation ϕ , under the assumption that \mathcal{H}_1 and \mathcal{H}_2 are random oracles.*

The basic idea of the proof is as follows.

- If the adversary happens to figure out a long-term secret key without issuing a **corrupt** query (event E), it must ask that value to a random oracle \mathcal{H}_1 , and we can immediately use that value to solve the CDH problem by having embedded one of the CDH challenge values in that public key.
- If the adversary is passive in the test session (event $\bar{E} \wedge M$) we can embed the CDH challenge values U, V as the ephemeral public keys X and Y of the test session. The adversary's view can be simulated perfectly unless the adversary asks either (\tilde{x}, a) or (\tilde{y}, b) as a query for \mathcal{H}_1 . But the freshness condition prevents the adversary from finding both elements of either pair. Therefore the adversary cannot do better than guess the session string unless it can compute σ . Here the CDH of U and V can be extracted from σ .
- If the adversary is active in the test session (event $\bar{E} \wedge \bar{M}$) we can embed the CDH challenge values in the long-term key of the partner of the test session and the ephemeral public key of the session. As before the simulation is perfect unless the adversary asks (\tilde{x}, a) as a query for \mathcal{H}_1 . Note that, since the adversary is active, the adversary cannot change or corrupt the secret long-term key of the peer. This time the value of σ is similar to a signature forgery and we can apply the Forking Lemma [3, 36] to extract the CDH of U and V .

Proof. Recall that cNR-eCK security means security in an ASICS model that omits the session-key query, so the allowed queries are $\mathcal{Q}_N \cup \{\text{corrupt}, \text{randomness}\}$. The freshness condition remains unchanged.

Recall further that the goal of the adversary is to *recovery* the session string; let S be the event that an algorithm \mathcal{M} computes the session string. The security proof largely follows the original proof of Ustaoglu that CMQV is eCK-secure, but can be simplified somewhat as the queries in the cNR-eCK game are restricted compared to full eCK security.

Consider the following two complementary events:

- E . There exists a certificate C' (created using hregister) such that \mathcal{M} , during its execution, queries $\mathcal{H}_1(*, b)$ (where $C'.\text{pk} = g^b$) before issuing any $\text{corrupt}(C'.\text{pk})$ query (if it issues one at all).

- \overline{E} . During its execution, for every certificate C' (created using `hregister`) for which \mathcal{M} queries $\mathcal{H}_1(*, b)$ (where $C'.\text{pk} = g^b$), it issued a `corrupt`($C'.\text{pk}$) query before the $\mathcal{H}_1(*, b)$ query.

Since the events are complementary, if \mathcal{M} succeeds in computing the session string, it succeeded either when E occurred or when \overline{E} occurred.

We will see how, when each event occurs, the required polynomial-time reduction exists.

E . Suppose event E occurs and \mathcal{M} succeeds in computing the session string.

Here, the simulator \mathcal{S} guesses one public key pk^* at random and assigns $\text{pk}^* \leftarrow V$, where (U, V) is the Diffie–Hellman challenge. All other public keys are generated according to the protocol specification.

For all sessions and queries where the session actor is not using pk^* , \mathcal{S} follows the protocol specification exactly.

For sessions where the session actor is using pk^* , \mathcal{S} responds to queries as follows:

- `hregister`(pk^*, \hat{P}): \mathcal{S} outputs a certificate as normal.
- `corrupt`(pk^*): \mathcal{S} aborts.
- `randomness`($s = (C, i)$) where $C.\text{pk} = \text{pk}^*$: Return s_{rand} .

For sessions where the session actor is using pk^* and is the initiator, \mathcal{S} responds to queries as follows:

- `create`($s = (C, i), \mathcal{I}, C'$) where $C.\text{pk} = \text{pk}^*$: \mathcal{S} selects $x \in_R \mathbb{Z}_q$, computes $X \leftarrow g^x$, and responds with X . Note that s_{rand} is not used in the calculation.
- `send`($s = (C, i), M$) where $s_{\text{acert}}.\text{pk} = \text{pk}^*$ and $s_{\text{role}} = \mathcal{I}$: \mathcal{S} does not need to simulate anything here, since there is not outgoing message required, and since the only variable updated is the session string ss but no session-key reveal query is allowed.

For sessions where the session actor is using pk^* and is the responder, \mathcal{S} responds to queries as follows:

- `create`($s = (C, i), \mathcal{R}, C'$) where $C.\text{pk} = \text{pk}^*$: no response required.
- `send`($s = (C, i), M$) where $s_{\text{acert}}.\text{pk} = \text{pk}^*$ and $s_{\text{role}} = \mathcal{R}$: \mathcal{S} selects $y \in_R \mathbb{Z}_q$, computes $Y \leftarrow g^y$, and responds with Y . Note that s_{rand} is not used in the calculation.

\mathcal{S} responds to \mathcal{H}_2 queries as normal. \mathcal{S} responds to $\mathcal{H}_1(*, b)$ queries as normal for all b such that $g^b \neq \text{pk}^*$. When \mathcal{M} queries $\mathcal{H}_1(*, b)$ where $g^b = \text{pk}^* = V$, \mathcal{S} outputs the solution to the Diffie–Hellman challenge (U, V) as U^b .

Note that \mathcal{S} 's simulation is perfect up until an abort event from the `corrupt` query occurs. Given that event E occurs, there exists some public key $\text{pk} = g^b$ for which the query $\mathcal{H}_1(*, b)$ occurs before any `corrupt`(pk) query occurs. With probability at least $1/n_{\text{kgen}}$, where n_{kgen} is the number of `kgen` queries made by \mathcal{M} , \mathcal{S} this condition holds for pk^* . When \mathcal{S} guesses correctly, \mathcal{M} will indeed query $\mathcal{H}_1(*, b)$ before any `corrupt`(pk^*) query, and thus \mathcal{S} will solve the computational Diffie–Hellman problem.

Thus, when event E occurs, there exists a polynomial-time reduction from a cNR-eCK adversary for the session string variant of CMQV to the computational Diffie–Hellman problem under the assumption that \mathcal{H}_1 is a random oracle, with a tightness factor of n_{kgen} .

\overline{E} . We divide this event into two complementary cases:

- M . The session s for which the adversary output the session string has an M2-matching session s' .
- \overline{M} . The session s for which the adversary output the session string does not have an M2-matching session.

When \overline{E} occurs, either M or \overline{M} must also occur.

$\overline{E} \wedge M$. Suppose event \overline{E} occurs and there is an M2-matching session s' for the target session s .

Here, the simulator guesses two sessions s and s' ; assume without loss of generality that $s_{\text{role}} = \mathcal{I}$ and $s'_{\text{role}} = \mathcal{R}$. \mathcal{S} responds to all `kgen`, `hregister`, `corrupt`, and `randomness` queries as specified by the protocol. For all sessions other than s and s' , \mathcal{S} responds to `create` and `send` as specified by the protocol. For s and s' , \mathcal{S} responds to `create` and `send` as follows:

- `create`($s = (C, i), \mathcal{I}, C'$): Return $X \leftarrow U$, where (U, V) is the Diffie–Hellman challenge. Note that s_{rand} is not used.
- `create`($s' = (C', i), \mathcal{R}, C$): No response required.
- `send`(s', M): Return $Y \leftarrow V$, where (U, V) is the Diffie–Hellman challenge. Note that s'_{rand} is not used.

\mathcal{S} responds to \mathcal{H}_2 queries as normal. \mathcal{S} responds to \mathcal{H}_1 queries as normal except for the queries (\tilde{x}, a) or (\tilde{y}, b) , where a and b are the secret keys corresponding to the public keys in sessions s and s' ; when this occurs, the simulation aborts.

Note that \mathcal{S} 's simulation is perfect unless a $\mathcal{H}_1(\tilde{x}, a)$ or $\mathcal{H}_1(\tilde{y}, b)$ query occurs. Because of event \overline{E} , \mathcal{M} issues a `corrupt`(g^a) query before any $\mathcal{H}_1(\tilde{x}, a)$ query, and a `corrupt`(g^b) query before any $\mathcal{H}_1(\tilde{y}, b)$ query. Since \tilde{x} and \tilde{y} are used in only one session and \mathcal{H}_1 is a random function, no information can be learned

about \tilde{x} and \tilde{y} without $\text{randomness}(s)$ or $\text{randomness}(s')$ queries. By the freshness condition, it cannot be that both $\text{randomness}(s)$ and $\text{corrupt}(g^a)$ occurred, or that both $\text{randomness}(s')$ and $\text{corrupt}(g^b)$ occurred. Thus, if \mathcal{S} correctly guess s and s' , the simulation is perfect and does not abort. This happens with probability at least $2/n_{\text{create}}^2$.

Assuming the simulation is perfect and does not abort, and that \mathcal{M} outputs the session string, \mathcal{S} can use this to solve the Diffie–Hellman problem. In particular, let σ be the shared secret in the session string output by \mathcal{M} . Then \mathcal{S} outputs $\sigma g^{-abed}U^{-be}V^{-ad}$ as the solution to the computational Diffie–Hellman challenge (U, V) .

Thus, when event $\overline{E} \wedge M$ occurs, there exists a polynomial-time reduction from a cNR-eCK adversary for the session string variant of CMQV to the computational Diffie–Hellman problem under the assumption that \mathcal{H}_1 is a random oracle, with a tightness factor of n_{create}^2 .

$\overline{E} \wedge \overline{M}$. Suppose event \overline{E} occurs but there is no M2-matching session for the target session s .

Here, the simulator guesses integers $j \in_R \{1, \dots, n_{\text{kgen}}\}$, and a session s^* . Assume without loss of generality that $s_{\text{role}}^* = \mathcal{I}$.

For the j th query to kgen , \mathcal{S} assigns $\text{pk}^* \leftarrow V$ from the Diffie–Hellman challenge (U, V) to be the public key; for all other kgen queries it responds as specified by the protocol.

All hregister queries are responded to as normal. All corrupt queries are responded to as normal, except for $\text{corrupt}(\text{pk}^*)$, in which case \mathcal{S} aborts.

Suppose that \mathcal{M} selects s^* as the target session and furthermore that $s_{\text{pcert}}^* \cdot \text{pk} = V$.

For all sessions and queries where the session actor or peer is not using pk^* , \mathcal{S} follows the protocol specification exactly.

For sessions where the session actor is using pk^* , \mathcal{S} responds as in event E .

For sessions where the session peer is using pk^* , \mathcal{S} responds as specified by the protocol, except for the target session s^* . In s^* , \mathcal{S} responds as follows:

- $\text{create}(s^*, \mathcal{I}, C')$: \mathcal{S} returns $X \leftarrow U$, where (U, V) is the Diffie–Hellman challenge. Note that s_{rand}^* is not used.
- $\text{send}(s^*, M)$: No response required.
- $\text{randomness}(s^*)$: Return s_{rand}^* .
- $\text{session-key}(s^*)$: \mathcal{S} aborts. Assuming that \mathcal{S} correctly guesses s^* as the target session, this abort will never occur.

\mathcal{S} responds to \mathcal{H}_2 queries as normal. \mathcal{S} responds to $\mathcal{H}_1(\tilde{x}, b)$ queries as normal except for the following two cases:

- If $g^b = \text{pk}^*$: \mathcal{S} outputs the solution to the Diffie–Hellman challenge (U, V) as U^b .
- If $\tilde{x} = s_{\text{rand}}^*$ and $g^a = s_{\text{acert}}^* \cdot \text{pk}$: \mathcal{S} aborts.

Note that \mathcal{S} 's simulation is perfect up until an abort event from the corrupt or the \mathcal{H}_2 query occurs. Given that s^* is fresh and no matching session exists, no $\text{corrupt}(s^*)$ query is allowed and hence \mathcal{S} does not abort for that reason. Given that event \overline{E} occurs, if \mathcal{M} queries $\mathcal{H}_1(s_{\text{rand}}^*, a)$ such that $g^a = s_{\text{acert}}^* \cdot \text{pk}$, \mathcal{M} must have issued a $\text{corrupt}(g^a)$ query first. But it is also the case that s^* is fresh, so \mathcal{M} cannot have also issued a $\text{randomness}(s^*)$ query, and thus cannot know s_{rand}^* unless it guessed it correctly, which can be done only with negligible probability.

Assume the simulation is perfect and does not abort, and that \mathcal{M} outputs the session string containing the correct shared secret $\sigma = g^{uy}g^{ady}g^{uve}g^{adev}$. \mathcal{S} can then compute $\eta = \sigma Y^{-ad}V^{-ade} = g^{uy+uve}$. But the peer's ephemeral secret key y was chosen by the adversary, so without y \mathcal{S} cannot directly compute g^{uv} from η .

Using the Forking Lemma, \mathcal{S} runs \mathcal{M} on the same input and the same random coins but with modified answers to \mathcal{H}_2 queries. Note that \mathcal{M} must have queried $\mathcal{H}_2(Y, s_{\text{acert}}^* \cdot \text{id}, s_{\text{pcert}}^* \cdot \text{id})$ to obtain e , because otherwise \mathcal{M} would be unable to compute σ except with negligible probability. For the second run of \mathcal{M} , \mathcal{S} responds to $\mathcal{H}_2(Y, s_{\text{acert}}^* \cdot \text{id}, s_{\text{pcert}}^* \cdot \text{id})$ with $e' \neq e$ selected uniformly at random.

If \mathcal{M} succeeds in the second run, it outputs $\sigma = g^{uy}g^{ad'y}g^{uve'}g^{ad'e'v}$. \mathcal{S} can then compute $\eta' = \sigma' Y^{-ad'} V^{-ad'e'} = g^{uy}g^{uve'}$. \mathcal{S} can furthermore compute

$$(\eta/\eta')^{1/(e-e')} = (g^{uy}g^{uve}g^{-uy}g^{-uve'})^{1/(e-e')} = g^{uv(e-e')/(e-e')} = g^{uv}$$

which is the solution to the computational Diffie–Hellman challenge (U, V) .

Thus, when event $\overline{E} \wedge \overline{M}$ occurs, there exists a polynomial-time reduction from a cNR-eCK adversary for the session string variant of CMQV to the computational Diffie–Hellman problem under the assumption that \mathcal{H}_1 and \mathcal{H}_2 are random oracles, with a tightness factor of $n_{\text{create}}n_{\text{kgen}}n_{\mathcal{H}_2}c$, where c is a constant from the Forking Lemma. \square

Remark 10. Because in the above lemma we do not have to prove full session key indistinguishability security of CMQV, instead proving the hardness of session string computation of a variant of CMQV, we can make a few simplifications from Ustaoglu’s original proof:

- We do not have to worry about key replication attacks (when the adversary causes two non-matching sessions to have the same session key (that is, session string), and then reveals the session key at one of the sessions) because there is no session-key query.
- In event E , we do not have to worry about setting the session string correctly for any session involving the user whose public key has been injected with the CDH challenge, because there is no session-key query. Thus we do not need a DDH oracle here.
- In event $\bar{E} \wedge M$, we do not have to use the DDH oracle to test which of the many \mathcal{H} random oracle queries is the solution we need: we simply output the CDH value derived directly from the output of \mathcal{M} .

Lemma 3. *The session string decision problem for CMQV is polynomial-time reducible to the decisional problem of the Diffie–Hellman relation ϕ .*

Proof. Let D be a polynomial-time algorithm that can distinguish real CMQV session strings $(g^{(y+eb)(x+da)} \parallel X \parallel Y \parallel \text{id} \parallel A \parallel \text{id}' \parallel B)$ from random session strings $(g^r \parallel g^x \parallel g^y \parallel \text{id} \parallel g^a \parallel \text{id}' \parallel g^b)$, for randomly chosen $a, b, x, y, r \in_R \mathbb{Z}_q$, id and id' are arbitrary binary strings, $d = \mathcal{H}_2(g^x \parallel \text{id} \parallel \text{id}')$, and $e = \mathcal{H}_2(g^y \parallel \text{id} \parallel \text{id}')$.

We claim that there exists an algorithm E that can distinguish real Diffie–Hellman triples (g^u, g^v, g^{uv}) from random triples (g^u, g^v, g^w) for randomly chosen $u, v, w \in_R \mathbb{Z}_q$.

First, note that $g^{(y+eb)(x+da)} = g^{xy+ady+beax+abde}$. Using D construct E_D as follows. Let (U, V, W) be a Diffie–Hellman challenge. Pick arbitrary id, id' . Do one of the following, each with equal probability:

1. Set $A \leftarrow U$ and $B \leftarrow V$. Choose $x, y \in_R \mathbb{Z}_q$.
Run D on the session string $(g^{xy} A^{dy} B^{ex} W^{de} \parallel g^x \parallel g^y \parallel \text{id} \parallel A \parallel \text{id}' \parallel B)$.
2. Set $A \leftarrow U$ and $Y \leftarrow V$. Choose $x, b \in_R \mathbb{Z}_q$.
Run D on the session string $(Y^x W^d g^{bex} A^{bde} \parallel g^x \parallel Y \parallel \text{id} \parallel A \parallel \text{id}' \parallel g^b)$.
3. Set $X \leftarrow U$ and $B \leftarrow V$. Choose $a, y \in_R \mathbb{Z}_q$.
Run D on the session string $(X^y g^{ady} W^e B^{ade} \parallel X \parallel g^y \parallel \text{id} \parallel g^a \parallel \text{id}' \parallel B)$.
4. Set $X \leftarrow U$ and $Y \leftarrow V$. Choose $a, b \in_R \mathbb{Z}_q$.
Run D on the session string $(WY^{ad} X^{be} g^{abde} \parallel X \parallel Y \parallel \text{id} \parallel g^a \parallel \text{id}' \parallel g^b)$.

E outputs the result of D .

Note that in each of the above cases, if (U, V, W) is a real Diffie–Hellman triple, then D is run on a real CMQV session string, whereas if (U, V, W) is a random triple, then D is run on a random session string. Thus, if D is a distinguisher for CMQV session strings, then E is a distinguisher for the Diffie–Hellman relation. \square