

# Stronger Security of Authenticated Key Exchange

Brian LaMacchia

Microsoft Corp.

1 Microsoft Way, Redmond, WA

bal@exchange.microsoft.com

Kristin Lauter

Microsoft Research

1 Microsoft Way, Redmond, WA

klauter@microsoft.com

Anton Mityagin

University of California, San Diego

9500 Gilman Dr., La Jolla, CA

amityagin@cs.ucsd.edu

## Abstract

Recent work by Krawczyk [13] and Menezes [17] has highlighted the importance of understanding well the guarantees and limitations of formal security models when using them to prove the security of protocols. In this paper we focus on security models for two-round authentic key exchange (AKE) protocols. We observe that there are several classes of attacks on AKE protocols that lie outside the boundary of the current class of security models. In an attempt to bring these attacks within the scope of analysis we extend the Canetti-Krawczyk model for AKE security by providing significantly greater powers to the adversary. We then introduce a new AKE protocol called NAXOS and prove that it is secure against these stronger adversaries.

## 1 Introduction

In this paper we extend the Canetti-Krawczyk [10, 13] security model for authenticated key exchange (AKE) to capture all possible attacks resulting from ephemeral and long-term data compromise. Our security model for authenticated key exchange is defined in the spirit of Bellare and Rogaway [4] and Canetti and Krawczyk [10] by an experiment in which the adversary is given many corruption powers for various key exchange sessions and must solve a challenge on a test session. We extend adversarial capabilities to the following extent: the only corruption powers we do not give an adversary in the experiment are those that would trivially break an AKE protocol. We also define a new AKE protocol which is secure in our new model.

More specifically, in an authenticated key exchange protocol, two parties exchange information and compute a secret key as a function of at least four pieces of secret information: their own long-term and ephemeral secret keys and the other party’s static and ephemeral secret keys. Of the four pieces of information, we allow an adversary to reveal any subset of the four which does not contain both the long-term and ephemeral secrets of one of the parties. To explain this more precisely, we divide AKE test sessions (sessions which are subject to attack by an adversary) into two types. Let  $\mathcal{A}$  and  $\mathcal{B}$  be the participants of the test session. In sessions of the first type (“passive” sessions), the adversary does not cancel or modify communications between the two parties. In sessions of the second type (“active” sessions), the adversary may forge the communication of the second party. Another way to phrase the distinction, as done by Krawczyk in the analysis of the HMQV protocol [13], is whether the adversary actively intervenes in the key exchange session or is a passive eavesdropper.

In addition to distinguishing between passive and active sessions, we identify which pieces of secret information the adversary can reveal without being able to trivially break the AKE protocol

(compute the session key for any AKE protocol). In both types of sessions, if an adversary can reveal the long-term and the ephemeral secret keys of one of the parties in the session, then the adversary can trivially compute a session key as it has all the secret information of one of the legitimate parties in the session.

For passive sessions, an adversary may reveal both ephemeral secret keys, both long-term secret keys, or one of each from the two different parties without trivially breaking the protocol. For example, Krawczyk [13] defines weak Perfect Forward Secrecy (wPFS) to be security against revelation of both long-term secret keys after the session is completed (without active adversarial intervention in the session establishment).

For active sessions, the adversary may forge communications from one of the parties. Thus, if the adversary can also reveal the long-term secret key of that same party, then the adversary can trivially compute the session key. The same argument was used by Krawczyk [13] to show that no 2-round AKE protocol can achieve full perfect forward secrecy (PFS). Still, an adversary can reveal a long-term secret key or ephemeral secret key of the other party without trivially breaking the session.

Considering attacks involving both types of sessions, it is natural to define a single security model which captures all of them. In our model, in passive test sessions we allow the adversary to reveal any subset of the four pieces of secret information which does not contain both the long-term and ephemeral secrets of one of the parties. In active test sessions, we allow the adversary to reveal only the long-term secret or the ephemeral secret key of the party which is executing the test session.

One attack scenario not covered by the original Canetti-Krawczyk model is key-compromise impersonation, where the adversary first reveals a long-term secret of a party and then impersonates others to this party. Our extension to the Canetti-Krawczyk model implies weak perfect forward secrecy, security against KCI attacks and security against a number of other attacks not covered by the Canetti-Krawczyk model which were not considered before (see Section 2.2).

We stress that our extension of the security model allows the adversary to register arbitrary public keys for adversary-controlled parties without any checks such as proof-of-possession done by the certificate authority. In contrast, some of the protocols in the literature [14, 15] were proved secure assuming that the key registration is done honestly. Namely, that initially a trusted party generates keys for all, even adversary-controlled parties.

Also we clarify what constitutes the ephemeral secret key of a party in a session and what information can be revealed by an adversary via a session-state reveal query. We require that the ephemeral secret key contains *all* session-specific secret information of a party. We illustrate the need for this requirement with the SIG-DH protocol [10], which provably satisfies Canetti-Krawczyk security if only part of the session-specific secret is revealed and which is vulnerable to adversaries who can reveal all session-specific secret information. Revelations of ephemeral data are motivated by practical scenarios, such as if the state-specific secret information is stored in insecure memory or if the random-number generator of a party is corrupted. In both of these cases, as a result of the specific vulnerability, the adversary might be able to gain access to the ephemeral data.

Finally, we present a new AKE protocol, called NAXOS, which provably meets our definition of AKE security. We prove the security of NAXOS either under the standard Gap Diffie-Hellman assumption. We also improve the concrete security of NAXOS under the related Pairing Diffie-Hellman assumption.

In Figure 1 we compare the efficiency and security of NAXOS with four other recent authen-

ticated key exchange protocols. Efficiency is given as the number of exponentiations executed by one party. Communication in all these protocols (except for Katz-Jeong-Lee) is the same as in the original Diffie-Hellman. The key registration column specifies whether adversary-controlled parties can register arbitrary public keys or if honest key-registration is assumed. The ephemeral column indicates whether an adversary is allowed to reveal ephemeral secret information of the parties. CK denotes Canetti-Krawczyk security without perfect forward secrecy, assuming that partnership is defined via matching conversations. The protocols of Jeong, Katz and Lee [14] and Kudla and Paterson [15]<sup>1</sup> use the Bellare-Rogaway model [4] (BR), which appears to be equivalent to the Canetti-Krawczyk model [9], where no ephemeral reveals are allowed and key-registration is done honestly. KCI denotes security against key-compromise impersonation and wPFS denotes weak perfect forward secrecy. Extended CK denotes our extension of the Canetti-Krawczyk model. Security assumptions include: RO – random oracle model [5], DDH – Decisional Diffie-Hellman, GDH – Gap Diffie-Hellman [19], PDH – Pairing Diffie-Hellman [16] and KEA1 – knowledge of exponent assumption [2].

Protocol	Effic.	Key Reg.	Ephemeral	Security	Assumptions
NAXOS	4	Arbitrary	yes	Extended CK	GDH (or PDH) + RO
HMQR	3	Arbitrary	yes	CK + wPFS + KCI	GDH + KEA1 + RO
KEA+	3	Arbitrary	yes	CK + wPFS + KCI	GDH (or PDH) + RO
Jeong-Katz-Lee	3	Honest	no	BR + wPFS	DDH + secure MACs
Kudla-Paterson	3	Honest	no	BR + KCI	GDH + RO

Figure 1: Comparison of recent AKE protocols.

We begin with a brief review in Section 2 of the Canetti-Krawczyk security model and discuss some attacks not covered by their definition. We introduce our extension of Canetti-Krawczyk in Section 3. In Section 4 we describe the NAXOS protocol and prove its security in the extended model.

## 2 Previous Models

### 2.1 Overview of the Canetti-Krawczyk model

The strongest security definition for AKE protocols was formalized by Canetti and Krawczyk [10]. We give a high-level overview of their model and introduce some notation which will be useful later in the paper. We remark that the model we describe differs from the original definition in that we use session identifiers defined via matching conversations. The same definition was used by Krawczyk when analyzing the security of the HMQR protocol [13] and it is now a commonly used variant of the Canetti-Krawczyk model.

The AKE security experiment involves multiple honest parties and an adversary  $\mathcal{M}$  connected via an unauthenticated network. The adversary selects parties to execute key-exchange sessions and selects an order in which the sessions will be executed. Actions the adversary is allowed to perform include taking full control of any party (a **Corrupt** query), revealing the session key of any session (a

<sup>1</sup>Kudla and Paterson [15] define partnership via matching session identifiers (computed by the parties), although for their protocol this appears to be equivalent to matching conversations.

Reveal query), or revealing session-specific secret information of any session (a **Session-State Reveal** query).

We stress that an AKE session is executed by a single party: since all communication is controlled by an adversary, a party executing a session cannot know for sure with whom it is communicating. The party executing the session is called the *owner* of the session and the other party is called the *peer*. The *matching session* to an AKE session (by the owner with the peer) is the corresponding AKE session which is supposed to be executed by the peer with the owner. The matching session might not exist if the communications were modified by the adversary. The *session identifier* of an AKE session consists of the parties' identities concatenated with messages they exchanged in the session. A completed session is “clean” if the session as well as its matching session (if it exists) is not corrupted (neither session key nor session state were revealed by  $\mathcal{M}$ ) and if none of the participating parties were corrupted.

At some point in the experiment, the adversary is allowed to make one **Test** query: it can select any clean completed session (called the *test session*) and it is given a challenge which consists either of the session key for that session or a randomly selected string. The adversary's goal is to guess correctly which of the cases was selected.

Additionally, the Canetti-Krawczyk [10] definition has an optional perfect forward secrecy (PFS) requirement. In the variant of Canetti-Krawczyk security with PFS, the adversary is allowed to corrupt a participant of the test session (either owner or peer) after the test session is completed. As noted by Krawczyk [13], the PFS requirement is not relevant for 2-round AKE protocols since no 2-round protocol can achieve PFS.

The Canetti-Krawczyk security model is the strongest of a family of models that includes those of Bellare and Rogaway [4, 6] and Bellare, Pointcheval and Rogaway [3]. We refer the reader to Choo et al. [9] for a concise summary of the differences among these various models.

More recently, Krawczyk [13] introduced the notion of *weak perfect forward secrecy* (wPFS) which can be achieved by 2-round protocols and which he demonstrated is achieved by HMQV [13]. Weak PFS guarantees perfect forward secrecy only for those AKE sessions where the adversary didn't modify communications between the parties. (Using the above terminology, the matching session exists for the test session and both test and matching sessions are clean.)

## 2.2 Attacks not Covered by the Existing Definitions

We point out several attacks which are not captured by the previous definitions and explain which components of the Canetti-Krawczyk model prohibit these attacks from being considered. First, we observe that although the adversary is allowed to reveal the session state of the parties, he is not allowed to make **Session-State Reveal** queries against the session he wants to attack (the test session). That is, existing security models do not provide any security guarantees for a session if the ephemeral secret key of either party has been leaked.

Second, when the adversary corrupts an honest party, he takes full control over this party and reveals all its secret information. This definition of **Corrupt** query does not allow attacks where the adversary reveals a long-term secret key of some party prior to the time when that party executes the test session.

Here we summarize some attacks which are not allowed by the Canetti-Krawczyk model but are permitted under our new definition:

- **Key-compromise impersonation (KCI) attack** [7, 13]: the adversary reveals a long-term secret

key of a party and then impersonates others to this party.

- An adversary reveals the ephemeral secret key of a party and impersonates others to this party.
- Two honest parties execute matching sessions, and the adversary reveals the ephemeral secret keys of both of the parties and tries to learn the session key.
- Two honest parties execute matching sessions, and the adversary reveals the ephemeral secret key of one party, the long-term secret key of the other party and tries to learn the session key.
- Two honest parties execute matching sessions, while the adversary reveals the long-term keys of both of the parties prior to the execution of the session and tries to learn the session key.

### 2.3 State-reveal Queries

The Canetti-Krawczyk [10] model gives the adversary the power to reveal the state-specific information of the parties (Session-State Reveal query) without revealing the long-term secret key. State reveal queries are motivated by practical scenarios, such as if the state-specific secret information is stored in insecure memory or if the random-number generator of a party is corrupted. In both of these cases, as a result of the specific vulnerability, the adversary might be able to gain access to the ephemeral data.

As an example of a protocol which is secure against state-reveal queries, Canetti and Krawczyk present the SIG-DH protocol, depicted in Figure 2, which uses the Diffie-Hellman protocol and any digital signature scheme. Both parties are assumed to have secret keys for this signature scheme and know each other’s registered public keys and a session identifier  $sid$ . Denote by  $SIG_{\mathcal{A}}(M)$  a signature of  $M$  under  $\mathcal{A}$ ’s private key. Canetti and Krawczyk [10] prove that SIG-DH is secure

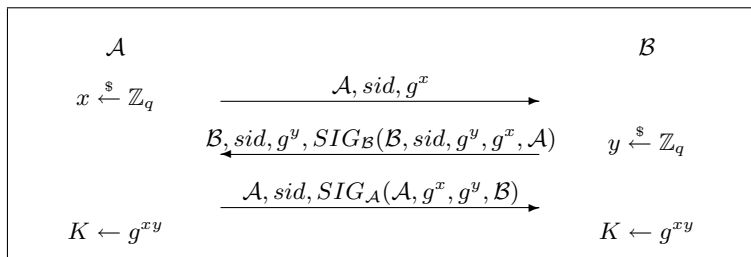


Figure 2: SIG-DH authenticated key-exchange protocol [10]

against an adversary who can reveal ephemeral secret keys ( $x$  or  $y$  from Figure 2) of the parties.

Note that the ephemeral secret keys  $x$  or  $y$  are not the only session-specific secret information used by the parties — they also use secret random coins in the signature generation. We observe that if the adversary reveals these random coins, it can break the security of the protocol. Specifically, for certain signature schemes such as Elgamal[11], DSA[18], GQ[12] or Schnorr[20], by revealing random coins used in the computation of any signature by any party the adversary can compute the long-term secret key of that party. Thus, the adversary would be able to impersonate the

honest party indefinitely. A similar deficiency in the encryption-based MT-authenticator of Bellare-Canetti-Krawczyk [1] was pointed out by Canetti and Krawczyk [10] and also discussed by Choo et al. [8].

In the security analysis of SIG-DH, note that the adversary is only allowed to reveal the ephemeral exponent but not random coins used in the signature generation. The reason for this somewhat surprising limitation lies in the ambiguity of the definition of a state reveal query: the Canetti-Krawczyk definition [10] leaves it up to a protocol to specify which information can be revealed by a state reveal query. As an extreme example, note that one can specify that the **Session-State Reveal** query does not reveal any information at all. Then any protocol secure against a weak adversary (who makes no state reveal queries) will also be secure against a strong adversary (who can make state reveal queries).

### 3 Definitions

#### 3.1 Motivation for Our Security Definition

We begin the presentation of the extended security model by clarifying the notion of what constitutes the “ephemeral secret key”. Informally, we require that the ephemeral secret key contain all the session-specific information used by a party in an AKE session. That is, all computations by a party must depend deterministically on that party’s ephemeral key, long-term secret key, and communication received from the other party. We observe that any AKE protocol can conform to this specification by setting all random coins used by a party in an AKE session as the ephemeral secret key of that session.

We modify the Canetti-Krawczyk model in the definition of adversarial power and in the notion of cleanness of the test session. Specifically, we replace the **Session-State Reveal** query with **Ephemeral Key Reveal** query which reveals the ephemeral secret key of the party. Additionally, we give the adversary the power to reveal a long-term secret key without corrupting the party by making a **Long-Term Key Reveal** query. We remove the **Corrupt** query as it is no longer necessary: the adversary can achieve the same result as the **Corrupt** query by revealing all the secret information of the party through **Long-Term Key Reveal**, **Ephemeral Key Reveal** and **Reveal** queries and by computing everything on behalf of that party.

We also modify the definition of a “clean session” by allowing the adversary to reveal the maximum possible amount of data. We disallow only those corruptions which allow the adversary to break any AKE protocol.

As above, we classify the test sessions as either “passive” or “active” depending on whether the adversary is able to cancel or modify the information sent between two honest participants. Formally, passive sessions are those where the matching session was completed at some point in the experiment (possibly after the test session), and active sessions are those where no matching session was completed at any time in the experiment.

For passive sessions we allow the adversary to reveal any subset of the four secret keys (each party’s ephemeral and long-term secret keys) which does not contain both the ephemeral and long-term secret keys of a single party. Note that the knowledge of both the ephemeral and long-term keys of one of the parties allows the adversary to compute the session key for any AKE protocol.

For active sessions the communication sent by the peer might be corrupted and thus we cannot define the ephemeral key of the peer. In this case we only allow the adversary to reveal either

the ephemeral or the long-term secret key of the owner, as revealing both keys would trivially compromise the protocol. Note that we cannot allow the adversary to reveal the long-term secret of the peer (even after the test session is completed), since Krawczyk [13] shows that in this case one can break any AKE protocol. (This is the same attack which shows the impossibility of the full perfect forward secrecy requirement.)

### 3.2 Security Experiment for Extended Canetti-Krawczyk

Assume that the identities of the parties are binary strings (they can be derived from the actual names of the parties). We will use letters  $\mathcal{A}$ ,  $\mathcal{B}$ ,  $\mathcal{C}$ ,  $\dots$ , both for referring to the parties and for their identities. The adversary is given the power to select each party's identity (the binary string) if it so chooses.

There are a number of honest parties which are connected to the certificate authority,  $\mathcal{CA}$ , and to the adversary,  $\mathcal{M}$ . That is, the communication between the parties is fully controlled by  $\mathcal{M}$ .  $\mathcal{M}$  is also connected to the certificate authority and can register fictitious parties. The adversary plays a central role in the experiment and is responsible for activating all other parties.

We call a particular instantiation of the AKE protocol executed by one of the parties an *AKE session*. Since all communication is controlled by the adversary, a party can never know if the second party actually exists and if the communication it receives was computed by an honest party or by the adversary. Legitimate execution of an AKE protocol by two parties  $\mathcal{A}$  and  $\mathcal{B}$  consists of two AKE sessions, matching sessions executed by  $\mathcal{A}$  and by  $\mathcal{B}$  respectively. Note that an instantiation of the AKE protocol is different depending on whether the executor is the initiator or the responder.

We do not assume the existence of explicit session identifiers. Instead, we define a session identifier to consist of the identities of the 2 participants and the information they exchanged. Specifically, a session identifier

$$sid = (role, ID, ID^*, comm_1, \dots, comm_n),$$

where  $ID \in \{0,1\}^*$  is the identity of the party executing the session,  $role \in \{I, R\}$  is its role (initiator/responder) in the protocol,  $ID^*$  is the identity of the other party and  $comm_i \in \{0,1\}^*$  is the  $i$ -th communication sent by the parties.

A party computes a communication  $comm_i$  as a function of its own ephemeral and long-term secret keys, its partner's public key and previous messages exchanged. Once a party receives all the communications, it computes a session key as a function of its own ephemeral and long-term secret keys, its partner's public key, and all communications, and completes the session.

The experiment proceeds as follows. Initially  $\mathcal{M}$  selects the identities of all honest parties (which can be arbitrary distinct binary strings) and honest parties generate and register their public keys with the  $\mathcal{CA}$ . The adversary can register arbitrary public keys (even the same as those of some honest parties) on behalf of adversary-controlled parties. Then the adversary makes any sequence of the following queries:

- $\text{Send}(\mathcal{A}, \mathcal{B}, comm)$ . Sends a message  $comm$  to  $\mathcal{A}$  on behalf of  $\mathcal{B}$ . Returns  $\mathcal{A}$ 's response to this message. This query allows  $\mathcal{M}$  to order  $\mathcal{A}$  to start an AKE session with  $\mathcal{B}$  and to provide communications from  $\mathcal{B}$  to  $\mathcal{A}$ .
- $\text{Long-Term Key Reveal}(\mathcal{A})$ . Reveals a long-term key of a party  $\mathcal{A}$ .

- **Ephemeral Key Reveal**( $sid$ ). Reveals an ephemeral key of a session  $sid$  (possibly incomplete).
- **Reveal**( $sid$ ). Reveals a session key of a completed session  $sid$ .

Eventually (at any time in the experiment),  $\mathcal{M}$  selects a completed session  $sid$ , makes a query **Test**( $sid$ ) and is given a challenge value  $C$ .  $\mathcal{M}$  continues the experiment after the **Test** query. The experiment terminates as soon as  $\mathcal{M}$  makes the **Guess**( $b'$ ) query. The experiment answers the adversary's queries as follows:

- **Test**( $sid$ ) // can be made only once.  
Pick  $b \xleftarrow{\$} \{0, 1\}$ . If  $b = 1$ , obtain  $C \leftarrow \text{Reveal}(sid)$ ; otherwise pick  $C \xleftarrow{\$} \{0, 1\}^\lambda$ . Return  $C$ .
- **Guess**( $b'$ ) //  $\mathcal{M}$  terminates after making this query.  
If  $b' = b$ , return 1, otherwise return 0.

An adversary  $\mathcal{M}$  *wins* the experiment if the selected test session is *clean* and if he guesses the challenge correctly (that is, if the **Guess** query returns 1).

We now define what it means for a test session to be *clean*. Let  $sid$  be an AKE session completed by a party  $\mathcal{A}$  with some other party  $\mathcal{B}$ , and denote by  $sid^*$  the matching session to  $sid$ , supposedly executed by  $\mathcal{B}$  ( $sid^*$  may not exist in the experiment). Denote by  $sk_{\mathcal{A}}$  and  $sk_{\mathcal{B}}$  long-term secret keys of  $\mathcal{A}$  and  $\mathcal{B}$ . Denote by  $esk_{\mathcal{A}}$  and  $esk_{\mathcal{B}}$  ephemeral secret keys generated by  $\mathcal{A}$  and  $\mathcal{B}$  in  $sid$  and  $sid^*$  (the latter is defined only if  $sid^*$  exists). We say that an AKE session  $sid$  is *not clean* if an adversary can trivially compute the session key. That is, a session  $sid$  is not clean if any of the following conditions hold:

- $\mathcal{A}$  or  $\mathcal{B}$  is an adversary-controlled party.
- $\mathcal{M}$  reveals the session key of  $sid$  or  $sid^*$  (if the latter exists).
- Session  $sid^*$  exists and  $\mathcal{M}$  reveals either both  $sk_{\mathcal{A}}$  and  $esk_{\mathcal{A}}$ , or both  $sk_{\mathcal{B}}$  and  $esk_{\mathcal{B}}$ .
- Session  $sid^*$  doesn't exist and  $\mathcal{M}$  reveals either  $sk_{\mathcal{B}}$  or both  $sk_{\mathcal{A}}$  and  $esk_{\mathcal{A}}$ .

A session  $sid$  is clean if *none* of these conditions hold. We remark that the cleanness of the test session can be identified only after the experiment is completed: the third and fourth conditions above can only be determined in the end of the experiment. That is, the adversary wins the experiment if he correctly guesses the challenge for the test session and this session remains clean until the end of the experiment.

**Definition 1 (Extended Canetti-Krawczyk security)** *The advantage of the adversary  $\mathcal{M}$  in the AKE experiment with AKE protocol  $\Pi$  is defined as*

$$\text{Adv}_{\Pi}^{\text{AKE}}(\mathcal{M}) = \Pr[\mathcal{M} \text{ wins}] - \frac{1}{2}.$$

*We say that an AKE protocol is secure (in the extended Canetti-Krawczyk model) if no efficient adversary  $\mathcal{M}$  has more than a negligible advantage in winning the above experiment.*



## 4 NAXOS AKE Protocol

### 4.1 Assumptions

All the arithmetic in this section is assumed to be in a mathematical group  $G$  of known prime order  $q$ . We denote by  $g$  a generator of  $G$  and write the group operation multiplicatively.

The discrete logarithm function  $DLOG(\cdot)$  in  $G$  takes input an element  $a \in G$  and returns  $x \in \mathbb{Z}_q$  such that  $a = g^x$ . The computational Diffie-Hellman (CDH) function  $CDH(\cdot, \cdot)$  takes as input a tuple of elements  $(a, b) \in G^2$  and returns  $g^{DLOG(a) \cdot DLOG(b)}$ . The Decisional Diffie-Hellman (DDH) function  $DDH(\cdot, \cdot, \cdot)$  takes as input a triple of elements  $(a, b, c) \in G^3$  and returns 1 if  $c = CDH(a, b)$  and 0 otherwise.

The *advantage* of an algorithm  $\mathcal{M}$  in solving the Discrete Logarithm problem,  $\mathbf{Adv}^{DLOG}(\mathcal{M})$ , is the probability that, given  $a \xleftarrow{\$} G$ ,  $\mathcal{M}$  correctly returns  $DLOG(a)$ . Similarly, the advantage of an algorithm  $\mathcal{M}$  in solving the Gap Diffie-Hellman (GDH) problem,  $\mathbf{Adv}^{GDH}(\mathcal{M})$ , is the probability that, given as input  $(a, b) \xleftarrow{\$} G^2$  and oracle access to  $DDH(\cdot, \cdot, \cdot)$ ,  $\mathcal{M}$  correctly outputs  $CDH(a, b)$ . We say that  $G$  satisfies the GDH assumption if no feasible adversary can solve the GDH problem with non-negligible probability. The GDH assumption was introduced by Okamoto and Pointcheval [19] and is now a standard cryptographic assumption used to establish the security of many protocols.

Let  $G'$  be another group of order  $q$ . A function  $e : G \times G \rightarrow G'$  is a bilinear pairing if it is non-degenerate and if for any pair  $g^a, g^b \in G$ ,  $e(g^a, g^b) = e(g, g)^{ab}$ . The Pairing Diffie-Hellman (PDH) problem recently introduced by Mityagin and Lauter [16] is to solve the CDH problem when given access to the pairing oracle  $e$ . The advantage  $\mathbf{Adv}^{PDH}(\mathcal{M})$  of an algorithm  $\mathcal{M}$  in solving the PDH problem is the probability that  $\mathcal{M}$ , given  $(a, b) \xleftarrow{\$} G^2$  and a pairing oracle  $e$ , computes  $CDH(a, b)$ . We say that  $G$  satisfies the PDH assumption if no feasible adversary solves the PDH problem with non-negligible probability. In groups which have a bilinear pairing, the PDH problem is equivalent to the original CDH problem, although one can also consider the PDH problem in groups where no efficient pairing operation is known. We find the Pairing Diffie-Hellman assumption to be as justified as the GDH assumption since the only known way to compute DDH in groups where CDH is hard is via a pairing function.

### 4.2 Protocol Description

The NAXOS AKE protocol uses a mathematical group  $G$  and two hash functions,  $H_1 : \{0, 1\}^* \rightarrow \mathbb{Z}_q$  and  $H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$  (for some constant  $\lambda$ ). A long-term secret key of a party  $\mathcal{A}$  is an exponent  $sk_{\mathcal{A}} \in \mathbb{Z}_q$ , and the corresponding long-term public key of  $\mathcal{A}$  is the power  $pk_{\mathcal{A}} = g^{sk_{\mathcal{A}}} \in G$ . In the following description of an AKE session of NAXOS executed between the parties  $\mathcal{A}$  and  $\mathcal{B}$  we assume that each party knows the other's public key and that public keys are in the group  $G$ .

The session execution proceeds as follows. The parties pick ephemeral secret keys  $esk_{\mathcal{A}}$  and  $esk_{\mathcal{B}}$  at random from  $\{0, 1\}^\lambda$ . Then the parties exchange values  $g^{H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})}$  and  $g^{H_1(esk_{\mathcal{B}}, sk_{\mathcal{B}})}$ , check if received values are in the group  $G$  and only compute the session keys if the check succeeds. The session key  $K \in \{0, 1\}^\lambda$  is computed as

$$H_2(g^{H_1(esk_{\mathcal{B}}, sk_{\mathcal{B}})sk_{\mathcal{A}}}, g^{H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})sk_{\mathcal{B}}}, g^{H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})H_1(esk_{\mathcal{B}}, sk_{\mathcal{B}})}, \mathcal{A}, \mathcal{B}).$$

The last two components in the hash are the identities of  $\mathcal{A}$  and  $\mathcal{B}$ , which we assume to be binary strings. Figure 3 depicts the protocol.




Figure 3: NAXOS AKE Protocol.

**Theorem 1** *NAXOS satisfies Extended Canetti-Krawczyk security if  $H_1$  and  $H_2$  are modeled by independent random oracles.*

For any AKE adversary  $\mathcal{M}$  against NAXOS that runs in time at most  $t$ , involves at most  $n$  honest parties and activates at most  $k$  sessions, we show that there exists a GDH solver  $\mathcal{S}$ , a PDH solver  $\mathcal{R}$  and a DLOG solver  $\mathcal{T}$  such that

$$\mathbf{Adv}^{GDH}(\mathcal{S}) = \mathbf{Adv}^{PDH}(\mathcal{R}) \geq \frac{1}{2} \min \left\{ \frac{2}{k^2}, \frac{1}{nk} \right\} \cdot \mathbf{Adv}_{NAXOS}^{AKE}(\mathcal{M}) - \mathbf{Adv}^{DLOG}(\mathcal{T}) - O\left(\frac{k^2}{2^\lambda}\right),$$

where  $\mathcal{S}$  runs in time  $O(tk)$ ,  $\mathcal{R}$  runs in time  $O(t \log t)$  and  $\mathcal{T}$  runs in time  $O(t)$ .

The outline of the security proof of Theorem 1 is given in Section 4.3.

### 4.3 Security Proof for NAXOS

Let  $\mathcal{A}$  be any AKE adversary against NAXOS. We start by observing that since the session key of the test session is computed as  $K = H_2(\sigma)$  for some 5-tuple  $\sigma$ , the adversary  $\mathcal{M}$  has only two ways to distinguish  $K$  from a random string:

1. Forging attack. At some point  $\mathcal{M}$  queries  $H_2$  on the same 5-tuple  $\sigma$ .
2. Key-replication attack.  $\mathcal{M}$  succeeds in forcing the establishment of another session that has the same session key as the test session.

If random oracles produce no collisions, the key-replication attack is impossible as equality of session keys implies equality of the corresponding 5-tuples (which are hashed to produce session keys). In turn, distinct AKE sessions must have distinct 5-tuples. Therefore, if random oracles produce no collisions (collisions happen with probability  $O(k^2/2^\lambda)$ ),  $\mathcal{M}$  must perform a forging attack. Next we show that if  $\mathcal{M}$  can mount a successful forging attack, then we can construct a Gap Diffie-Hellman solver  $\mathcal{S}$  which uses  $\mathcal{M}$  as a subroutine. Most of the remaining proof is devoted to the construction of  $\mathcal{S}$ .

$\mathcal{S}$  takes as input a GDH challenge  $(X_0, Y_0)$ . Then  $\mathcal{S}$  executes the Extended Canetti-Krawczyk (ECK) experiment with  $\mathcal{M}$  the adversary against the NAXOS protocol, and modifies the data returned by the honest parties in such a way that if  $\mathcal{M}$  breaks the ECK security of NAXOS, then  $\mathcal{S}$  can reveal the solution to the GDH problem from  $\mathcal{M}$ .

## Matching Session Exists

Assume that  $\mathcal{M}$  always selects a test session for which the matching session exists. Then  $\mathcal{S}$  modifies the experiment as follows.  $\mathcal{S}$  selects at random matching sessions executed by some honest parties  $\mathcal{A}$  and  $\mathcal{B}$  (in fact,  $\mathcal{S}$  selects two sessions at random and continues only if they are matching). Denote by  $comm_A$  and  $comm_B$  the communications sent by the respective parties in these matching sessions. When either of these sessions is activated,  $\mathcal{S}$  does not follow the protocol. Instead,  $\mathcal{S}$  generates  $esk_A$  and  $esk_B$  normally but sets  $comm_A \leftarrow X_0$  (in place of  $g^{H_1(sk_A, esk_A)}$ ) and  $comm_B \leftarrow Y_0$  (in place of  $g^{H_1(sk_B, esk_B)}$ ).

With probability  $1/k^2$   $\mathcal{M}$  picks one of the selected sessions as the test session and another as its matching session. We claim that if  $\mathcal{M}$  wins in the forging attack,  $\mathcal{S}$  can solve the CDH challenge. Indeed, the supposed session key for the selected session is  $H_2(\sigma)$ , where the 5-tuple  $\sigma$  includes the value  $CDH(X_0, Y_0)$ . To win,  $\mathcal{M}$  must have queried  $\sigma$  to the random oracle  $H_2$ .

If the selected session is indeed the test session,  $\mathcal{M}$  is allowed to reveal  $sk_A, sk_B, esk_A$  and  $esk_B$  but it is not allowed to reveal both  $(sk_A, esk_A)$  or both  $(sk_B, esk_B)$ . We observe that in this case, the only way that  $\mathcal{M}$  can distinguish this simulated ECK experiment from a true ECK experiment is if  $\mathcal{M}$  queries  $(sk_A, esk_A)$  or  $(sk_B, esk_B)$  to  $H_1$  (this way,  $\mathcal{M}$  will find out that  $comm_A$  and  $comm_B$  were not computed correctly). However,  $\mathcal{M}$  cannot do this unless he computes the discrete logarithm of either  $g^{sk_A}$  or  $g^{sk_B}$ . This corresponds to a hypothetical discrete logarithm adversary  $\mathcal{T}$  in the security statement.

## No Matching Session

Assume now that  $\mathcal{M}$  always selects a test session such that the matching session doesn't exist. In this case  $\mathcal{S}$  modifies the experiment as follows.  $\mathcal{S}$  selects a random party  $\mathcal{B}$  and sets  $pk_B \leftarrow X_0$ . Note that  $\mathcal{S}$  doesn't know the secret key corresponding to this public key and thus it cannot properly simulate ECK sessions executed by  $\mathcal{B}$ .  $\mathcal{S}$  handles ECK sessions executed by  $\mathcal{B}$  as follows (assume that  $\mathcal{B}$  is the initiator).  $\mathcal{S}$  randomly selects  $esk_B$ , picks  $h$  at random from  $\mathbb{Z}_q$  and sets  $comm_B = g^h$  instead of  $g^{H_1(esk_B, DLOG(X_0))}$ .  $\mathcal{S}$  sets a session key  $K$  (which is supposed to be  $H_2(CDH(X_0, comm_C), pk_C^h, comm_C^h, \mathcal{B}, \mathcal{C}))$  to be a random value. Note that  $\mathcal{S}$  can handle session key and ephemeral secret key reveals by revealing  $K$  and  $esk_B$ , but cannot handle long-term secret key reveals.

Note that if  $\mathcal{C}$  is a fictitious party,  $\mathcal{M}$  can compute the session key on its own, reveal  $K$  and detect that it is fake. To address this issue,  $\mathcal{S}$  watches  $\mathcal{M}$ 's random oracle queries and if  $\mathcal{M}$  queries  $(Z, pk_C^h, comm_C^h, \mathcal{B}, \mathcal{C})$  to  $H_2$  (for some  $Z \in G$ ),  $\mathcal{S}$  checks if  $DDH(X_0, comm_C, Z) = 1$  and if yes, replies with the key  $K$ . Similarly, on the computation of  $K$ ,  $\mathcal{S}$  checks if  $K$  should equal any previous response from the random oracle.

$\mathcal{M}$  cannot detect that it is in the simulated ECK experiment unless it either queries  $(esk_B, DLOG(X_0))$  to  $H_1$  (by doing this,  $\mathcal{M}$  can detect that  $comm_B$  is computed incorrectly) or tries to reveal a long-term secret key of  $\mathcal{B}$ . The first event will reveal  $DLOG(X_0)$  and will clearly allow  $\mathcal{S}$  to solve the CDH problem. As we note below, the second event is impossible if  $\mathcal{S}$  correctly guesses the test session.

Now  $\mathcal{S}$  also randomly selects an ECK session in which  $\mathcal{B}$  is the peer. Denote the owner of this session by  $\mathcal{A}$ . When the selected session is activated,  $\mathcal{S}$  follows the protocol only partially:  $\mathcal{S}$  generates  $esk_A$  normally but sets  $comm_A \leftarrow Y_0$  (in place of  $g^{H_1(sk_A, esk_A)}$ ).

With probability at least  $1/nk$  ( $1/n$  to pick the correct party  $\mathcal{B}$  and  $1/k$  to pick the correct session),  $\mathcal{M}$  picks the selected session as the test session, and if it wins, it solves the CDH problem. The supposed session key for the selected session is  $H_2(\sigma)$ , where the 5-tuple  $\sigma$  includes the value  $CDH(X_0, Y_0)$ . To win,  $\mathcal{M}$  must have queried  $\sigma$  to the random oracle  $H_2$ .

If the selected session is indeed the test session,  $\mathcal{M}$  is not allowed to reveal both  $sk_{\mathcal{A}}$  and  $esk_{\mathcal{A}}$  and is not allowed to corrupt  $\mathcal{B}$ . In this case, the only way that  $\mathcal{M}$  can distinguish this simulated ECK experiment from a true ECK experiment is if  $\mathcal{M}$  queries  $(sk_{\mathcal{A}}, esk_{\mathcal{A}})$  to  $H_1$ . However,  $\mathcal{M}$  cannot do this unless he computes the discrete logarithm of  $g^{sk_{\mathcal{A}}}$ .

## Analysis

We observe that the running time of  $\mathcal{S}$  is  $O(kt)$ . For each session key computation done by  $\mathcal{B}$  (let  $Y$  be the incoming communication in that session) the solver  $\mathcal{S}$  has to go over all previous  $H_2$  queries and for each  $H_2$  query of the form  $(\dots, Z, \dots)$  check if  $DDH(X_0, Y, Z) = 1$ . Similarly, on each DDH query of the form  $(\dots, Z, \dots)$   $\mathcal{S}$  has to go over all previous session key computations done by  $\mathcal{B}$  and for each such computation (let  $Y$  the incoming communication in that session)  $\mathcal{S}$  checks if  $DDH(X_0, Y, Z)$ . Since  $\mathcal{M}$  can activate at most  $k$  sessions and make at most  $t$   $H_2$  queries, total running time is  $O(tk)$ .

The running time of the solver can be improved if the solver has access to the pairing oracle instead of to the DDH oracle. We construct the PDH solver  $\mathcal{R}$  in the same way as  $\mathcal{S}$  with the only difference being that  $\mathcal{R}$  must also handle the checks discussed above. Note that  $DDH(X_0, Y, Z) = 1$  if and only if  $e(Z, g) = e(X_0, Y)$ . Therefore  $\mathcal{R}$  can store corresponding values  $e(Z, g)$  in a balanced binary tree and on each session executed by  $\mathcal{B}$  check for  $X_0, Y$  by computing  $e(X_0, Y)$  and searching for this value in the binary tree (which can be done in  $\log t$  steps). Therefore,  $\mathcal{R}$  has the same advantage as  $\mathcal{S}$  and runs in time  $O(t \log t)$ .

## References

- [1] M. Bellare, R. Canetti and H. Krawczyk, *A modular approach to the design and analysis of authentication and key exchange protocols*, STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing, ACM Press, 1998
- [2] M. Bellare, A. Palacio, *The Knowledge-of-Exponent Assumptions and 3-Round Zero-Knowledge Protocols*, Advances in Cryptology — CRYPTO '04, pp. 273–289, Springer-Verlag, 2004
- [3] M. Bellare, D. Pointcheval, P. Rogaway, *Authenticated Key Exchange Secure Against Dictionary Attacks*, Advances in Cryptology — Eurocrypt '00, pp. 139–155, Springer-Verlag, 2000
- [4] M. Bellare and P. Rogaway, *Entity Authentication and Key Distribution*, Advances in Cryptology — CRYPTO '93, pp. 110–125, Springer-Verlag, 1993
- [5] M. Bellare and P. Rogaway, *Random Oracles are Practical: A Paradigm for Designing Efficient Protocols*, ACM Conference on Computer and Communications Security 1993, pp. 62–73

- [6] M. Bellare and P. Rogaway, *Provably Secure Session Key Distribution: the Three Party Case*, STOC '95: Proc. 27th Annual Symposium on the Theory of Computing, ACM, 1995
- [7] S. Blake-Wilson, D. Johnson, and A. Menezes, *Key Agreement Protocols and their Security Analysis*, 6th IMA International Conference on Cryptography and Coding, LNCS 1355, pp. 30-45, Springer-Verlag, 1997
- [8] K.-K. R. Choo, C. Boyd and Y. Hitchcock, *Errors in Computational Complexity Proofs for Protocols*, Advances in Cryptology — Asiacrypt '05, LNCS 3788, pp. 624–643, Springer-Verlag, 2005
- [9] K.-K. R. Choo, C. Boyd and Y. Hitchcock, *Examining Indistinguishability-Based Proof Models for Key Establishment Protocols*, Advances in Cryptology — Asiacrypt '05, Springer-Verlag, 2005
- [10] R. Canetti and H. Krawczyk, *Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels*, Advances in Cryptology — EUROCRYPT '01, pp. 453–474, Springer-Verlag, 2001
- [11] T. Elgamal, *A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms*, IEEE Transactions on Information Theory, 1985
- [12] L. Guillou and J. J. Quisquater, *A paradoxical identity-based signature scheme resulting from zero-knowledge*, Advances in Cryptology — CRYPTO 88, LNCS 403, Springer-Verlag, 1988
- [13] H. Krawczyk, *HMQR: A High-Performance Secure Diffie-Hellman Protocol*, Advances in Cryptology — CRYPTO '05, LNCS 3621, pp. 546–566, Springer-Verlag, 2005
- [14] I. R. Jeong, J. Katz, D. H. Lee, *One-Round Protocols for Two-Party Authenticated Key Exchange*, ACNS'04, 2004
- [15] C. Kudla and K. G. Paterson, *Modular Security Proofs for Key Agreement Protocols*, Advances in Cryptology — ASIACRYPT '05, pp. 549–565, Springer-Verlag, 2005
- [16] K. Lauter and A. Mityagin, *Security Analysis of KEA Authenticated Key Exchange*, IACR Eprint archive, <http://eprint.iacr.org/2005/265>, 2005
- [17] A. Menezes, *Another look at HMQR*, IACR Eprint archive, <http://eprint.iacr.org/2005/205>, 2005
- [18] Federal Information Processing Standards Publication (FIPS PUB) 186, *Digital Signature Standard*, May 19, 1994.
- [19] T. Okamoto and D. Pointcheval, *The Gap Problems: A New Class of Problems for the Security of Cryptographic Schemes*, Public Key Cryptology — PKC '01, LNCS 1992, pp. 104–118, Springer-Verlag, 2001
- [20] C. P. Schnorr, *Efficient signature generation by smart cards*, Journal of Cryptology 4, 3, pp. 161-174, 1991.