On the Security of the Pre-shared Key Ciphersuites of TLS

Yong Li¹, Sven Schäge^{2,*}, Zheng Yang¹, Florian Kohlar¹, and Jörg Schwenk¹

¹ Ruhr-Universität Bochum, Germany

{yong.li,zheng.yang,florian.kohlar,joerg.schwenk}@rub.de

² University College London, UK

s.schage@ucl.ac.uk

Abstract. TLS is by far the most important protocol on the Internet for negotiating secure session keys and providing authentication. Only very recently, the standard ciphersuites of TLS have been shown to provide provably secure guarantees under a new notion called Authenticated and Confidential Channel Establishment (ACCE) introduced by Jager *et al.* at CRYPTO'12. In this work, we analyse the variants of TLS that make use of pre-shared keys (TLS-PSK). In various environments, TLS-PSK is an interesting alternative for remote authentication between servers and constrained clients like smart cards, for example for mobile phone authentication, EMV-based payment transactions or authentication via electronic ID cards. First, we introduce a new and strong definition of ACCE security that covers protocols with pre-shared keys. Next, we prove that all ciphersuite families of TLS-PSK meet our strong notion of ACCE security. Our results do not rely on random oracles nor on any non-standard assumption.

Keywords: TLS, TLS-PSK, ACCE, Pre-Shared Keys, Authenticated Key Exchange, Secure Channels.

1 Introduction

TLS is undeniably the most prominent key exchange protocol in use today. While the security of most web applications relies on the classical Diffie-Hellman and RSA-based ciphersuites of TLS, there also exist several important applications that make use of one of the less common ciphersuites [31,1,29]. One such application is (remote) authentication of resource-restricted clients like smart-cards. In these scenarios, computational efficiency and low power consumption often are one of the most important system features. Instead of using the public-key based ciphersuites of TLS, applications can apply a variant of TLS that assumes pre-shared symmetric keys between client and server. The corresponding ciphersuite family is termed TLS with pre-shared keys (TLS-PSK) and available in many TLS releases and libraries, for example [28,24,7].

* Supported by EPSRC grant number EP/J009520/1.

RELATED WORK: ON THE SECURITY OF TLS. Since the introduction of its predecessor SSL, the security of TLS has often been the focus of security researchers and attackers worldwide. Over the time, several attacks on TLS have been published. Most of these attacks do not directly attack the cryptographic core of TLS, but rather exploit side-channels or vulnerabilities in associated technologies, like the famous Bleichenbacher attack [6], or attacks on the domain name system or the public-key infrastructure [20,10,26]. However, despite that no serious attacks on the cryptographic core of the current TLS protocol are known, determining exactly what security guarantees TLS provides has been an elusive problem for many years. This is partly due to the fact that the popular TLS ciphersuites provably do not provide security in the sense of authenticated key exchange (AKE) protocols [3], the classical and very strong standard notion of security of key exchange protocols (which requires that the session key remains indistinguishable from random even if the adversary obtains the communication transcript of the session). Until recently only security analyses of modified versions of TLS were published [18,16,27]. At CRYPTO 2012, Jager, Kohlar, Schäge, and Schwenk (JKSS) [17] were the first to present a detailed security analysis of the *unmodified* version of one of TLS's ciphersuite families. They showed that the cryptographic core of ephemeral Diffie-Hellman with mutual authentication is a provably secure authenticated and confidential channel establishment (ACCE) protocol in the standard model. ACCE is a new security notion that is particularly well suited to capture what protocols like TLS intuitively want to achieve: the establishment of a secure channel between client and server. Among its features, it not only formalizes confidentiality and integrity of messages exchanged between client and server, but also covers replay and re-ordering attacks. Very recently, Krawczyk, Paterson, and Wee (KPW) [22] and independently Kohlar, Schäge, Schwenk (KSS) [19] presented, while relying on different cryptographic assumptions and security models¹, extensions of the JKSS result to the remaining ciphersuite families. In particular, they show that TLS-RSA and TLS-DH also constitute ACCE protocols when used for mutual authentication setting and that TLS-RSA, TLS-DH, and TLS-DHE are ACCE secure in the practically important setting of server-only authentication (for which they provide new formal security definitions).

Unfortunately, all previous results on the (ACCE) security of TLS are based on either i) new, non-standard security assumption like the PRF-ODH assumption introduced in [17] and refined in [22,19] or ii) strong idealizations such as the modeling of TLS's key derivation function as a random oracle [2] or assuming that the public-key encryption scheme in TLS-RSA is replaced with an IND-CCA secure one. Looking somewhat ahead, for the TLS-PSK ciphersuites, fortunately the situation is different, i.e. security can be based on standard assumptions only.

¹ The security models and complexity assumptions differ mainly with respect to the capabilities granted to the adversary when corrupting and registering new parties and the application of the random oracle model.

TLS WITH PRE-SHARED KEYS. The original specifications of the TLS protocol [11,12,13] do not explicitly include ciphersuites that support authentication and key exchange using pre-shared keys. However, since 2005 there exists an extension called “Pre-Shared Key Ciphersuites for Transport Layer Security” (TLS-PSK) which specifically describes such ciphersuites in RFC 4279 [14]. The TLS-PSK standard specifies three ciphersuites, `TLS_PSK`, `TLS_RSA_PSK` and `TLS_DHE_PSK`, each of which derives the master secret in a different way. In `TLS_PSK`, the master secret is solely based on the secret pre-shared keys. In the remaining ciphersuites the computation of the master secret is additionally dependent on freshly exchanged secrets via encrypted key transport in `TLS_RSA_PSK` or Diffie-Hellman key exchange in `TLS_DHE_PSK`. The intuition is that as long as either the pre-shared key or the freshly exchanged secret is not compromised, then the protocol produces a secure application key. All three ciphersuites assume that the client only has a pre-shared key for authentication. Although it is not as widespread as TLS with RSA key transport, several interesting and important scenarios for TLS with pre-shared keys exist where its efficiency makes TLS-PSK a much more attractive alternative than, for example, TLS with self-signed certificates.

- Since November 2010, the new electronic German ID (eID) card supports online remote authentication of the eID card holder to some online service (eService). Here TLS-PSK is applied to perform mutual authentication between the two parties [15].
- As a second example, we mention the application of TLS-PSK in the Generic Authentication Architecture, the 3GPP mobile phone standard for UMTS and LTE. According to ETSI TR 133 919 V11.0.0 (2012-11), TLS-PSK can be used to secure the communication between server and user equipment.
- An IETF draft from 2009 for EMV smart cards describes an authentication protocol based on TLS-PSK [29]. EMV chips are widely deployed and are used commonly for secure payment transactions [9].

CONTRIBUTION. In this paper, we provide a security analysis of all three TLS-PSK ciphersuites. Similar to classical TLS, it is provably impossible to show that the keys produced by TLS-PSK are indistinguishable from random. Therefore, as one of our main contributions, we introduce the first definition of ACCE security for authentication protocols with pre-shared keys. We do not propose a separate model but rather an extension of the ACCE model of JKSS to also cover authentication via pre-shared keys. Next, we introduce a strengthened variant of this definition called *asymmetric perfect forward secrecy*, that captures that protocol sessions of ACCE protocols with pre-shared keys may retain a strong level of confidentiality even if the long-term secrets of the client are exposed after the protocol run. Asymmetric perfect forward secrecy is a strong security notion that can hold for protocols that do not fulfill the standard notion of perfect forward secrecy. This allows us to prove the security of such protocols in a stronger security model than was previously possible. We show that `TLS_PSK` is ACCE secure (without forward secrecy), `TLS_RSA_PSK` is ACCE secure with

asymmetric perfect forward secrecy and `TLS_DHE_PSK` is secure with (classical) perfect forward secrecy. Informally, our results say that `TLS_PSK` guarantees confidentiality and integrity of all messages exchange between client and server, unless the adversary has learned the pre-shared key or corrupted one of the parties to learn the application/session key. In `TLS_DHE_PSK` the communication remains confidential even if the adversary corrupts the pre-shared secret later on. In contrast, in `TLS_RSA_PSK` the communication remains confidential even if the adversary manages to corrupt the pre-shared key or the server's long-term key later on, but not both of them.

DOUBLE PRFS AND FORWARD SECRECY. To prove `TLS_RSA_PSK` and `TLS_DHE_PSK`, we introduce a variant of pseudo-random functions (PRFs), called double pseudo-random function (DPRF). Roughly, a DPRF takes as input two keys only one of which is generated randomly and kept secret from the attacker (as in classical PRFs). However, when the adversary makes its queries, not only the message but also the other key can entirely be specified by the adversary. Our notion of DPRF nicely abstracts the crucial mechanism in `TLS_PSK` that is required to guarantee (asymmetric) perfect forward secrecy. In our security proofs, we assume that `TLS`'s key derivation function provides a suitable DPRF in the standard model. Existing results on the security of HMAC support this assumption for `TLS` 1.1 when the pre-shared key has a specific bit length. Our new DPRF notion may be of independent interest beyond the scope of this work.

Note also, that for the `TLS_PSK` and `TLS_DHE_PSK` ciphersuites we neither have to rely on non-standard assumptions like the PRF-ODH assumption of JKSS to give a proof nor on idealized setup assumptions like the random oracle model. We can show that `TLS_RSA_PSK` is secure under our basic notion of ACCE security without any assumption on the public key encryption system used in `TLS`. This is because under the basic ACCE definition security can be derived solely from secrecy of the pre-shared keys. However, if we want to prove the ACCE security of `TLS_RSA_PSK` with asymmetric perfect forward secrecy in the *standard model* we need to assume that the public key encryption scheme is IND-CCA secure², similar to [22,19]. Thus, we do not consider `TLS-RSA` with `RSA-PKCS` encryption as it is currently used in practice. We remark that [22] were also able to prove security of the classical `TLS` ciphersuites based on `RSA` key transport with `RSA-PKCS` encryption in the random oracle model. It would be interesting to show that the results of KPW on `TLS-RSA` can be transferred to show that `TLS_PSK` with `RSA-PKCS` based key transport provides asymmetric perfect forward secrecy in the random oracle model.

LIMITATIONS. In our work, we give a dedicated security analysis for `TLS_PSK`. We believe that it is possible to give a more modularized analysis, similar to KPW [22] who analyzed the classical ciphersuites of `TLS` by abstracting the handshake phase into a Constrained-CCA-secure (CCCA) KEM that is combined with a secure authenticated encryption scheme. The benefit of the KPW analysis is re-usability: once the security proof is established for a generic

² KPW call this `TLS-CCA`.

CCCA-secure KEM, all that remains is to show that each of the ciphersuites indeed provides such a KEM.

2 Security Assumptions

To state our results, we will rely on standard security definitions for the Decisional Diffie-Hellman assumption (DDH), collision-resistant cryptographic hash functions, IND-CCA secure public key encryption schemes, (plain) pseudo-random functions (PRF), and stateful length-hiding authenticated encryption (sLHAE) schemes as recently defined in [30]. However, we will sometimes also rely on a new class of PRFs called double pseudo-random functions.

DOUBLE PSEUDO-RANDOM FUNCTIONS. Double pseudo-random functions can be thought of as a class of pseudo-random functions with two keys. Let $\text{DPRF} : \mathcal{K}_{\text{DPRF}_1} \times \mathcal{K}_{\text{DPRF}_2} \times \mathcal{M}_{\text{DPRF}} \rightarrow \mathcal{R}_{\text{DPRF}}$ denote a family of deterministic functions, where $\mathcal{K}_{\text{DPRF}_1}, \mathcal{K}_{\text{DPRF}_2}$ is the key space, $\mathcal{M}_{\text{DPRF}}$ is the domain and $\mathcal{R}_{\text{DPRF}}$ is the range of PRF.

Intuitively, security requires that the output of the DPRF is indistinguishable from random as long as one key remains hidden from the adversary even if the adversary is able to adaptively specify the second key and the input message. To formalize security we consider the following security game played between a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$. Let $\text{RF}_{\text{DPRF}}(\cdot, \cdot)$ denote an oracle implemented by $\mathcal{C}$, which takes as input a key $k_j \in \mathcal{K}_{\text{DPRF}_j}$ (where j is specified by the adversary via an Init query) and message $m \in \mathcal{M}_{\text{DPRF}}$ and outputs a random value $z \in \mathcal{R}_{\text{DPRF}}$.

1. The adversary first runs $\text{Init}(j)$ with $j \in \{1, 2\}$ to specify the key $k_j \in \mathcal{K}_{\text{DPRF}_j}$ that he wants to manipulate.
2. The challenger $\mathcal{C}$ samples $\hat{b} \xleftarrow{\$} \{0, 1\}$, and sets $u = (j \bmod 2) + 1$. If $\hat{b} = 0$, the challenger samples $k_u \in_R \mathcal{K}_{\text{DPRF}_u}$ and assigns $\text{RF}_{\text{DPRF}}(\cdot, \cdot)$ to either $\text{DPRF}(\cdot, k_2, \cdot)$ or $\text{DPRF}(k_1, \cdot, \cdot)$ depending on the value of u . For instance, if $u = 2$ then the random function RF_{DPRF} is assigned to $\text{DPRF}(\cdot, k_2, \cdot)$, and the $\mathcal{A}$ is allowed to specify k_1 arbitrarily in each query. If $\hat{b} = 1$, the challenger assigns RF_{DPRF} to $\text{RF}(\cdot, \cdot)$ which is a truly random function that takes as input key k_j and message m and outputs a value in the same range $\mathcal{R}_{\text{DPRF}}$ as $\text{DPRF}(\cdot, \cdot, \cdot)$.
3. The adversary may adaptively make queries $k_{j,i}, m_i$ for $1 \leq i \leq q$ to oracle RF_{DPRF} and receives the result of $\text{RF}_{\text{DPRF}}(k_{j,i}, m_i)$, where $k_{j,i}$ denotes the i -th key k_j chosen by $\mathcal{A}$.
4. Finally, $\mathcal{A}$ outputs its guess $\hat{b}' \in \{0, 1\}$ of $\hat{b}$. If $\hat{b} = \hat{b}'$, $\mathcal{A}$ wins.

Definition 1. We say that DPRF is a (t, ϵ) -secure double pseudo-random function, if any adversary running in time t has at most an advantage of ϵ to distinguish the DPRF from a truly random function, i.e.

$$\Pr [\hat{b} = \hat{b}'] \leq 1/2 + \epsilon.$$

The number of allowed queries q is upper bounded by t .

3 A Brief Introduction to TLS-PSK

This section describes the three sets of ciphersuites specified in TLS-PSK: `TLS_PSK`, `TLS_RSA_PSK` and `TLS_DHE_PSK`. In each of these ciphersuites, the master secret is computed using pre-shared keys which are symmetric keys shared in advance among the communicating parties. The main differences are in the way the master secret is computed. The following description is valid for *all* `TLS_PSK` versions. We only describe the cryptographically relevant messages and only those that deviate from the classical TLS ciphersuites. A detailed description can be found in the full version.

Fig. 1. Handshake in TLS-PSK

SERVERCERTIFICATE. For `TLS_PSK` and `TLS_DHE_PSK`, the message is not included. In `TLS_RSA_PSK` `certS` contains a public key that is bound to the server's identity.

SERVERKEYEXCHANGE. Since clients and servers may have pre-shared keys with many different parties, in the `ServerKeyExchange` message m_4 , the server provides a `PSK identity hint` pointing to the PSK used for authentication. However, for ephemeral Diffie-Hellman key exchange, the Diffie-Hellman (DH) key exchange parameters are also contained in the `ServerKeyExchange` messages including information on the DH group (e.g. a large prime number $p \in \{0, 1\}^{poly(\kappa)}$, where κ is the security parameter, and a generator $\langle g \rangle$ for a prime-order q subgroup of $\mathbb{Z}_p^*$), and the DH share T_S ($T_S = g^{t_S}$, where t_S is a random value in $\mathbb{Z}_q$). (We implicitly assume that the client checks whether the received parameters are valid, in particular if T_S is indeed in the group generated by g .)

CLIENTKEYEXCHANGE. Message m_6 is called `ClientKeyExchange`. We describe the contents of this message for the ciphersuites `TLS_DHE_PSK`, `TLS_PSK` and `TLS_RSA_PSK` separately:

- For `TLS_PSK`, the message is not included.
- For ephemeral Diffie-Hellman key exchange `TLS_DHE_PSK`, it contains the Diffie-Hellman share T_C of the client, i.e. $T_C = g^{t_C}$.
- For the RSA-based key exchange `TLS_RSA_PSK` the client selects a 46-byte random value R and sends a 2-byte version number V and the 46-byte random value R encrypted under the server's RSA public key to the server.

Also, the client sends an identifier for the pre-shared key it is going to use when communicating with the server. This information is called `PSK identity`.

COMPUTING THE MASTER SECRET. According to the original specification, released as RFC 4279 [14], the key derivation function of TLS, denoted here as PRF_{TLS} , is used when constructing the master secret. PRF_{TLS} takes as input a secret, a seed, and an identifying label and produces an output of arbitrary length. We first describe the generic computation of the master secret ms for all ciphersuites using pre-shared keys. Then, a detailed description of all cases (TLS_PSK, TLS_DHE_PSK, and TLS_RSA_PSK) is provided. The *master secret* ms is computed as follows:

$$ms := \text{PRF}_{\text{TLS}}(pms, label_1 || r_C || r_S) \quad (1)$$

- TLS_PSK case: For TLS_PSK, the client/server is able to compute the *master secret* ms using the pre-master secret pms , from which all further secret values are derived. If the PSK is N bytes long, the pms consists of the 2-byte representation (uint16) of the integer value N , N zero bytes, the 2-byte representation of N once again, and the PSK itself, i.e. $pms := N || 0 \dots 0 || N || \text{PSK}$. Since the first half of pms is constant for any PSK we get for TLS_PSK that the entire security of PRF_{TLS} only relies on the second half of pms .
- TLS_DHE_PSK case: Let Z be the value produced for DH-based ciphersuites, i.e. $Z = g^{t_{stc}} = T_C^{t_s} = T_S^{t_C}$. The pms consists of a concatenation of four values: the uint16 len_Z indicating the length of Z , Z itself, the uint16 len_{PSK} showing the length of the PSK, and the PSK itself: $pms := len_Z || Z || len_{PSK} || \text{PSK}$.
- TLS_RSA_PSK case: First, the pre-master secret concatenates the uint16 constant $C = 48$, the 2-byte version number V , a 46-byte random value R , the uint16 len_{PSK} containing the length of the PSK, and the PSK itself, i.e. $pms := C || V || R || len_{PSK} || \text{PSK}$.

3.1 On the Security of PRF_{TLS}

In our security proof of TLS_PSK, we assume that the pseudo-random function of TLS (PRF_{TLS}) that is used for the computation of the master-secret constitutes a secure PRF in the standard model when applied with pms as the key. However to prove (asymmetric) perfect forward secrecy in TLS_DHE_PSK and TLS_RSA_PSK, we assume that PRF_{TLS} constitutes a secure DPRF (in the standard model) where the key space of the DPRF consists of the key space of the pre-shared key and the key space of the freshly generated RSA or Diffie-Hellman secret. Unfortunately, existing results do not *directly* prove that PRF_{TLS} as used in TLS-PSK is a secure DPRF. Nevertheless, they might in some cases serve as a strong indicator of the security of PRF_{TLS} . We provide a more detailed analysis of the plausibility of this assumption in the full version.

4 ACCE Protocols

In this section, we present an extension of the formal security model for two party authenticated and confidential channel establishment (ACCE) protocols introduced by JKSS [17] to also cover scenarios with pre-shared, symmetric keys. Additionally, we extend the model to also address PKI-related attacks that exploit

that the adversary does not have to prove knowledge of the secret key when registering a new public key [5]. (In [25] such attacks are generally called strong-key substitution attacks.) For better comparison with JKSS we will subsequently use boxes to highlight state variables that are essentially new in our model.

In this model, while emulating the real-world capabilities of an active adversary, we provide an ‘execution environment’ for adversaries following the tradition of the seminal work of Bellare and Rogaway [3] and its extensions [4,8,21,23,17]. Let $\mathcal{K}_0 = \{0, 1\}^\kappa$ be the key space of the session key and $\mathcal{K}_1 = \{0, 1\}^\kappa$ be the key space of the pre-shared keys.

Execution Environment. In the following let $\ell, d \in \mathbb{N}$ be positive integers. In the execution environment, we fix a set of ℓ honest parties $\{P_1, \dots, P_\ell\}$. Each party is either identified by index i in the security experiment or a unique, fixed-length string id_i (which might appear in the protocol flows).

To cover authentication with symmetric keys, we extend the state of each party to also include pre-shared keys. Each party holds (symmetric) pre-shared keys with all other parties. We denote with $\text{PSK}_{i,j} = \text{PSK}_{j,i}$ the symmetric key shared between parties P_i and P_j . Each party P_i with $i \in \{1, \dots, \ell\}$ also has access to a long-term public/private key pair (pk_i, sk_i) . Formally, each party maintains the state variables given in Table 1.

Table 1. Internal States of Parties

Variable	Description
sk_i	stores the secret key of a public key pair (pk_i, sk_i)
PSK_i	a vector which contains an entry $\text{PSK}_{i,j}$ per party P_j
τ_i	denotes, that sk_i was corrupted after the τ_i -th query of $\mathcal{A}$
f_i	a vector denoting the freshness of all pre-shared keys, containing one entry $f_{i,j} \in \{\text{exposed}, \text{fresh}\}$ for each entry in PSK_i

The first two variables, sk_i and PSK_i , are used to store keys that are used in the protocol execution while the remaining variables are solely used to define security. (When defining security the latter are additionally managed and updated by the challenger.) The variables of each party P_i will be initialized according to the following rules:

- The long-term key pair (pk_i, sk_i) and pre-shared key vector PSK_i are chosen randomly from the key space. For all parties P_i, P_j with $i, j \in \{1, \dots, \ell\}$ and with $i \neq j$, and pre-shared keys PSK_i it holds that $\text{PSK}_{i,j} = \text{PSK}_{j,i}$ and $\text{PSK}_{i,i} := \emptyset$.
- All entries in f_i are set to **fresh**.
- τ_i is set to $\tau_i := \infty$, which means that all parties are initially not corrupted.

In the following, we will call party P_i uncorrupted iff $\tau_i = \infty$. Thus, we do not consider a dedicated variable that holds the corruption state of the secret key

sk_i . Each honest party P_i can sequentially and concurrently execute the protocol multiple times. This is modeled by a collection of oracles $\{\pi_i^s : i \in [\ell], s \in [d]\}$. Oracle π_i^s behaves as party P_i carrying out a process to execute the s -th protocol instance with some partner P_j (which is determined during the protocol execution). All oracles of P_i have access to the long-term keys sk_i and PSK_i with $j \in \{1, \dots, \ell\}$. Moreover, we assume each oracle π_i^s maintains a list of independent internal state variables with the semantics given in Table 2. The variables Φ_i^s ,

Table 2. Internal States of Oracles

Variable	Description
Φ_i^s	denotes π_i^s 's execution-state in $\{\text{negotiating}, \text{accept}, \text{reject}\}$
Pid_i^s	stores the identity of the intended communication partner
ρ_i^s	denotes the role $\rho_i^s \in \{\text{Client}, \text{Server}\}$
$K_i^s = (k_{\text{enc}}, k_{\text{dec}})$	stores the application keys K_i^s
$\text{St}_i^s = (u, v, st_e, st_d, C)$	stores the current states of the sLHAE scheme
T_i^s	records the transcript of messages sent and received by π_i^s
kst_i^s	denotes the freshness $\text{kst}_i^s \in \{\text{exposed}, \text{fresh}\}$ of the session key
b_i^s	stores a bit $b \in \{0, 1\}$ used to define security

Pid_i^s , ρ_i^s , K_i^s , st_e , st_d , and T_i^s are used by the oracles to execute the protocol. The remaining variables are only used to define security. The variables of each oracle π_i^s will be initialized by the following rules:

- The execution-state Φ_i^s is set to **negotiating**.
- The variable kst_i^s is set to **fresh**.
- The bit b_i^s is chosen at random.
- The counters u, v are initialized to 0.
- All other variables are set to only contain the empty string $\emptyset$.

At some point, each oracle π_i^s completes the execution with a decision state $\Phi_i^s \in \{\text{accept}, \text{reject}\}$. Furthermore, we will always assume (for simplicity) that $K_i^s = \emptyset$ if an oracle has not reached **accept**-state (yet).

Matching Conversations. To formalize the notion that two oracles engage in an on-line communication, we define partnership via *matching conversations* as proposed by Bellare and Rogaway [3]. We use the variant by JKSS.

Definition 2. We say that an oracle π_i^s has a matching conversation to oracle π_j^t , if

- π_i^s has sent all protocol messages and T_j^t is a prefix of T_i^s , or
- π_j^t has sent all protocol messages and $\text{T}_i^s = \text{T}_j^t$.

To keep our definition of ACCE protocols general we do not consider protocol-specific definitions of partnership like for example [22] who define partnership of TLS sessions using only the first three messages exchanged in the handshake phase.

Adversarial Model. An adversary $\mathcal{A}$ in our model is a PPT taking as input the security parameter 1^κ and the public information (e.g. generic description of above environment), which may interact with these oracles by issuing the following queries.

Send^{pre}(π_i^s, m): This query sends message m to oracle π_i^s . The oracle will respond with the next message m^* (if there is any) that should be sent according to the protocol specification and its internal states.

After answering a **Send^{pre}** query, the variables $(\Phi_i^s, \text{Pid}_i^s, \rho_i^s, K_i^s, T_i^s)$ will be updated depending on the protocol specification. This query is essentially defined as in JKSS.

RegisterParty($\mu, pk_\mu, [psk]$): This query allows $\mathcal{A}$ to register a new party with a new identity μ and a static public key (pk_μ) to be used for party P_μ . In response, if the same identity μ is already registered (either via a **RegisterParty**-query or $\mu \in [\ell]$), a failure symbol $\perp$ is returned. Otherwise, a new party P_μ is added with the static public key pk_μ . The secret key sk_μ is set to a constant. The parties registered by this query are considered corrupted and controlled by the adversary. If **RegisterParty** is the τ' -th query of the adversary, P_μ is initialized with $\tau_\mu = \tau'$. If the adversary also provides a pre-shared key psk , then this key will be implemented for every party P_i with $i \in [\ell]$ as key $\text{PSK}_{i,\mu}$.³ Otherwise, the simulator chooses a random key $psk \xleftarrow{\$} \{0, 1\}^\kappa$ and sets $\text{PSK}_{i,\mu} = \text{PSK}_{\mu,i} := psk$ for all parties P_i before outputting psk . The corresponding entries $f_{i,\mu}$ in the vectors of the other parties P_i with $i \in [\ell]$ are set to **exposed**. Via this query we extend the ACCE model of JKSS to also model key registration.

RevealKey(π_i^s): Oracle π_i^s responds to a **RevealKey**-query with the contents of variable K_i^s , the application keys. At the same time the challenger sets $\text{kst}_i^s = \text{exposed}$. If at the point when $\mathcal{A}$ issues this query there exists another oracle π_j^t having matching conversation to π_i^s , then we also set $\text{kst}_j^t = \text{exposed}$ for π_j^t . This query slightly deviates from JKSS.⁴

Corrupt($P_i, [P_j]$): Depending on the second input parameter, oracle π_i^1 responds with certain long-term secrets of party P_i . This query extends the corruption capabilities of JKSS to symmetric keys.

- If $\mathcal{A}$ queries **Corrupt**(P_i) or **Corrupt**($P_i, \emptyset$)⁵, oracle π_i^1 returns the long-term secret key sk_i of party P_i . If this query is the τ -th query issued by $\mathcal{A}$, then we say that P_i is τ -corrupted and π_i^1 sets $\tau_i := \tau$.
- If $\mathcal{A}$ queries **Corrupt**(P_i, P_j), oracle π_i^1 returns the symmetric pre-shared key $\text{PSK}_{i,j}$ stored in PSK_i and sets $f_{i,j} := \text{exposed}$.
- If $\mathcal{A}$ queries **Corrupt**($P_i, \top$), oracle π_i^1 returns the vector PSK_i and sets $f_{i,j} := \text{exposed}$ for all entries $f_{i,*} \in f_i$.

³ This is just for simplicity. Modeling different pre-shared keys between the registered party and every other party is equivalent to registering multiple parties with a single shared key each.

⁴ JKSS implicitly located the specification of when to set $\text{kst}_j^t = \text{exposed}$ into the security definition.

⁵ The party P_i is not adversarially controlled.

Encrypt($\pi_i^s, m_0, m_1, \text{len}, H$): This query takes as input two messages m_0 and m_1 , length parameter len , and header data H . If $\Phi_i^s \neq \text{accept}$ then π_i^s returns $\perp$. Otherwise, it proceeds as depicted in Figure 2, depending on the random bit $b_i^s \xleftarrow{\$} \{0, 1\}$ sampled by π_i^s at the beginning of the game and the internal state variables of π_i^s . This query is essentially defined as in JKSS.

Decrypt(π_i^s, C, H): This query takes as input a ciphertext C and header data H . If π_i^s has $\Phi_i^s \neq \text{accept}$ then π_i^s returns $\perp$. Otherwise, it proceeds as depicted in Figure 2. This query is essentially defined as in JKSS.

Encrypt ($\pi_i^s, m_0, m_1, \text{len}, H$): $u := u + 1$ $(C^{(0)}, st_e^{(0)}) \xleftarrow{\$} \text{StE.Enc}(k_{\text{enc}}^\rho, \text{len}, H, m_0, st_e)$ $(C^{(1)}, st_e^{(1)}) \xleftarrow{\$} \text{StE.Enc}(k_{\text{enc}}^\rho, \text{len}, H, m_1, st_e)$ If $C^{(0)} = \perp$ or $C^{(1)} = \perp$ then return $\perp$ $(C_u, H_u, st_e) := (C^{(b)}, H, st_e^{(b)})$ Return C_u	Decrypt (π_i^s, C, H): $v := v + 1$ If $b_i^s = 0$, then return $\perp$ $(m, st_d) = \text{StE.Dec}(k_{\text{dec}}^\rho, H, C, st_d)$ If $v > u$ or $C \neq C_v$ or $H \neq H_v$, then phase := 1 If phase = 1 then return m
---	---

Here $u, v, b_i^s, \rho, k_{\text{enc}}^\rho, k_{\text{dec}}^\rho, C$ denote the values stored in the internal variables of π_i^s .

Fig. 2. Encrypt and Decrypt oracles in the ACCE security experiment

Definition 3 (Correctness). We say that an ACCE protocol Π is correct, if for any two oracles π_i^s, π_j^t that have matching conversations with $\text{Pid}_i^s = j$ and $\text{Pid}_j^t = i$ and $\Phi_i^s = \text{accept}$ and $\Phi_j^t = \text{accept}$ it always holds that $K_i^s = K_j^t$.

Secure ACCE Protocols. We define security via an experiment played between a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$.

SECURITY GAME. Assume there is a global variable *pinfo* which stores the role information of each party for the considered protocol Π .⁶ In the game, the following steps are performed:

1. Given the security parameter κ , $\mathcal{C}$ implements the collection of oracles $\{\pi_i^s : i, j \in [\ell], s \in [d]\}$ with respect to Π and *pinfo*. In this process, $\mathcal{C}$ generates long-term keys PSK_i for all parties $i \in [\ell]$. Next it additionally generates long-term key pairs (pk_i, sk_i) for all parties $i \in [\ell]$ that require them (e.g. if the corresponding party is a server in the `TLS_RSA_PSK` protocol). Finally, $\mathcal{C}$ gives all identifiers $\{id_i\}$, all public keys (if any), and *pinfo* to $\mathcal{A}$.
2. Next the adversary may start issuing `Sendpre`, `RevealKey`, `Corrupt`, `Encrypt`, `Decrypt`, and `RegisterParty` queries.
3. At the end of the game, the adversary outputs a triple (i, s, b') and terminates.

⁶ This information is simply used to determine which party also holds asymmetric key pairs besides the shared symmetric keys.

In the following, we provide a general security definition for ACCE protocols. It will subsequently be referred to when formalizing specific definitions for ACCE protocols that provide no forward secrecy, perfect forward secrecy or asymmetric perfect forward secrecy. We have tried to keep the details of the execution environment and the definition of security close to that of JKSS. Intuitively, our security definition mainly differs from JKSS in that it considers adversaries that also have access to the new **RegisterParty** query and the extended **Corrupt** query.

Definition 4 (ACCE Security). *We say that an adversary (t, ϵ) -breaks an ACCE protocol, if $\mathcal{A}$ runs in time t , and at least one of the following two conditions holds:*

1. *When $\mathcal{A}$ terminates, then with probability at least ϵ there exists an oracle π_i^s such that*
 - π_i^s ‘accepts’ with $\text{Pid}_i^s = j$ when $\mathcal{A}$ issues its τ_0 -th query, and
 - both P_i and the intended partner P_j ⁷ are not corrupted throughout the security game and
 - π_i^s has internal state $\text{kst}_i^s = \text{fresh}$, and
 - there is no unique oracle π_j^t such that π_i^s has a matching conversation to π_j^t .

If an oracle π_i^s accepts in the above sense, then we say that π_i^s accepts maliciously.

2. *When $\mathcal{A}$ terminates and outputs a triple (i, s, b') such that*
 - π_i^s ‘accepts’ – with a unique oracle π_j^t such that π_i^s has a matching conversation to π_j^t – when $\mathcal{A}$ issues its τ_0 -th query, and
 - $\mathcal{A}$ did not issue a **RevealKey**-query to oracle π_i^s nor to π_j^t , i.e. $\text{kst}_i^s = \text{fresh}$, and
 - P_i is τ_i -corrupted and P_j is τ_j -corrupted,*then the probability that b' equals b_i^s is bounded by*

$$|\Pr[b_i^s = b'] - 1/2| \geq \epsilon.$$

If adversary $\mathcal{A}$ outputs (i, s, b') with $b' = b_i^s$ and the above conditions are met, we say that $\mathcal{A}$ answers the encryption-challenge correctly.

We say that the ACCE protocol is (t, ϵ) -secure, if there exists no adversary that (t, ϵ) -breaks it.

Let us now define security more concretely. We consider three levels of forward secrecy. We start with a basic security definition for protocols that do not provide any form of forward secrecy.

Definition 5 (ACCE Security without Forward Secrecy). *We say that an ACCE protocol is (t, ϵ) -secure without forward secrecy (NoFS), if it is (t, ϵ) -secure with respect to Definition 4 and $\tau_i = \tau_j = \infty$.*

⁷ The party P_j is not adversarially corrupted, i.e. $j \in [\ell]$. This means that P_j has not been registered by a **RegisterParty** query. Otherwise $\mathcal{A}$ may obtain all corresponding secure keys and trivially make oracle π_i^s accept.

The next definition considers PFS in the classical sense for both, client and server, as in JKSS.

Definition 6 (ACCE Security with Perfect Forward Secrecy). *We say that an ACCE protocol is (t, ϵ) -secure with perfect forward secrecy (PFS), if it is (t, ϵ) -secure with respect to Definition 4 and $\tau_i, \tau_j \geq \tau_0$.*

In the following, we provide our new definition of asymmetric perfect forward secrecy which is similar to that of classical perfect forward secrecy except that only the client is allowed to be corrupted after it has accepted.

Definition 7 (ACCE Security with Asymmetric Perfect Forward Secrecy). *We say that an ACCE protocol is (t, ϵ) -secure with asymmetric perfect forward secrecy (APFS), if it is (t, ϵ) -secure with respect to Definition 4 and it holds that $\tau_i = \infty$ and $\tau_j \geq \tau_0$ if π_i^s has internal state $\rho = \text{Server}$ or $\tau_i \geq \tau_0$ and $\tau_j = \infty$ if π_i^s has internal state $\rho = \text{Client}$.*

5 Security Analysis of Pre-shared Key Ciphersuites for Transport Layer Security

In this section, we present our results for each of the TLS-PSK ciphersuites. Due to space restrictions, the proofs are given in the full version.

Theorem 1. *Let μ be the output length of PRF_{TLS} and let λ be the length of the nonces. Assume that PRF_{TLS} is a $(t, \epsilon_{\text{PRF}})$ -secure PRF when keyed with the pre-master secret $pms := N || 0 \dots 0 || N || \text{PSK}$ or the master secret ms . Suppose the hash function H is (t, ϵ_H) -secure, and the sLHAE scheme is $(t, \epsilon_{\text{StE}})$ -secure. Then for any adversary that $(t', \epsilon_{\text{tls}})$ -breaks the TLS_PSK protocol in the sense of Definition 5 with $t \approx t'$ it holds that*

$$\epsilon_{\text{tls}} \leq (d\ell)^2 \left(\frac{1}{2^{\lambda-1}} + 3\epsilon_{\text{DPRF}} + 3\epsilon_{\text{PRF}} + 2\epsilon_H + \frac{1}{2^{\mu-1}} + 6\epsilon_{\text{StE}} \right).$$

Theorem 2. *Let μ be the output length of PRF_{TLS} and let λ be the length of the nonces. Assume that PRF_{TLS} is a $(t, \epsilon_{\text{DPRF}})$ -secure DPRF when keyed with the pre-master secret $pms := \text{len}_Z || Z || \text{len}_{\text{PSK}} || \text{PSK}$ (that consists of the pre-shared secret PSK and the Diffie-Hellman value Z). Assume that PRF_{TLS} is a $(t, \epsilon_{\text{PRF}})$ -secure PRF when keyed with the master secret ms . Suppose the hash function H is (t, ϵ_H) -secure, the DDH-problem is $(t, \epsilon_{\text{DDH}})$ -hard in the group G used to compute Z , and the sLHAE scheme is $(t, \epsilon_{\text{StE}})$ -secure. Then for any adversary that $(t', \epsilon_{\text{tls}})$ -breaks the TLS_DHE_PSK protocol in the sense of Definition 6 with $t \approx t'$ we get*

$$\epsilon_{\text{tls}} \leq (d\ell)^2 \left(\frac{1}{2^{\lambda-1}} + 3\epsilon_{\text{DPRF}} + 3\epsilon_{\text{PRF}} + 2\epsilon_H + \frac{1}{2^{\mu-1}} + \epsilon_{\text{DDH}} + 6\epsilon_{\text{StE}} \right).$$

Theorem 3. *Let μ be the output length of PRF_{TLS} and let λ be the length of the nonces. Assume that PRF_{TLS} is a $(t, \epsilon_{\text{DPRF}})$ -secure DPRF when keyed with the pre-master secret $pms := C || V || R || \text{len}_{\text{PSK}} || \text{PSK}$ (that consists of the pre-shared key*

PSK and the random key R that is exchanged between client and server). Assume that PRF_{TLS} is a $(t, \epsilon_{\text{PRF}})$ -secure PRF when keyed with the master secret ms . Suppose the hash function H is (t, ϵ_H) -secure, the public key encryption scheme PKE is $(t, \epsilon_{\text{PKE}})$ -secure (IND-CCA). Suppose that the sLHAE scheme is $(t, \epsilon_{\text{StE}})$ -secure. Then for any adversary that $(t', \epsilon_{\text{tls}})$ -breaks the TLS_RSA_PSK protocol (where the key transport mechanism is implemented via PKE) in the sense of Definition 7 with $t \approx t'$ it holds that

$$\epsilon_{\text{tls}} \leq (d\ell)^2 \left(\frac{1}{2^{\lambda-1}} + \epsilon_{\text{PKE}} + 3\epsilon_{\text{DPRF}} + 3\epsilon_{\text{PRF}} + 2\epsilon_H + \frac{1}{2^{\mu-1}} + 6\epsilon_{\text{StE}} \right).$$

TECHNICAL OVERVIEW OF THE SECURITY PROOFS. At a high level, the security proofs are similar to that of JKSS. From a technical standpoint, the security proof of TLS_PSK is simpler than that of the classical ciphersuites of TLS as security only relies on the secrecy of the pre-shared secrets. Roughly, in the proofs of the classical TLS ciphersuites one additionally has to establish that the key exchange mechanism produces a shared secret in the first place. To prove TLS_RSA_PSK and TLS_DHE_PSK we exploit the DPRF-security of PRF_{TLS} . The challenge is to show that the master secret is indistinguishable from random although the adversary may reveal the pre-shared secret or a freshly generated ephemeral secret. Intuitively, if only one of these values remains unrevealed by the adversary, then at least one input key to the DPRF PRF_{TLS} is (indistinguishable from) random. Therefore, PRF_{TLS} computes a random-looking master secret which in turn can be used to derive secure application keys.

Acknowledgements. We would like to thank Kenny Paterson and the anonymous referees for their valuable comments and suggestions.

References

1. Badra, M., Urien, P.: Toward SSL integration in SIM smartcards. In: WCNC, pp. 889–893. IEEE (2004)
2. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM Conference on Computer and Communications Security, pp. 62–73. ACM (1993)
3. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)
4. Blake-Wilson, S., Johnson, D., Menezes, A.: Key agreement protocols and their security analysis. In: Darnell, M. (ed.) Cryptography and Coding 1997. LNCS, vol. 1355, pp. 30–45. Springer, Heidelberg (1997)
5. Blake-Wilson, S., Menezes, A.: Unknown key-share attacks on the Station-to-Station (STS) protocol. In: Imai, H., Zheng, Y. (eds.) PKC 1999. LNCS, vol. 1560, pp. 154–170. Springer, Heidelberg (1999)
6. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998)

7. BouncyCastle Software Developers. Bouncy Castle Crypto APIs (2013), <http://www.bouncycastle.org/>
8. 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)
9. Chen, C., Tang, S., Mitchell, C.J.: Building general-purpose security services on EMV payment cards. In: Keromytis, A.D., Di Pietro, R. (eds.) SecureComm 2012. LNCS, vol. 106, pp. 29–44. Springer, Heidelberg (2013)
10. Dacosta, I., Ahamad, M., Traynor, P.: Trust no one else: Detecting MITM attacks against SSL/TLS without third-parties. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 199–216. Springer, Heidelberg (2012)
11. Dierks, T., Allen, C.: The TLS Protocol Version 1.0. RFC 2246 (Proposed Standard). Obsoleted by RFC 4346, updated by RFCs 3546, 5746 (January 1999)
12. Dierks, T., Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.1. RFC 4346 (Proposed Standard). Obsoleted by RFC 5246, updated by RFCs 4366, 4680, 4681, 5746 (April 2006)
13. Dierks, T., Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (Proposed Standard). Updated by RFCs 5746, 5878 (August 2008)
14. Eronen, P., Tschofenig, H.: Pre-Shared Key Ciphersuites for Transport Layer Security (TLS). RFC 4279 (Proposed Standard) (December 2005)
15. German Federal Office for Information Security (BSI). TR-03112, Das eCard-API-Framework (2005), https://www.bsi.bund.de/ContentBSI/Publikationen/TechnischeRichtlinien/tr03112/index_htm.html
16. Gajek, S., Manulis, M., Pereira, O., Sadeghi, A.-R., Schwenk, J.: Universally Composable Security Analysis of TLS. In: Baek, J., Bao, F., Chen, K., Lai, X. (eds.) ProvSec 2008. LNCS, vol. 5324, pp. 313–327. Springer, Heidelberg (2008)
17. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 273–293. Springer, Heidelberg (2012)
18. Jonsson, J., Kaliski Jr., B.S.: On the security of RSA encryption in TLS. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 127–142. Springer, Heidelberg (2002)
19. Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DH and TLS-RSA in the standard model. IACR Cryptology ePrint Archive, 2013:367 (2013)
20. Kohlar, F., Schwenk, J., Jensen, M., Gajek, S.: Secure bindings of SAML assertions to TLS sessions. In: ARES, pp. 62–69. IEEE Computer Society (2010)
21. 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)
22. Krawczyk, H., Paterson, K.G., Wee, H.: On the security of the TLS protocol: A systematic analysis. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 429–448. Springer, Heidelberg (2013)
23. 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, Heidelberg (2007)
24. Mavrogianopoulos, N., Josefsson, S.: The GnuTLS Transport Layer Security library, <http://gnutls.org> (last updated March 22, 2013)
25. Menezes, A., Smart, N.P.: Security of signature schemes in a multi-user setting. Des. Codes Cryptography 33(3), 261–274 (2004)
26. Meyer, C., Schwenk, J.: Lessons learned from previous SSL/TLS attacks - a brief chronology of attacks and weaknesses. IACR Cryptology ePrint Archive, 2013:49 (2013)

27. Morrissey, P., Smart, N.P., Warinschi, B.: The TLS handshake protocol: A modular analysis. *Journal of Cryptology* 23(2), 187–223 (2010)
28. OpenSSL. The OpenSSL project (2013), <http://www.openssl.org>
29. Urien, L.C.P., Martin, P.: EMV support for TLS-PSK. draft-urien-tls-psk-emv-02 (February 2011)
30. Paterson, K.G., Ristenpart, T., Shrimpton, T.: Tag size *does* matter: Attacks and proofs for the TLS record protocol. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 372–389. Springer, Heidelberg (2011)
31. Urien, P.: Introducing TLS-PSK authentication for EMV devices. In: Smari, W.W., McQuay, W.K. (eds.) CTS, pp. 371–377. IEEE (2010)