

Authenticated Key Exchange Secure under the Computational Diffie-Hellman Assumption

Jooyoung Lee and Je Hong Park

The Attached Institute of Electronics and Telecommunications Research Institute
Yuseong-gu, Daejeon, Korea 305-390
{jlee05, jhpark}@ensec.re.kr

Abstract. In this paper, we present a new authenticated key exchange(AKE) protocol and prove its security under the random oracle assumption and the computational Diffie-Hellman(CDH) assumption. In the extended Canetti-Krawczyk model, there has been no known AKE protocol based on the CDH assumption. Our protocol, called NAXOS+, is obtained by slightly modifying the NAXOS protocol proposed by LaMacchia, Lauter and Mityagin. We establish a formal security proof of NAXOS+ in the extended Canetti-Krawczyk model using as a main tool the trapdoor test presented by Cash, Kiltz and Shoup.

Key words: authenticated key exchange, trapdoor test, computational Diffie-Hellman assumption, extended Canetti-Krawczyk model

1 Introduction

Key exchange protocols are cryptographic protocols that allow two parties that share no secret information in common to establish a secret key via public communication. The most famous example is the classic Diffie-Hellman(DH) key exchange protocol that marked the birth of modern cryptography [9]. However, the original DH protocol did not provide authentication of the communicating parties, suffering from active attacks such as a man-in-the-middle attack. The quest for an authenticated key exchange(AKE) protocol, that not only allows parties to compute a shared key but also ensures authenticity of the parties, has resulted in a considerable amount of research both in security models and protocols [3–5, 10–15, 21, 23].

Security models. A well known security model for authenticated key exchange is the Canetti-Krawczyk(CK) model [5]. The CK model lies in a family of indistinguishability-based proof models including the Bellare-Rogaway model [4], Bellare-Pointcheval-Rogaway model [3], and their variants. Choo, Boyd and Hitchcock compared these various security models, and concluded that the CK model offers the strongest definition of security with some modification of the model [7]. However, all these models commonly based on a strong assumption that the adversary is not allowed to obtain certain secret information about the session that is being attacked. As a result, AKE protocols secure in the CK model still might be vulnerable to key-compromise impersonation attacks or the leakage of ephemeral private keys. Especially, the latter attack might be practical in a scenario where the random number generator used by a party is compromised.

Krawczyk later provided a stronger version of the CK model that not only guarantees resilience to the above attacks, but also achieves *weak* perfect forward secrecy [11]. However, this stronger version still does not include attacks such as revelation of both ephemeral private keys or both static private keys. Recently, LaMacchia, Lauter and Mityagin proposed an

elegant security model, called henceforth eCK model, that captures all these security properties [13]. Informally speaking, the only corruption powers that an adversary is not allowed for in the eCK model are those that would trivially break an authentication protocol. For this reason, the eCK model is currently regarded as the strongest security model.

AKE protocols and security assumptions. One of the most famous and standardized AKE protocols is the MQV protocol proposed by Law, Menezes, Qu, Solinas and Vanstone [15]. However, the construction of the MQV protocol was not based on a formal security proof. In [11], Krawczyk slightly modified the MQV protocol to obtain the HMQV protocol. He presented a formal security proof of the HMQV protocol in the stronger version of the CK model, under the random oracle assumption, the gap Diffie-Hellman(GDH) assumption and the knowledge of exponent assumption [2]. The NAXOS protocol is the first AKE protocol whose security is established in the eCK model [13]. One shortcoming of the protocol is that it requires 4 exponentiations per party compared to 2.5 exponentiations for MQV and HMQV. Motivated by this observation, Ustaoglu presented the CMQV protocol that achieves both efficiency and security [23]. The security proof of NAXOS and CMQV uses only the random oracle assumption and the GDH assumption, which are weaker than the security assumptions of HMQV.

Focusing on the attempts to weaken the security assumptions in the analysis of AKE protocols, Okamoto proposed an AKE protocol that is proven secure without the random oracle assumption [21]. However, the proof depends upon a rather strong assumption of the existence of π PRFs. The AKE protocol $\mathcal{TS3}$ proposed by Jeong, Katz and Lee [10] uses only the decisional Diffie-Hellman(DDH) assumption and the existence of secure message authentication codes, while we note that the security model is much weaker than the eCK model.

Contribution. We observe that AKE protocols that are proven secure in the eCK model commonly use the GDH assumption. The gap Diffie-Hellman problem, first introduced in [22], is to solve the computational Diffie-Hellman problem with a decision Diffie-Hellman oracle. Obviously, the GDH assumption is not weaker than the CDH assumption, while the equivalence of the two hardness assumptions still remains open. Motivated by this observation, we present a new authenticated key exchange protocol such that the only security assumptions are the random oracle assumption and the CDH assumption. We obtain our protocol, called NAXOS+, by slightly modifying the NAXOS protocol. Compared to the original NAXOS protocol, NAXOS+ requires only one more exponentiation per party. On the other hand, our protocol enjoys a simple and tight security reduction compared to those of HMQV and CMQV. The security proof is very similar to that of NAXOS. One main difference is that we implement the decisional Diffie-Hellman oracle using the trapdoor test, recently presented by Cash, Kiltz and Shoup [8].

Organization. In the next section, we present a formal description of the security assumptions and the eCK model. We describe the NAXOS+ protocol in Section 3, and prove its security in Section 4. In Section 5, we provide one-pass and three-pass variants of NAXOS+ as well as discussions on certain security aspects.

2 Preliminaries

2.1 Assumption

Let $\mathbb{G} = \langle g \rangle$ denote a multiplicative cyclic group of prime order q , generated by $g \in \mathbb{G}$. The discrete logarithm function $\text{DLOG}(\cdot)$ in \mathbb{G} takes as input an element $U \in \mathbb{G}$ and returns $u \in \mathbb{Z}_q$ such that $U = g^u$. The computational Diffie-Hellman function $\text{CDH}(\cdot, \cdot)$ takes as input a pair of elements $(U, V) \in \mathbb{G}^2$ and returns $g^{\text{DLOG}(U) \cdot \text{DLOG}(V)}$.

Let \mathbb{G}^* denote the set of non-identity elements in \mathbb{G} . The *advantage* of an algorithm \mathcal{S} in solving the discrete logarithm (DLOG) problem, $\text{Adv}^{\text{dlog}}(\mathcal{S})$, is defined to be the probability that, given $U \xleftarrow{\$} \mathbb{G}^*$, \mathcal{S} correctly returns $\text{DLOG}(U)$. Similarly, the advantage of an algorithm \mathcal{S} in solving the computational Diffie-Hellman (CDH) problem, $\text{Adv}^{\text{cdh}}(\mathcal{S})$, is the probability that, given as input $(U, V) \xleftarrow{\$} (\mathbb{G}^*)^2$, \mathcal{S} correctly returns $\text{CDH}(U, V)$. It is obvious that any DLOG solver can be used in solving the CDH problem. We say that \mathbb{G} satisfies the CDH assumption if no feasible adversary can solve the CDH problem with non-negligible probability.

2.2 Extended Canetti-Krawczyk Security Model

In this section, we describe the eCK model focusing on two pass AKE protocols in the random oracle model. We slightly modify the original model by introducing *explicit* session identifiers. We believe that the explicit identifiers would render a simpler description of AKE protocols, especially for those of three or more passes. We are using the same framework in [17]. For further details on the eCK model, we refer to [13, 23].

Parties. We fix a set \mathbf{P} of n parties, each of which is modeled as a probabilistic Turing machine that makes queries to a set $\mathbf{H} = (H_1, \dots, H_l)$ of independent random oracles. Each party $\mathcal{A} \in \mathbf{P}$ is represented by a distinct string of a fixed length, say μ . When $\mathcal{A} \in \mathbf{P}$ appears in a protocol flow or as an argument to a function, we mean the string which names the party. We assume that each party \mathcal{A} stores a static public/private key pair $(pk_{\mathcal{A}}, sk_{\mathcal{A}})$, together with a certificate that binds the public key to that party. However, we do not assume that the certification authority (CA) requires parties to prove possession of their static private keys.

An instance of the protocol executed by a party is called a *session*. When a party \mathcal{A} creates a new session, \mathcal{A} chooses a distinct index i from a set \mathbf{I} of indices, and assigns a session identifier $sid = (\mathcal{A}, i) \in \mathbf{SID}$ to the session, where we denote $\mathbf{SID} = \{0, 1\}^\mu \times \mathbf{I}$. Creation of a session is made via an incoming message that has one of the forms $(\mathcal{A}, \mathcal{B}) \in \{0, 1\}^{2\mu}$ and $(\mathcal{A}, (\mathcal{B}, j), Y) \in \{0, 1\}^\mu \times \mathbf{SID} \times \{0, 1\}^*$. When \mathcal{A} creates a session on a message of the form $(\mathcal{A}, \mathcal{B})$, \mathcal{A} is called the *initiator* of the session, otherwise the *responder* of the session. The behavior of a party \mathcal{A} is illustrated in Fig. 1. Note that each session sid is associated with the following three variables.

- esk_{sid} (sometimes denoted $esk_{\mathcal{A}}$). Stores an ephemeral private key of λ bits used in the computation of the session key for a constant λ .
- T_{sid} . Records the transcript of the session. The variable takes its value from the set $\mathbf{T} = (\{0, 1\}^\mu)^2 \times \{init, resp\} \times (\{-\} \cup \{0, 1\}^*)^2$, where “ $-$ ” is a special symbol not in $\{0, 1\}^*$ that means there is no message associated with the position. $T_{sid} = (\mathcal{A}, \mathcal{B}, role, X, Y)$ means that \mathcal{A} executes the session with \mathcal{B} as the second party. $role = init$ means that \mathcal{A}

is an initiator, and $role = resp$ a responder. X represents an outgoing message from \mathcal{A} , and Y an incoming message to \mathcal{A} . We call \mathcal{A} the *owner* of the session, and \mathcal{B} the *peer* of the session.

- K_{sid} . Stores the session key of λ bits.

We say that a session sid is *complete* if T_{sid} does not contain the symbol “–”. For a complete session sid with $T_{sid} = (\mathcal{A}, \mathcal{B}, role, X, Y)$, its *matching session* is defined as a session sid^* such that $T_{sid^*} = (\mathcal{B}, \mathcal{A}, role', Y, X)$ or $(\mathcal{B}, \mathcal{A}, role', Y, -)$, and $role \neq role'$. An *equivalent session* of sid is defined as a session sid^e such that $T_{sid^e} = T_{sid}$.

Remark 1. For most of existing AKE protocols, the probability that a session sid has an equivalent session that is not sid itself is negligible. Similarly, each session has more than one matching session only with negligible probability.

```

if  $msg\_in = (\mathcal{A}, \mathcal{B})$  then
  Create a session  $sid = (\mathcal{A}, i)$ , Generate  $esk_{sid}$ 
  Compute an ephemeral public key  $X$ ,  $T_{sid} \leftarrow (\mathcal{A}, \mathcal{B}, init, X, -)$ 
  return  $(\mathcal{B}, (\mathcal{A}, i), X)$ 
else if  $msg\_in = (\mathcal{A}, (\mathcal{B}, j), Y)$  then
  Create a session  $sid = (\mathcal{A}, i)$ , Generate  $esk_{sid}$ 
  Compute an ephemeral public key  $X$ ,  $T_{sid} \leftarrow (\mathcal{A}, \mathcal{B}, resp, X, Y)$ 
  Compute  $K_{sid}$ 
  return  $((\mathcal{B}, j), (\mathcal{A}, i), X)$ 
else if  $msg\_in = ((\mathcal{A}, i), (\mathcal{B}, j), Y)$  and  $T_{(\mathcal{A}, i)} = (\mathcal{A}, \mathcal{B}, init, X, -)$  for some  $X$ 
then
   $T_{(\mathcal{A}, i)} \leftarrow (\mathcal{A}, \mathcal{B}, init, X, Y)$ 
  Compute  $K_{sid}$ 

```

Fig. 1. Behavior of a party \mathcal{A} on an incoming message msg_in

Adversaries. An adversary \mathcal{M} is modeled as a probabilistic Turing machine that makes oracle queries to honest parties as well as the set \mathbf{H} of random oracles. It means that the adversary is able to control all communications between parties by sending arbitrary messages to a party on behalf of another party, obtaining responses from the party and/or making decisions about their delivery. In order to capture possible leakage of private information, we assume that an adversary is allowed to make the following additional oracle queries.

- **EstablishParty**(\mathcal{A}). Registers an additional party \mathcal{A} not in \mathbf{P} together with a static public key on the adversary’s own choice. The parties registered by **EstablishParty**(\cdot) queries are called *dishonest*, while the parties in \mathbf{P} called *honest*. We assume that the adversary is able to totally control dishonest parties.
- **StaticKeyReveal**(\mathcal{A}). Reveals a static private key of an honest party \mathcal{A} .
- **EphemeralKeyReveal**(sid). Reveals an ephemeral private key of a session sid owned by an honest party.
- **SessionKeyReveal**(sid). Reveals a session key of a complete session sid owned by an honest party.

Experiment. Initially, the adversary \mathcal{M} is given a set \mathbf{P} of honest parties. \mathcal{M} makes any sequence of the oracle queries described above. At any time in the experiment, \mathcal{M} selects a complete session $sid \in \mathbf{SID}$ owned by an honest party and makes a query $\text{Test}(sid)$. On the query, that is made only once during the experiment, the experiment answers $C \leftarrow \text{SessionKeyReveal}(sid)$ or $C \xleftarrow{\$} \{0, 1\}^\lambda$ with the same probability. At the end of the experiment, \mathcal{M} guesses whether the challenge C is random or not. If the test session is *clean* and if \mathcal{M} makes a correct guess, then we say that \mathcal{M} *wins* the experiment.

Definition 1. Let sid be a complete session with $T_{sid} = (\mathcal{A}, \mathcal{B}, \text{role}, X, Y)$, and let sid^* and sid^e denote a matching session and an equivalent session of sid , respectively. sid is called *clean* if \mathcal{A} and \mathcal{B} are honest parties, and if none of the following conditions hold.

1. \mathcal{M} makes a $\text{SessionKeyReveal}(sid^e)$ query,
2. \mathcal{M} makes a $\text{SessionKeyReveal}(sid^*)$ query, if sid^* exists,
3. \mathcal{M} makes $\text{StaticKeyReveal}(\mathcal{A})$ and $\text{EphemeralKeyReveal}(sid^e)$ queries,
4. \mathcal{M} makes $\text{StaticKeyReveal}(\mathcal{B})$ and $\text{EphemeralKeyReveal}(sid^*)$ queries, if sid^* exists,
5. \mathcal{M} makes a $\text{StaticKeyReveal}(\mathcal{B})$ query, if sid^* doesn't exist.

Definition 2 (eCK security [13]). The advantage of the adversary \mathcal{M} in the AKE experiment with AKE protocol Π 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 eCK model if matching sessions compute the same session keys and no efficient adversary \mathcal{M} has more than a negligible advantage in winning the above experiment.

2.3 Trapdoor test

Since we use the *trapdoor test* as a basic tool for our security proof, we state the theorem without proof. For further details on the proof and its applications, we refer to [8].

Theorem 1 (Trapdoor test [8]). Let \mathbb{G} be a cyclic group of prime order q , generated by $g \in \mathbb{G}$. Suppose U_1, r, s are independent random variables, where U_1 takes values in \mathbb{G} , and each of r, s is uniformly distributed over \mathbb{Z}_q , and define the random variable $U_2 = g^s / U_1^r$. Further, suppose that Y, Z_1, Z_2 are random variables taking values in \mathbb{G} , each of which is independent of r . Then we have

1. U_2 is uniformly distributed over \mathbb{G} ,
2. U_1 and U_2 are independent,
3. the probability that the truth value of

$$Z_1^r Z_2 = Y^s$$

does not agree with the truth value of

$$(\text{CDH}(Y, U_1) = Z_1) \wedge (\text{CDH}(Y, U_2) = Z_2)$$

is at most $1/q$.

3 NAXOS+ AKE protocol

The NAXOS+ AKE protocol uses a group $\mathbb{G} = \langle g \rangle$ of prime order q such that the CDH assumption holds, and two hash functions $H_1 : \{0, 1\}^\lambda \times \mathbb{Z}_q^* \rightarrow \mathbb{Z}_q^*$ and $H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, where $\mathbb{Z}_q^* = \mathbb{Z}_q \setminus \{0\}$, λ is a constant such that $q = O(2^\lambda)$, and H_1 and H_2 are modeled as independent random oracles. For the simplicity of description, we assume there is a natural embedding from \mathbb{G} to $\{0, 1\}^*$. Using the framework of an AKE protocol illustrated in Section 2.2, the NAXOS+ AKE protocol is described as follows.

Static public/private key pair. For each party \mathcal{A} , a static private key $sk_{\mathcal{A}} \in \mathbb{Z}_q^*$ is assigned uniformly at random. The corresponding public key is defined as $pk_{\mathcal{A}} = g^{sk_{\mathcal{A}}} \in \mathbb{G}^*$. When a party registers its static public key to the CA, the CA checks if the public key is contained in \mathbb{G}^* , and issues the associated certificate. We assume that each party learns the other’s public key via the certificate whenever a new session is initiated.

Ephemeral public/private key pair. For each session sid , a party \mathcal{A} generates an ephemeral private key $esk_{sid} \in \{0, 1\}^\lambda$ uniformly at random. The corresponding public key is defined as $X = g^{H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})} \in \mathbb{G}$. We assume that each party checks if the other’s ephemeral public key is contained in \mathbb{G}^* on receipt of the key. Note that the value $H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})$ is not regarded as an ephemeral key, since $H_1(esk_{\mathcal{A}}, sk_{\mathcal{A}})$ can be computed from $esk_{\mathcal{A}}$ and $sk_{\mathcal{A}}$ whenever needed.

Session key. In the end of a session sid such that $T_{sid} = (\mathcal{A}, \mathcal{B}, role, X, Y)$, its session key is defined as $K_{sid} = H_2(\sigma)$, where

$$\sigma = (\text{CDH}(pk_{\mathcal{A}}, pk_{\mathcal{B}}), \text{CDH}(pk_{\mathcal{A}}, Y), \text{CDH}(X, pk_{\mathcal{B}}), \text{CDH}(X, Y), \mathcal{A}, \mathcal{B}),$$

if $role = init$, and

$$\sigma = (\text{CDH}(pk_{\mathcal{A}}, pk_{\mathcal{B}}), \text{CDH}(X, pk_{\mathcal{B}}), \text{CDH}(pk_{\mathcal{A}}, Y), \text{CDH}(X, Y), \mathcal{B}, \mathcal{A}),$$

if $role = resp$. Fig. 2 depicts the protocol in a simplified form.

Remark 2. The “public key validation” prevents potential leakage of private information as shown in invalid-curve attacks [1] and small subgroup attacks [16, 19]. Especially, Menezes and Ustaoglu demonstrates the importance of public key validation by investigating the effects of omitting public-key validation in HMQV and MQV [18, 19]. As long as the underlying group is cryptographically strong, public key validation provides some assurance that computations involving its static and ephemeral private keys do not reveal any information about the key itself [18, 19, 23].

4 Security of NAXOS+

In this section, we present a formal security proof of NAXOS+ protocol under the CDH assumption and the random oracle assumption.

Let \mathcal{M} be an AKE adversary against NAXOS+ who runs in time t and involves n honest parties. We let k denote the number of sessions that \mathcal{M} activates in honest parties, and let




Fig. 2. NAXOS+ AKE protocol

m_1 and m_2 respectively denote the number of queries that \mathcal{M} makes to H_1 and H_2 . We now start by observing that since the session key of the test session is computed as $K = H_2(\sigma)$ for some 6-tuple σ , the adversary has only two ways to distinguish K from a random string.

- $A1$: At some point, \mathcal{M} queries H_2 on the same 6-tuple σ .
- $A2$: \mathcal{M} forces the establishment of another session key that is computed as $K = H_2(\sigma)$ for the same 6-tuple σ , and then reveal the session key.

The input σ determines the transcript of the session; The session should be either a matching session or an equivalent session of the test session. Since \mathcal{M} is not allowed to make queries to reveal any session key of the matching sessions or equivalent sessions, event $A2$ could not occur. Therefore we conclude that

$$\text{Adv}_{NAXOS+}^{\text{ake}}(\mathcal{M}) \leq \Pr[A1] + \frac{1}{2} (1 - \Pr[A1]) - \frac{1}{2}, \quad (1)$$

or

$$2 \text{Adv}_{NAXOS+}^{\text{ake}}(\mathcal{M}) \leq \Pr[A1]. \quad (2)$$

Now we define the following three events.

- $E1$: There exists an honest party \mathcal{B} such that \mathcal{M} makes a random oracle query $H_1(*, sk_{\mathcal{B}})$, either before or without making a query $\text{StaticKeyReveal}(\mathcal{B})$.
- $E2$: $E1$ does not occur, $A1$ occurs and the test session has a matching session.
- $E3$: $E1$ does not occur, $A1$ occurs and the test session has no matching session.

Then we have

$$\Pr[A1] \leq \Pr[E1] + \Pr[E2] + \Pr[E3]. \quad (3)$$

Case 1. In this case, we suppose $\Pr[E1] = \max(\Pr[E1], \Pr[E2], \Pr[E3])$, and construct a DLOG solver \mathcal{S} that uses the adversary \mathcal{M} as a subroutine. Given a challenge $V \in \mathbb{G}$, \mathcal{S} prepares n parties. One party, we call \mathcal{B} , is selected at random and assigned static public key $pk_{\mathcal{B}} = V$. The remaining $n - 1$ parties are assigned random static key pairs. The solver \mathcal{S} initiates \mathcal{M} on this set of parties and awaits the actions of \mathcal{M} . When \mathcal{M} makes a query to any party except \mathcal{B} , \mathcal{S} follows the protocol description. \mathcal{S} is also able to respond to the queries $\text{EstablishParty}(\cdot)$ and $\text{Test}(\cdot)$ faithfully. Thus we only describe how \mathcal{S} simulates the party \mathcal{B} and responds to the other oracle queries including the random oracles H_1 and H_2 .

- Simulation of \mathcal{B} : On the incoming messages $(\mathcal{B}, \mathcal{A})$ and $(\mathcal{B}, (\mathcal{A}, j), X)$, \mathcal{S} creates a new session $sid = (\mathcal{B}, i)$ and selects a random ephemeral key $esk_{(\mathcal{B}, i)} \in \{0, 1\}^\lambda$. \mathcal{S} selects $r_i, s_i \in \mathbb{Z}_q$ independently and uniformly at random, except in the case of $esk_{(\mathcal{B}, i)} = esk_{(\mathcal{B}, i')}$ for some previous session (\mathcal{B}, i') . For the exception, \mathcal{S} sets $r_i = r_{i'}$ and $s_i = s_{i'}$. Then \mathcal{S} presents \mathcal{M} with the outgoing messages $(\mathcal{A}, (\mathcal{B}, i), g^{s_i}/V^{r_i})$ for $(\mathcal{B}, \mathcal{A})$ and $((\mathcal{A}, j), (\mathcal{B}, i), g^{s_i}/V^{r_i})$ for $(\mathcal{B}, (\mathcal{A}, j), X)$, respectively. On the incoming message $((\mathcal{B}, i), (\mathcal{A}, j), X)$, \mathcal{S} makes no response.
- $\text{StaticKeyReveal}(\mathcal{A})$: If $\mathcal{A} = \mathcal{B}$, then \mathcal{S} aborts. Otherwise, \mathcal{S} submits the static private key that \mathcal{S} has generated in the initialization phase.
- $\text{EphemeralKeyReveal}(sid)$: \mathcal{S} submits the value esk_{sid} selected for sid .
- $\text{SessionKeyReveal}(sid)$: \mathcal{S} submits the value K_{sid} selected as follows.
 - If sid is not owned by \mathcal{B} , then \mathcal{S} is able to compute the input σ such that $K_{sid} = H_2(\sigma)$. Then \mathcal{S} computes $H_2(\sigma)$ following the recipe of H_2 described below.
 - If sid is owned by \mathcal{B} , then \mathcal{S} checks if there is a previous session with the same transcript as sid , in which case the previous session key is selected as K_{sid} .
 - Otherwise, let $T_{sid} = (\mathcal{B}, \mathcal{A}, role, g^{s_i}/V^{r_i}, X)$ and $role = resp$. Then the session key should be computed as $K_{sid} = H_2(\sigma)$, where

$$\sigma = \left(\text{CDH}(pk_{\mathcal{A}}, V), \text{CDH}(pk_{\mathcal{A}}, \frac{g^{s_i}}{V^{r_i}}), \text{CDH}(X, V), \text{CDH}(X, \frac{g^{s_i}}{V^{r_i}}), \mathcal{A}, \mathcal{B} \right).$$

\mathcal{S} checks if H_2 was queried with $(W_1, W_2, Z_1, Z_2, \mathcal{A}, \mathcal{B})$ such that

$$W_1^{r_i} W_2 = pk_{\mathcal{A}}^{s_i} \text{ and } Z_1^{r_i} Z_2 = X^{s_i},$$

in which case the answer to that query is selected as the session key. If no such query was made, then \mathcal{S} assigns a random value to K_{sid} . For the case $role = init$, K_{sid} is selected in an analogous way.

- $H_1(x, y)$: \mathcal{S} checks if $g^y \stackrel{?}{=} V$. If this is the case, \mathcal{S} succeeds in computing $\text{DLOG}(V)$. Otherwise, \mathcal{S} submits a random value by lazy sampling.
- $H_2(\sigma)$: \mathcal{S} submits the value $H_2(\sigma)$ selected as follows.
 - \mathcal{S} checks if H_2 was queried with the same input before, in which case the answer to the previous query is selected as $H_2(\sigma)$.
 - If $\sigma = (W_1, W_2, Z_1, Z_2, \mathcal{A}, \mathcal{B})$ for some \mathcal{A}, W_j and $Z_j, j = 1, 2$, then \mathcal{S} checks if $\text{SessionKeyReveal}(\cdot)$ was queried with any sid of $T_{sid} = (\mathcal{B}, \mathcal{A}, resp, g^{s_i}/V^{r_i}, X)$ such that

$$W_1^{r_i} W_2 = pk_{\mathcal{A}}^{s_i} \text{ and } Z_1^{r_i} Z_2 = X^{s_i},$$

in which case the session key is selected as $H_2(\sigma)$. For the case $\sigma = (W_1, Z_1, W_2, Z_2, \mathcal{B}, \mathcal{A})$ for some \mathcal{A}, W_j and $Z_j, j = 1, 2$, $H_2(\sigma)$ is selected in an analogous way.

- Otherwise, \mathcal{S} selects a random value as $H_2(\sigma)$.

We observe the probability that some trapdoor test yields an incorrect answer is at most $2km_2/q$. If every trapdoor test is correct, \mathcal{S} perfectly simulates the \mathcal{M} 's environment until \mathcal{M} queries $\text{StaticKeyReveal}(\mathcal{B})$. Since in a perfect simulation \mathcal{M} will query $H_1(*, \text{DLOG}(V))$ without first issuing a $\text{StaticKeyReveal}(\mathcal{B})$ query with probability $\Pr[E1]/n$, we can estimate the success probability of \mathcal{S} as follows.

$$\text{Adv}^{\text{dlog}}(\mathcal{S}) \geq \frac{\Pr[E1]}{n} - \frac{2km_2}{q} \geq \frac{2}{3n} \text{Adv}_{\text{NAXOS}^+}^{\text{ake}}(\mathcal{M}) - \frac{2km_2}{q}. \quad (4)$$

Note that the solver \mathcal{S} runs in time

$$t_{\mathcal{S}} = O\left(t + \underbrace{n\lambda + k\lambda}_{\text{preparation of public keys}} + \underbrace{m_1\lambda + m_1^2 + k^2}_{\text{simulation of } H_1} + \underbrace{m_2k\lambda + m_2^2 + k^2}_{\text{simulation of } H_2 \text{ and SessionKeyReveal}(\cdot)}\right), \quad (5)$$

since a group exponentiation takes time $O(\lambda)$.

Case 2. In this case, we suppose $\Pr[E2] = \max(\Pr[E1], \Pr[E2], \Pr[E3])$, and construct a Diffie-Hellman solver \mathcal{S} that uses the adversary \mathcal{M} as a subroutine. Given challenges $U, V \in \mathbb{G}$, \mathcal{S} prepares n parties and assigns random static key pairs. \mathcal{S} faithfully responds to any query from \mathcal{M} , with the only difference being that \mathcal{S} chooses two sessions sid and sid^* independently and uniformly at random, and sets their ephemeral public keys as

$$U = g^{H_1(esk_{sid}, sk_{\mathcal{A}})} \text{ and } V = g^{H_1(esk_{sid^*}, sk_{\mathcal{B}})}. \quad (6)$$

Here we suppose sessions sid and sid^* are owned by \mathcal{A} and \mathcal{B} , respectively.

With probability $2\Pr[E2]/k^2$, one of the two sessions is the test session, the other its matching session, and \mathcal{M} queries H_2 with $\sigma = (*, *, *, Z, *, *)$ such that $Z = \text{CDH}(U, V)$. Since \mathcal{S} cannot check if $Z \stackrel{?}{=} \text{CDH}(U, V)$, \mathcal{S} selects one query uniformly at random from the set of m_2 queries that it makes to H_2 . Now the only remaining concern is that \mathcal{M} distinguishes the simulated experiment from a true experiment. The only way to do so is to query $(esk_{sid}, sk_{\mathcal{A}})$ or $(esk_{sid^*}, sk_{\mathcal{B}})$ to H_1 , and check if the equalities (6) hold. Since sid is a clean session, \mathcal{M} should query for at most one of the values in each pair $(esk_{sid}, sk_{\mathcal{A}})$ and $(esk_{sid^*}, sk_{\mathcal{B}})$. If \mathcal{M} reveals esk_{sid} via an `EphemeralKeyReveal`(\cdot) query, then \mathcal{M} cannot query $(esk_{sid}, sk_{\mathcal{A}})$ to H_1 since we exclude the possibility of event $E1$. Without an `EphemeralKeyReveal`(\cdot) query, \mathcal{M} cannot obtain any information about esk_{sid} since esk_{sid} is used in only one session. The same argument applying to esk_{sid^*} , we conclude that the probability that \mathcal{M} correctly makes a random guess for one of esk_{sid} and esk_{sid^*} , and then query the random oracle H_1 with $(esk_{sid}, sk_{\mathcal{A}})$ or $(esk_{sid^*}, sk_{\mathcal{B}})$ is at most $(k + m_1)/2^{\lambda-1}$. Now the success probability and the run time of \mathcal{S} are estimated as

$$\text{Adv}^{\text{cdh}}(\mathcal{S}) \geq \frac{2\Pr[E2]}{m_2k^2} \left(1 - \frac{k + m_1}{2^{\lambda-1}}\right) \geq \frac{4\text{Adv}_{\text{NAXOS}^+}^{\text{ake}}(\mathcal{M})}{3m_2k^2} \left(1 - \frac{k + m_1}{2^{\lambda-1}}\right), \quad (7)$$

and

$$t_{\mathcal{S}} = O\left(t + n\lambda + k\lambda + m_1^2 + m_2^2 + k^2\right), \quad (8)$$

respectively.

Case 3. In this case, we suppose $\Pr[E3] = \max(\Pr[E1], \Pr[E2], \Pr[E3])$, and construct a CDH solver \mathcal{S} that uses the adversary \mathcal{M} as a subroutine. Given challenges $U, V \in \mathbb{G}$, \mathcal{S} prepares n parties. One party, we call \mathcal{B} , is selected at random and assigned static public key $pk_{\mathcal{B}} = V$. The remaining $n - 1$ parties are assigned random static key pairs. \mathcal{S} also chooses a session sid uniformly at random, and sets its ephemeral public key as $U = g^{H_1(esk_{sid}, sk_{\mathcal{A}})}$. The simulation of \mathcal{S} is exactly the same as Case 1, with the only difference being that \mathcal{S} does not check if it obtained $\text{DLOG}(V)$ on a query $H_1(x, y)$, and \mathcal{S} aborts if `SessionKeyReveal`(\cdot) is queried with sid .

Now \mathcal{S} perfectly simulates the random oracle H_2 except with probability $2km_2/q$. In a perfect simulation of \mathcal{M} 's environment, with probability $\Pr[E3]/nk$, sid is the test session with a peer \mathcal{B} , and \mathcal{M} queries H_2 with $\sigma = (*, *, Z, *, \mathcal{A}, \mathcal{B})$ or $\sigma = (*, Z, *, *, \mathcal{B}, \mathcal{A})$ such that $Z = \text{CDH}(U, V)$. Since \mathcal{S} cannot check if $Z \stackrel{?}{=} \text{CDH}(U, V)$, \mathcal{S} selects one query uniformly at random from the set of m_2 queries that it makes to H_2 . The only remaining possibility that we should exclude is that \mathcal{M} queries $(esk_{sid}, sk_{\mathcal{A}})$ to H_1 , and check if the equality $U = g^{H_1(esk_{sid}, sk_{\mathcal{A}})}$ holds. This event occurs with probability at most $(k + m_1)/2^{\lambda-1}$, on the condition that event $E3$ occurs with the test session sid and \mathcal{B} is the peer of sid . Therefore we can estimate the success probability of \mathcal{S} as follows.

$$\begin{aligned} \text{Adv}^{\text{cdh}}(\mathcal{S}) &\geq \frac{\Pr[E3]}{m_2nk} \left(1 - \frac{k + m_1}{2^{\lambda-1}}\right) - \frac{2km_2}{q} \\ &\geq \frac{2 \text{Adv}_{\text{NAXOS+}}^{\text{ake}}(\mathcal{M})}{3m_2nk} \left(1 - \frac{k + m_1}{2^{\lambda-1}}\right) - \frac{2km_2}{q}. \end{aligned} \quad (9)$$

The solver \mathcal{S} runs in time

$$t_{\mathcal{S}} = O(t + n\lambda + k\lambda + m_1^2 + m_2k\lambda + m_2^2 + k^2). \quad (10)$$

To summarize, we obtain the following theorem.

Theorem 2. *If H_1 and H_2 are independent random oracles, and \mathbb{G} is a group where the CDH assumption holds, then the NAXOS+ AKE protocol satisfies eCK security.*

5 Discussion

One-pass variant of NAXOS+ Here we present a natural one-pass variant of NAXOS+. For a detailed description of the eCK model for one-pass AKE protocols, we refer to [23]. Figure 3 depicts the protocol in a simplified form. The security is proved under the GDH assumption, while the security under the CDH assumption is open.




Fig. 3. One-pass variant of NAXOS+

Three-pass variant of NAXOS+ The two-pass NASOX+ protocol does not provide assurance that the intended communicating peer has computed the session key. To address this issue, we add one more pass of message to NAXOS+ to obtain a three-pass variant, called

NAXOS+C. NAXOS+C involves a message authentication code(MAC) to construct a confirmation message. The MAC is generated by an additional hash function $H_3 : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, independent from H_1 and H_2 . The NAXOS+C protocol also provides for *full* perfect forward secrecy. Figure 4 depicts the protocol in a simplified form.




Fig. 4. Three-pass variant of NAXOS+

Precomputation of static pairwise keys Any pair of parties \mathcal{A} and \mathcal{B} might precompute their pairwise static key $\text{CDH}(pk_{\mathcal{A}}, pk_{\mathcal{B}})$ in order to reduce the on-line computational overload required during a session. In this case, the possibility of revelation of the pairwise static key should be captured in the security model. By slightly modifying the eCK model and the security proof of NAXOS+, we can prove that *if every party is honest*, then NAXOS+ remains secure against revelation of the pairwise static key.

Post-specified peer model In general, AKE protocols are analyzed in the pre-specified peer model, wherein it is assumed that a party knows the identifier of its intended communicating peer when it commences a run of the protocol. On the other hand, in the post-specified peer model which is introduced by Canetti and Krawczyk [6], a party initiates a new session knowing only a destination address of the communicating peer, and learns the peer's identifier during the execution of the protocol. Recently, Menezes and Ustaoglu compared the Canetti-Krawczyk pre- and post-specified peer models, and demonstrated that security in one model does not guarantee security or even executability in the other model [20]. As an example, they presented an unknown key-share attack on the HMQV protocol in the post-specified peer model.

The NAXOS+ protocol is easily adapted with minor changes to be executable in the post-specified peer model. Furthermore, the resulting protocol resists the unknown key share attack presented in [20] since the identifiers of the communicating parties are included in the key derivation function.

References

1. Antipa A., Brown D., Menezes A., Struik R., Vanstone S.: Validation of elliptic curve public keys. In: Desmedt Y.G. (ed.) PKC 2003. LNCS, vol. 2567, pp. 211–223. Springer, Heidelberg (2003)
2. Bellare M., Palacio A.: The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In: Franklin M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 273–289. Springer, Heidelberg (2004)
3. Bellare M., Pointcheval D., Rogaway P.: Authenticated key exchange secure against dictionary attacks. In: Pfitzmann B. (ed.) EUROCRYPT 2001. LNCS, vol. 1807, pp. 139–155. Springer, Heidelberg (2001)
4. Bellare M., Rogaway P.: Entity authentication and key distribution. In: Stinson D.R. (ed.) CRYPTO 1993, LNCS, vol. 773, pp. 110–125. Springer, Heidelberg (1994)
5. 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, Heidelberg (2001). Full version available at <http://eprint.iacr.org/2001/040>
6. Canetti R., Krawczyk H.: Security analysis of IKE's signature based key-exchange protocol. In: Yung M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 143–161. Springer, Heidelberg (2002). Full version available at <http://eprint.iacr.org/2002/120>
7. Choo K.K., Boyd C., Hitchcock Y.: Examining indistinguishability-based proof models for key establishment protocols. In: Roy B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 585–604. Springer, Heidelberg (2005)
8. Cash D., Kiltz E., Shoup V.: The twin Diffie-Hellman problem and applications. In: Smart N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008)
9. Diffie W., Hellman H.: New directions in cryptography. *IEEE transactions of Information Theory*, vol. 22(6): 644–654 (1976)
10. 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 2004. LNCS, vol. 3089, pp. 220–232. Springer, Heidelberg (1999)
11. Krawczyk H.: HMQV: A high-performance secure Diffie-Hellman protocol. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 546–566. Springer, Heidelberg (2005). Full version available at <http://eprint.iacr.org/2005/176>.
12. Kudla, C., Paterson, K.C.: Modular security proofs for key agreement protocols. Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 549–565 (2005)
13. LaMacchia K., 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, Heidelberg (2007)
14. Lauter K., Mityagin, A.: Security analysis of KEA authenticated key exchange. In: Yung M., Dodis Y., Kiayias A., Malkin T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 378–394. Springer, Heidelberg (2006)
15. Law L, Menezes A, Qu M., Solinas J., Vanstone S.: An efficient protocol for authenticated key agreement. *Designs, Codes and Cryptography*, vol. 28:119–134 (2003)
16. Lim C., Lee P.: A key recovery attack on discrete log-based schemes using a prime order subgroup. In: Desmedt Y.G. (ed.) CRYPTO 1994. LNCS, vol. 1294, pp. 249–263. Springer, Heidelberg (1994)
17. Lee J., Park C.: An Efficient Authenticated Key Exchange Protocol with a Tight Security Reduction. Preprint (2008)
18. Menezes A.: Another look at HMQV. *Journal of Mathematical Cryptology*, vol. 1(1): 148–175 (2007)
19. Menezes A., Ustaoglu B.: On the importance of public-key validation in the MQV and HMQV key agreement protocols. In: Lange T. (ed.) INDOCRYPT 2006. LNCS, vol. 4329, pp. 133–147. Springer, Heidelberg (2006)
20. Menezes A., Ustaoglu B.: Comparing the pre- and post-specified peer models for key agreement. In: Mu Y., Susilo W., Seberry J. (eds.) ACISP 2008. LNCS, vol. 5107, pp. 53–68. Springer, Heidelberg (2008)
21. Okamoto T.: Authenticated key exchange and key encapsulation in the standard model. Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 474–484 (2007)
22. Okamoto T., Pointcheval D.: The Gap-Problems: A new class of problems for the security of cryptographic schemes. In: Kim K. (ed.) PKC 2001, LNCS, vol. 1992, pp. 104–118. Springer, Heidelberg (2001)
23. Ustaoglu B.: Obtaining a secure and efficient key agreement protocol for (H)MQV and NAXOS. *Designs, Codes and Cryptography*, vol. 46(3): 329–342 (2008)