

Secure Modular Password Authentication for the Web Using Channel Bindings

Mark Manulis¹, Douglas Stebila², and Nick Denham²

¹ Surrey Centre for Cyber Security, University of Surrey, UK
mark@manulis.eu

² Queensland University of Technology, Brisbane, Australia
{stebila,n.denham}@qut.edu.au

Abstract. Secure protocols for password-based user authentication are well-studied in the cryptographic literature but have failed to see widespread adoption on the Internet; most proposals to date require extensive modifications to the Transport Layer Security (TLS) protocol, making deployment challenging. Recently, a few modular designs have been proposed in which a cryptographically secure password-based mutual authentication protocol is run inside a confidential (but not necessarily authenticated) channel such as TLS; the password protocol is bound to the established channel to prevent active attacks. Such protocols are useful in practice for a variety of reasons: security no longer relies on users' ability to validate server certificates and can potentially be implemented with no modifications to the secure channel protocol library.

We provide a systematic study of such authentication protocols. Building on recent advances in modelling TLS, we give a formal definition of the intended security goal, which we call *password-authenticated and confidential channel establishment* (PACCE). We show generically that combining a secure channel protocol, such as TLS, with a password authentication protocol, where the two protocols are bound together using either the transcript of the secure channel's handshake or the server's certificate, results in a secure PACCE protocol. Our prototype based on TLS is available as a cross-platform client-side Firefox browser extension and a server-side web application which can easily be installed on deployed web browsers and servers.

Keywords: password authentication, Transport Layer Security, channel binding.

1 Introduction

Authentication using passwords is perhaps the most prominent and human-friendly user authentication mechanism widely deployed on the Web. In this ubiquitous approach, which we refer to as *HTML-forms-over-TLS*, the user's password is sent encrypted over an established server-authenticated Transport Layer Security (TLS, previously known as Secure Sockets Layer (SSL)) channel in response to a received HTML form. This approach is subject to many threats:

the main problems with this technique are that security fully relies on a functional X.509 public key infrastructure (PKI) and on users correctly validating the server's X.509 certificate. In practice, these assumptions are unreliable due to a variety of reasons: the many reported problems with the trustworthiness of certification authorities (CAs), inadequate deployment of certificate revocation checking, ongoing threats from phishing attacks, and the poor ability of the users to understand and validate certificates [1,2]. Hypertext Transport Protocol (HTTP) basic and digest access authentication [3] has been standardized, and digest authentication offers limited protection for passwords, but usage is rare. Public-key authentication of users, e.g. using X.509 certificates, is also rare.

1.1 Password-Authenticated Key Exchange (PAKE)

Password-authenticated key exchange (PAKE) protocols, which were introduced by Bellovin and Merritt [4], and the security of which was formalized in several settings [5,6,7], could mitigate many of the risks of the HTML-forms-over-TLS approach as they do not rely on any PKI and offer stronger protection for client passwords against server impersonation attacks, such as phishing. PAKE protocols allow two parties determine whether they both know a particular string while cryptographically hiding any information about the string. They are resistant to offline-dictionary attacks: an adversary who observes or participates in the protocol cannot test many passwords against the transcript. Successful execution of a PAKE protocol also provides parties with secure session keys which can be used for encryption.

Despite the many benefits of PAKE, and the presence of a variety of existing protocols in the academic literature and in standards [8,9,10], PAKE-based approaches for client authentication have not been adopted in practice. There is no PAKE standard that has been agreed upon and implemented in existing web browser and server technologies. This is due to several practical obstacles, including: patents covering PAKE in general (some of which have recently expired in the US), patents on proposed standards such as the Secure Remote Password (SRP) protocol [11], lack of agreement on the appropriate layer within the networking stack for the integration of PAKE [12], complexity of backwards-compatible deployment with TLS, and user-interface challenges.

There have been a few proposals to integrate PAKE into TLS by adding password-based ciphersuites as an alternative to public-key authenticated ciphersuites. SRP has been standardized as a TLS ciphersuite [13] and has several reference implementations but none in major web browsers or servers. Abdalla et al. [14] proposed the provably secure Simple Open Key Exchange (SOKE) ciphersuite, which uses a variant of the PAKE protocol from [15] that is part of the IEEE-P1363.2 standard [10]. The J-PAKE protocol [16] is used in a few custom applications. Common to all PAKE ciphersuite approaches is that the execution of PAKE becomes part of the TLS handshake protocol: the key output by PAKE is treated as the TLS pre-master secret, which is then used to derive further encryption keys according to the TLS specification. An advantage of this approach is that secure password authentication could subsequently be used in

any application that makes use of TLS, and that standard TLS mechanisms for key derivation and secure record-layer communication can continue to be used. However, a major disadvantage is that any new ciphersuites in TLS require substantial vendor-side modifications of the web browser and server software. This is problematic for modern web server application architectures within large organizations, where a TLS accelerator immediately handles the TLS handshake and encryption, then hands the plaintext off to the first of many application servers; requiring the TLS accelerator to have access to the list of valid usernames and passwords may mean a substantial re-architecting. Moreover, using solely PAKE in TLS means abandoning the web public key infrastructure.

1.2 Running PAKE at the Application Layer

A better approach for realizing secure password-based authentication on the web may be to rely on existing TLS implementations to provide confidential communication between clients and servers, and integrate application-level PAKE for password-based authentication, without requiring any modifications to the TLS specification or implementation; in particular, without proposing any new TLS ciphersuites or changing any of the steps of TLS handshake protocol.

However, if the TLS channel is only assumed to provide confidentiality, not authentication, then one must use an alternative mechanism to rule out man-in-the-middle attacks on the TLS channel. Since it is the password-based protocol that provides mutual authentication, there should be a binding between the TLS channel and the password-based protocol. There are several potential values which might be used for binding: the transcript of the TLS handshake protocol, the TLS master secret key (or a value derived from it, such as the TLS `Finished` message), or even the server's certificate. A recent standard [17] describes three TLS channel bindings, two of which are relevant to us: `tls-unique` in which the binding string is the `Finished` message, and `tls-server-end-point` in which the binding string is the hash of the server's certificate. Notably, TLS channel bindings do not change the TLS protocol itself: all TLS protocol messages, ciphersuites, data transmitted, and all other values are entirely unchanged. Rather, TLS channel bindings expose an additional value to the application that can be obtained locally, thereby requiring minimal changes to TLS implementations.

The high-level approach of running PAKE at the application level is given in Figure 1. Using PAKE at the application-level supplements, rather than replaces, the use of public key certificates.

Several recent works have proposed protocols of this form. Oiwa et al. [18,19,20,21] published an Internet-Draft that employs an ISO-standardized PAKE protocol (KAM3 [8, §6.3],[22]) and binds it to the TLS channel using either the server's certificate or the TLS master secret key, but no formal justification is given for security of the combined construction.

Dacosta et al. [23] proposed the DVCert protocol which aims to achieve direct validation of the TLS server certificates by using a modification of the protocol from [24] for secure server-to-client password-based authentication. Dacosta et al. used an automated cryptographic protocol verifier, ProVerif, to demonstrate

that their protocol does not leak password information without addressing any further security properties of secure channels. In particular, the analysis carried out in [23] is insufficient for showing the actual benefit of using PAKE protocols to strengthen the security of the TLS channel: Dacosta et al. simply show that the protocol does not leak password information, which is not surprising since they build on a PAKE protocol which already did not leak information. Rather, the security goals of the overall channel must be fully modelled to provide a complete analysis. This is where our model for PACCE fills the gap: it explains the expected security goals from the combination of PAKE and TLS. While we have no reason to expect that the protocol from [23] is insecure, our model would enable such an analysis of that protocol.

Outside of the world of TLS, the Off-the-Record Messaging protocol [25,26] uses a PAKE-like password authentication functionality that is based on secure two-party computation techniques to authenticate the long-term public keys used in establishing confidential (yet deniable) channel, but with no justification for the security of the combined construction.

1.3 Contributions

We analyze the modular approach to secure password authentication on the web, in which a secure channel protocol such as TLS is combined—in a black box way—with a password-authentication protocol. The black-box approach enables smooth integration of PAKE functionality with secure channels such as TLS without requiring any modification to the original channel protocol specification, nor requiring abandoning public key certificates for server authentication.

At a high level, in our approach, a normal secure channel is established with no assumptions on correct validation of certificates; then a PAKE protocol is run *within* the secure channel to demonstrate (a) mutual knowledge of the password, and (b) absence of a man-in-the-middle attack on the channel; once the PAKE succeeds, the parties continue application communication. Notably, the session key established by the PAKE is not used for symmetric encryption; while it *could* replace the secure channel session key, in practice, such as in TLS, there is no standardized mechanism to do so. Moreover, using the existing channel session key is fine provided we bind the execution of the PAKE and the original channel establishment in a provably secure way. We show how to realize this binding using available TLS standards.

Our formal approach is as follows.

First, we define formally in Section 2 the security property we aim for in our protocol: since the most suitable definition for the security of the combined TLS handshake and record layer protocols is the *authenticated and confidential channel establishment (ACCE)* model of Jager et al. [27], we give a corresponding *password-based ACCE (PACCE)* notion. We apply our PACCE model to analyzing strong passwords in TLS; as ACCE has been used to analyze many real world protocols (SSH [28], EMV [29], and QUIC), PACCE should be suitable for strong password variants of those.

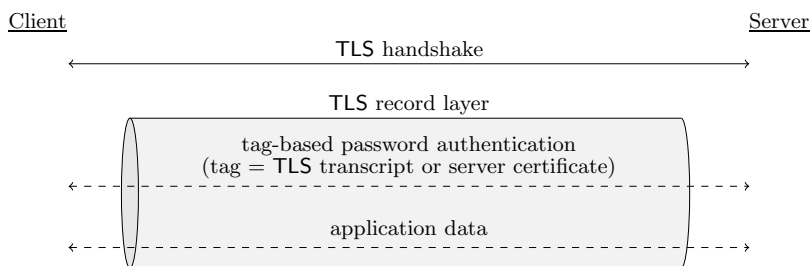


Fig. 1. High-level approach for combining secure password authentication with a TLS channel to establish a single password-authenticated secure channel

Next, we define the various primitives we employ to achieve this goal; these include the original ACCE notion, as well as an unauthenticated *confidential channel establishment (CCE)* protocol. Rather than using a PAKE protocol, we actually employ a newly defined *tag-based password authentication (tPAuth)* protocol, which provides mutual authentication based on knowledge of a shared password, with acceptance only if both parties input the same, possibly public, *tag* to the protocol.

Then, when we run the tPAuth protocol inside the established confidential channel, we bind the two protocols together by setting the tag to be either the transcript of the channel establishment or the long-term public key used by the server in the channel establishment. We prove in Sections 4 and 5 that both of these tags are sufficient to achieve the end result of a password-based authenticated and confidential channel establishment (PACCE) protocol. Our two security theorems provide a qualitative distinction between those two tags as binding mechanisms because they use slightly different assumptions on the confidential channel establishment protocol: when the tag used is the transcript, it suffices to use a CCE protocol, but when the tag used is the server’s public key, we require the underlying protocol to be an ACCE protocol. (These results give the first security justification of two standardized TLS channel binding mechanisms, `tls-unique` and `tls-server-end-point` [17].)

Applicability of results to other approaches. Our results, by employing the CCE/ACCE frameworks, are generic and could be applicable to constructions employing a wide variety of protocols, not just TLS. These results justify the *general approach* of the recent proposals of Oiwa et al.’s protocol [18,19,20,21] and Dacosta et al.’s DVCert protocol [23]. We caution that our theorems do not immediately imply security of those particular protocols for several reasons.

Our theorems depend on a formal construction called a tag-based password authentication (tPAuth) protocol. Only the tSOKE protocol we describe in Appendix B is known to be a secure tPAuth. For most PAKE protocols, such as the PAK protocol employed by Dacosta et al., it seems not hard to modify the PAKE security proof to demonstrate tPAuth, but this requires additional work for each particular protocol.

Oiwa et al. provide three channel binding mechanisms: the server’s TLS certificate, the server host string (e.g., `http://www.example.com:80`), and the server’s TLS master secret key.

- Our results do not address the cryptographic security of server host string binding; indeed, Oiwa et al. intend for this mode mainly to be used when TLS is not employed.
- Section 5.3 provides a justification for the use of the TLS server certificate, though as noted our results only apply when the PAKE is also a tPAuth.
- Section 4.4 provides a justification for the use of the TLS `Finished` message, whereas Oiwa et al. allow use of the TLS master secret key. Since the TLS handshake will not complete unless both parties compute the same `Finished` messages, our results also justify the use of the TLS master secret key for channel binding, though security of the whole construction again is not implied by our results due to the caveat above that our results only apply with the PAKE is also shown to be a tPAuth.

Dacosta et al.’s DV Cert protocol only provides server-to-client password-based authentication, whereas our PACCE notion provides mutual authentication.

Implementation. In Section 6, we demonstrate the practical merit of our approach with a reference implementation for the Mozilla Firefox web browser and the Apache web server that takes advantage of the modularity of the construction.

- On the client side, our implementation is achieved entirely as a Firefox *extension*: it is a cross-platform Javascript-based bundle that can be installed by the user at run-time, without any modifications to the source code of the Firefox browser or its TLS library, Network Security Services (NSS).
- On the server side, our implementation is achieved entirely as a cross-platform PHP application: it can be added at run-time without any modifications to the source code of the Apache web server or its TLS library, OpenSSL.

Binding the password authentication protocol to the TLS channel is achieved using the server’s TLS certificate as the tag; both Firefox and Apache have APIs exposing the server’s certificate to the extension and PHP application, respectively; the server certificate is one of two channel binding mechanisms standardized for TLS [17]. The source code size of our implementation is quite small. Our pure Javascript implementation is completely cross-platform and provides tolerable performance, with total round-trip time under half a second on a laptop, while our native C implementation using OpenSSL libraries provides high client and server performance with a total protocol execution time, including network latency in a corporate network, around 109ms. Our implementation is available for immediate download.

Though the ultimate goal of this line of work would be for such a protocol to be built into the browser, our Javascript extension may be amenable to gradual deployment to seed adoption while still achieving good performance, especially when making use of native libraries.

2 Password-Authenticated Confidential Channels

The security goal for our main construction is that it be a secure *password-authenticated and confidential channel establishment (PACCE)* protocol, which is a new password-based variant of the ACCE model of Jager et al. [27]. ACCE seems to be the most suitable for describing the security requirements of real-world secure channel protocols such as TLS [27,30,31,32] and SSH [28], and so it is natural to adapt it to the password setting.

A PACCE protocol is a two-party protocol that proceeds in two stages: in the *handshake* stage both participants perform an initial cryptographic handshake to establish session keys which are then used in the *record layer* stage to authenticate and encrypt the transmitted session data.¹ At some time during execution, the parties may accept the session as being legitimately authenticated, or reject. The main difference in PACCE compared to the original ACCE model is the use of passwords instead of long-term public keys for authentication.

At a high level, a PACCE protocol is secure if the adversary cannot break authentication, meaning it cannot cause a party to accept without having interacted with its intended partner, and cannot break the confidential channel, meaning it cannot read or inject ciphertexts.

We consider the standard client-server communication model where a party is either a *client* C or a *server* S . For each client-server pair (C, S) there exists a corresponding password $\text{pw}_{C,S}$ drawn from a dictionary \mathcal{D} .

An instance of party $U \in \{C, S\}$ in a session s is denoted as Π_U^s . Each instance Π_U^s records several variables:

- $\Pi_U^s.\text{pid}$: the *partner identity* with which Π_U^s believes to be interacting in the protocol session.
- $\Pi_U^s.\rho \in \{\text{init}, \text{resp}\}$: the *role* of this instance in the session, either initiator or responder. $\text{init}(\Pi_U^s)$ and $\text{resp}(\Pi_U^s)$ denote Π_U^s 's view of who the initiator and responder are in the session, namely $(U, \Pi_U^s.\text{pid})$ when $\Pi_U^s.\rho = \text{init}$, and $(\Pi_U^s.\text{pid}, U)$ when $\Pi_U^s.\rho = \text{resp}$.
- $\Pi_U^s.T$: a *transcript* composed of all messages sent and received by the instance in temporal order.
- $\Pi_U^s.\alpha \in \{\text{active}, \text{accept}, \text{reject}\}$: the *status* of this instance.
- $\Pi_U^s.k$: the *session key* computed by this stage; initially set to empty \emptyset ; when non-empty, it consists of two symmetric keys $\Pi_U^s.k^{\text{enc}}$ and $\Pi_U^s.k^{\text{dec}}$ for encryption and decryption with some stateful length-hiding authenticated encryption scheme [27] used to provide confidentiality in the record layer stage. If $\Pi_U^s.\alpha = \text{accept}$, then $\Pi_U^s.k \neq \emptyset$
- $\Pi_U^s.b \in \{0, 1\}$: a randomly sampled bit used in the **Encrypt** oracle.

¹ In the original ACCE model, these stages were called the *pre-accept* and *post-accept* stages respectively. In PACCE, the parties may start sending encrypted data before accepting, so we have renamed the stages to *handshake* and *record layer*, which is suggestive of TLS, but of course can be used to model any appropriate protocol.

Two instances Π_U^s and $\Pi_{U'}^s$ are said to be *partnered* if and only if $\Pi_U^s.\text{pid} = U'$, $\Pi_{U'}^s.\text{pid} = U$, $\Pi_U^s.\rho \neq \Pi_{U'}^s.\rho$, and their transcripts form *matching conversations* [27], denoted $\Pi_U^s.T \approx \Pi_{U'}^s.T$.

The adversary \mathcal{A} controls all communications and can interact with parties using certain oracle queries. Normal operation of the protocol is modelled by the following queries:

- $\text{Send}^{\text{pre}}(\Pi_U^s, m)$: This query is answered as long as $\Pi_U^s.k = \emptyset$. In response the incoming message m is processed by Π_U^s and any outgoing message which is generated as a result of this processing is given to \mathcal{A} . Special messages $m = (\text{init}, U')$ and $m = (\text{resp}, U')$ are used to initialize the instance as initiator or responder, respectively, and to specify the identity of the intended partner U' . Note that processing of m may eventually lead to the end of the handshake stage, in which case Π_U^s either computes $\Pi_U^s.k$ and switches to the record layer stage or terminates with a failure.
- $\text{Encrypt}(\Pi_U^s, m_0, m_1, \text{len}, \text{head})$,
- $\text{Decrypt}(\Pi_U^s, C, \text{head})$: The Encrypt and Decrypt queries proceed as defined in the original ACCE definition [27] and are omitted due to space restrictions. Note that, compared with ACCE, we allow protocol messages to be sent on the encrypted channel: If $\Pi_U^s.\alpha = \text{active}$, then the returned plaintext message m is processed as a protocol message; the resulting outgoing message m' is encrypted using $\text{Encrypt}(\Pi_U^s, m', m', \text{len}(m'), \text{head})$ and the resulting ciphertext C is returned to \mathcal{A} . Otherwise, when $\Pi_U^s.\alpha = \text{accept}$, the output of Decrypt is returned to \mathcal{A} .

Furthermore, the adversary may obtain some secret information:

- $\text{RevealSK}(\Pi_U^s)$: Return $\Pi_U^s.k$.
- $\text{Corrupt}(C, S)$: Return $\text{pw}_{C,S}$.

Note that, compared with AKE models like the eCK model [33] that use public key authentication, password-based protocols cannot tolerate ephemeral key leakage while maintaining resistance to offline dictionary attacks, hence we do not include an ephemeral key leakage query.

Definition 1 (PACCE security). *An adversary \mathcal{A} is said to (t, ϵ) -break a PACCE protocol if \mathcal{A} runs in time t and at least one of the following two conditions hold:*

1. \mathcal{A} breaks authentication: *When \mathcal{A} terminates, then with probability at least $\epsilon + O(n/|\mathcal{D}|)$ where n is the number of initialized PACCE instances there exists an instance Π_U^s such that*
 - (a) $\Pi_U^s.\alpha = \text{accept}$, and
 - (b) \mathcal{A} did not issue $\text{Corrupt}(\text{init}(\Pi_U^s), \text{resp}(\Pi_U^s))$ before Π_U^s accepted, and
 - (c) \mathcal{A} did not issue $\text{RevealSK}(\Pi_U^s)$ or $\text{RevealSK}(\Pi_{U'}^s)$ for any $\Pi_{U'}^s$ that is partnered to Π_U^s , and
 - (d) there is no unique instance $\Pi_{U'}^s$ that is partnered to Π_U^s .

2. \mathcal{A} breaks authenticated encryption: When \mathcal{A} terminates and outputs a triple (U, s, b') such that conditions (a)–(c) from above hold, then we have that $|\Pr[b' = \Pi_U^s \cdot b] - \frac{1}{2}| \geq \epsilon + O(n/|\mathcal{D}|)$.

A PACCE protocol is (t, ϵ) -secure if there is no \mathcal{A} that (t, ϵ) -breaks it; it is secure if it is (t, ϵ) -secure for all polynomial t and negligible ϵ in security parameter κ .

Observe that Definition 1 accounts for online dictionary attacks against PACCE protocols by using a lower bound $\epsilon + O(n/|\mathcal{D}|)$ for the adversarial success probability, which models \mathcal{A} 's ability to test at most one password (or a constant number) from the dictionary \mathcal{D} in a single session.

3 Generic Construction

Our generic construction for secure PACCE between a client C and a server S from password-based authentication and a secure channel is as follows. First, the channel establishment protocol (CCE or ACCE) is run until it accepts. Then, using the secure channel, the two parties run a tag-based password authentication (tPAAuth) protocol where the tag is a binding value from the secure channel; when the tPAAuth protocol accepts, then the parties accept in the overall PACCE protocol, and then continue to use the channel for communication. Our two constructions in Sections 4 and 5 differ only in the way the tPAAuth is bound to the established channel: the tag is either the transcript of the channel establishment protocol or the long-term public key of the server.

We now concretely describe the combination of a tPAAuth protocol with TLS, as detailed in Figure 2. C and S first build up a standard TLS channel: that is, they execute a normal TLS handshake, then exchange `ChangeCipherSpec` messages to start authenticated encryption within the TLS record layer, and then exchange their `Finished` messages for explicit key confirmation. Once `Finished` messages are successfully exchanged, the parties continue using the authenticated encryption mechanism of the TLS record layer to communicate messages of the tag-based password-authentication protocol tPAAuth. In our first construction this binding is achieved by using the `Finished` messages as the tag; note that `Finished` messages depend on the (hash of the) entire TLS handshake transcript. In our second construction the tag is the server's certificate (which includes the server's public key) that was communicated by S in its `Certificate` message of the TLS handshake. Upon successful completion of the password authentication phase both parties continue using session keys and authenticated encryption mechanism of the established TLS channel for secure communication.

4 Construction #1: Binding Using CCE Transcript

Our first generic PACCE protocol $\Gamma_T := \Gamma_T(\pi, \xi)$ is constructed as in Section 3 from a confidential channel establishment (CCE) protocol π and a tag-based password authentication (tPAAuth) protocol ξ where the tag τ used is the transcript T of the CCE handshake stage. We will see that, because we are using

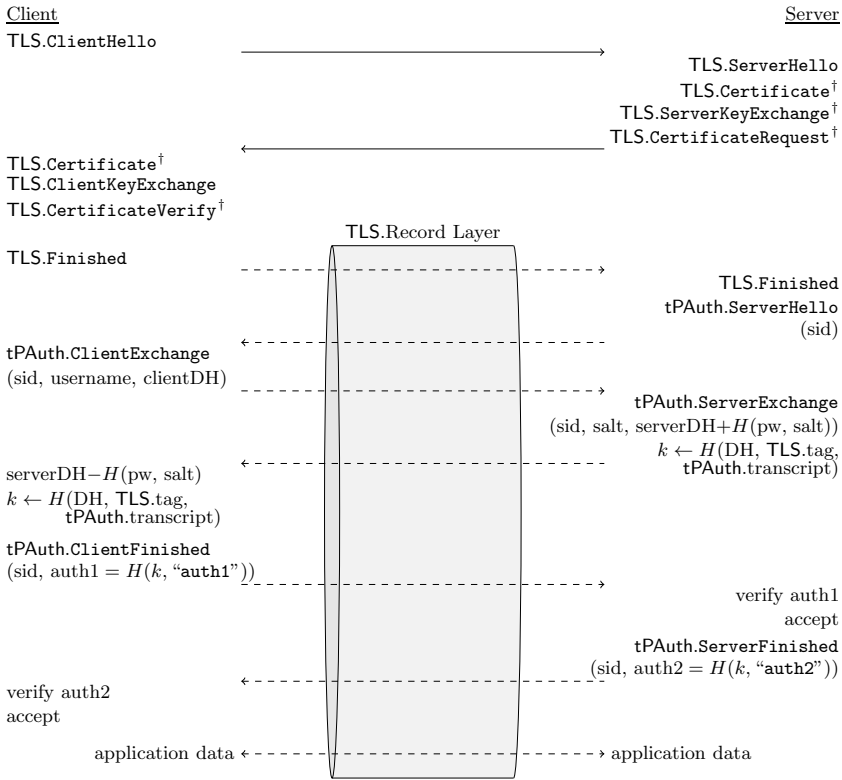


Fig. 2. Protocol message diagram for TLS with tunnelled tag-based password authentication, with example messages for the tSOKE protocol. TLS.tag is either TLS.Finished or TLS.Certificate . \dagger denotes optional messages.

the full transcript from the channel establishment to bind the two protocols together, we need not rely on any authenticity properties of the channel, and thus can use a CCE protocol, not an ACCE protocol.

4.1 Building Block: CCE

As a building block in our analysis we use the notion of *confidential channel establishment (CCE)* that differs from (P)ACCE in that it is supposed to guarantee only confidentiality (and integrity) of the established channel, but not authentication of partners; hence, security is only assured for sessions in which the adversary \mathcal{A} remains passive during the handshake stage. We thus model CCE by slightly modifying the PACCE model from Section 2. The differences are explained in Appendix A. Every secure (P)ACCE protocol is also CCE-secure: if we ignore the authentication aspects, then we still get confidential channel establishment in sessions where the adversary is passive during the handshake.

TLS does, when public keys are managed and used properly, provide strong authentication based on public keys, and can be proven ACCE-secure [27]. But,

as we observed in the introduction, practice suggests we cannot rely on the web PKI to provide ideal authentic distribution and mapping of public keys to identities. TLS can in practice be seen as a CCE protocol: even though long-term public keys may be used in TLS, we are not confident in their authenticity, so we only take TLS to provide CCE security.

4.2 Building Block: Tag-Based Password Authentication

Tag-based authentication [34] accounts for the use of auxiliary, possibly public, strings (tags) in authentication protocols — each party uses a tag, in addition to the authentication factor, and the protocol guarantees that if parties accept then their tags match. This concept was introduced in [34] for public key-based authentication protocols and then generalized in [35] for other types of authentication factors, including passwords and biometrics. In our analysis we will use a *tag-based password authentication* protocol, denoted tPAuth.

The model of tPAuth can be described using the setting of PACCE protocols from Section 2. A tPAuth session is executed between a client instance Π_C^s and a server instance $\Pi_S^{s'}$ on input the corresponding password $\text{pw}_{C,S}$ from the dictionary \mathcal{D} and some tag $\tau \in \{0, 1\}^*$. A tPAuth session is successful if both instances use the same password $\text{pw}_{C,S}$ and tag τ as their input. The security of tPAuth protocols, extends the traditional password authentication requirement that accounts for online dictionary attacks with the requirement of tag equality, and is achieved formally by an extended definition of partnering. The details of tPAuth security appear in Appendix B.

In Appendix B, we present tSOKE, a tag-based variant of the Simple Open Key Exchange (SOKE) variant [14]. An example of the tSOKE message flow can be seen in Figure 2. Since SOKE is a password-authenticated key exchange protocol, it does establish a secure session key; however, we only use its properties of *mutual authentication* and *resistance to dictionary attacks*. The tag is inserted into the key used for explicit key confirmation; it does not need resistance to dictionary attacks since it is known to the adversary, and thus is not used in the Diffie–Hellman portion like the password.

4.3 Security Analysis of Construction #1

Theorem 1 (CCE + tPAuth $_{\tau=T_{\text{CCE}}} \implies \text{PACCE}$). *The generic construction of a PACCE protocol $\Gamma_T(\pi, \xi)$ from a CCE protocol π and a tPAuth protocol ξ , with the tag equal to the transcript T_{CCE} from the CCE handshake stage, is PACCE-secure, assuming the underlying protocols are secure.*

The proof consists of a sequence of games. In the first game, the simulator continues to simulate the CCE portion of the protocol but undetectably replaces the tPAuth simulation with that of a real tPAuth challenger. Next, the simulator aborts if any of its instances accept without a partnered instance existing; this will correspond to a violation of authentication in the underlying tPAuth challenger. In the third game, the simulator now undetectably replaces the CCE

simulation with that of a real CCE challenger. An adversary who can win against the resulting PACCE simulator can be used to win against the underlying CCE challenger. The proof appears in the full version [36].

4.4 Using `tls-unique` Channel Binding

The `tls-unique` channel binding mechanism [17] can be used to instantiate construction #1. For `tls-unique` the channel binding string is the first `Finished` message, which is the output of a pseudorandom function on the hash of the TLS handshake transcript.

It is straightforward to see that if the `Finished` message is used as the tag for channel binding instead of the full transcript in an analogous generic construction $\Gamma_{fin}(\pi, \xi)$, and the hash function H is collision-resistant and the pseudorandom function PRF is secure, then Γ_{fin} is a secure PACCE. This follows by noting that, except with negligible probability, the parties must use the same transcript in order to arrive at the same tag, and then the proof of Theorem 1 applies.

5 Construction #2: Binding Using Server Public Key

Our second generic PACCE protocol $\Gamma_{pk} := \Gamma_{pk}(\pi, \xi)$ is constructed as in Section 3 from an authenticated and confidential channel establishment (ACCE) protocol π and a `tPAAuth` protocol ξ where the tag τ used is the long-term public key used by the server in the ACCE protocol.

Because we are using only the server’s long-term public key, and not the full transcript from the channel establishment, to bind the two protocols together, we now must rely on some authenticity properties of the channel. However, we will not be relying on *users* to correctly validate the server’s public key or decide which long-term server public key corresponds with which password: from the external perspective, the protocol is still a PACCE protocol, with authentication only coming from passwords, not from long-term server public keys.

5.1 Building Block: ACCE (With Key Registration)

ACCE [27] is currently the most complete model for the security properties of the core TLS protocol. We use a variant of ACCE as a building block in our generic construction. The first variation is that we allow for either server-only or mutual authentication, as in Giesen et al. [32]; when server-only authentication is used, only client instances are legitimate targets for breaking authentication. The second variation is in how static public keys are distributed. In typical AKE and ACCE models, it is simply assumed that parties have authentic copies of all static public keys, abstracting the problem away. Since we will use ACCE as a building block under the assumption that the “static” public keys are not to be trusted as authentic, we allow the adversary to cause any public key to be accepted as a static public key using a `Register` query; only sessions where the key is not an adversary-registered key are legitimate targets for breaking. The formal differences between the ACCE model with key registration and the original ACCE model appear in the full version [36].

5.2 Security Analysis of Construction #2

Theorem 2 ($\text{ACCE} + \text{tPAAuth}_{\tau=pk} \implies \text{PACCE}$). *The generic construction of a PACCE protocol $\Gamma_T(\pi, \xi)$ from an ACCE protocol π and a tPAAuth protocol ξ , with tag set to the server’s public key pk from the ACCE handshake stage, is PACCE-secure, assuming the underlying protocols are secure.*

In the first game, the simulator simulates the ACCE portion of the protocol and undetectably replaces the tPAAuth portion using messages that it obtains from a real tPAAuth challenger. Next, the simulator aborts if any of its instances accept without an instance whose tPAAuth transcripts and the input tags match, which corresponds to an attack against the tPAAuth protocol. In the third game, the simulator will use messages obtained from a real ACCE challenger such that it can use any adversary who wins against the resulting PACCE simulator to break the security of the ACCE protocol; when the adversary uses its own long-term public keys (which is allowed since they are not authenticated), we use the key registration functionality of the (modified) ACCE challenger. Due to space constraints, the proof appears in the full version [36].

5.3 Using `tls-server-end-point` Channel Binding

The `tls-server-end-point` channel binding mechanism [17] can be used to instantiate construction #2. For `tls-server-end-point` the channel binding string is the hash of the server’s X.509 certificate. Note that the certificate contains the server’s public key as a canonically identifiable substring.

It is straightforward to see that if the hash of the certificate is used as the tag for channel binding instead of the raw public key in an analogous generic construction $\Gamma_{cert}(\pi, \xi)$, and the hash function is second-preimage-resistant, then Γ_{cert} is a secure PACCE. This follows by noting that an active adversary must use a certificate that hashes to the same value as the server’s certificate.

6 Implementation

As an important motivation for our modular protocol design was the ability to modularly implement the protocol, we produced a prototype to demonstrate this.

For the tag-based password authentication protocol, we propose tSOKE, a tag-based version of the SOKE protocol [14], a highly efficient Diffie–Hellman based PAKE; see Appendix B for details of tSOKE. We used the NIST P-192 elliptic curve group [37].

We implemented the client side of the protocol as an extension for Mozilla Firefox (320 lines of custom Javascript, plus libraries), and the server side of the protocol as a PHP application (210 lines of custom PHP code, plus libraries) on an Apache web server. No modifications to the source code of the underlying web browser (Firefox) or the underlying web server (Apache with OpenSSL)

were required—in particular, we did not have to alter the SSL/TLS implementation and we did not have to recompile Firefox or Apache. Since the mechanism that the server code uses to obtain the certificate of the TLS connection is an Apache CGI (Common Gateway Interface) variable, any server-side language would work, not just PHP.

Our implementation is available online under an open-source license at <http://www.douglas.stebila.ca/research/papers/MSD14/> and <http://eprints.qut.edu.au/76270>.

6.1 Firefox Extension

The client-side Firefox extension is written in Javascript and uses an existing Firefox API to obtain the certificate of the TLS connection. The client implementation of our protocol (excluding underlying cryptographic primitives) is just 320 lines of Javascript code. Cryptographic operations can be done in either pure Javascript (relying on about 1400 lines of code from Wu’s Javascript elliptic curve cryptography and big integer arithmetic implementation² and about 6KB of minified Javascript from the Stanford Javascript Crypto Library³ for the PBKDF2 algorithm) or can make use of native C OpenSSL libraries using Firefox’s js-ctypes API⁴.

When the extension detects (using an appropriate triggering mechanism; see the full version [36]) a page that supports the protocol, it displays a notification bar that secure password authentication is supported (Figure 3(a)). The user then clicks on the “Login” button in the notification bar to bring up the password entry dialog box (Figure 3(b)). Note that the notification bar is displayed using Firefox’s API for notifications, similar to how alerts are rendered for missing plugins. By using the standard notification mechanism, we provide a trusted UI path to the notification bar, and then, through the login button on that bar, a trusted UI path to the dialog box, somewhat mitigating concerns about the difficulty of providing a trusted UI path in browser-based secure password authentication [38,12]. The status of the mutual authentication is displayed in the notification bar (Figures 3(c), (d)); if successful the browser is redirected to the URL indicated by the server.

At present, Firefox is the only web browser whose extension APIs offer partial implementation of the channel bindings for TLS from RFC5929 [17], providing access to the certificate of the page’s TLS connection. The APIs for Google Chrome and Apple Safari extensions do not seem to permit this ability so far, nor does the API for Microsoft Internet Explorer browser toolbars. However, our modular approach is still validated, in that Chrome, Safari, and IE would only need to implement the recommendations from [17] such that our extension could do the rest of the protocol, rather than requiring the full protocol be implemented within the core browser source code as many other approaches require.

² <http://www-cs-students.stanford.edu/~tjw/jsbn/>

³ <http://crypto.stanford.edu/sjcl/>

⁴ <https://developer.mozilla.org/en-US/docs/Mozilla/js-ctypes>

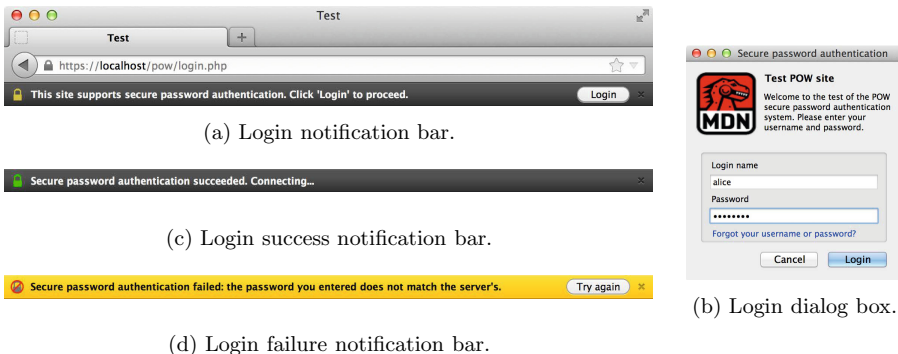


Fig. 3. User interface for Mozilla Firefox extension

Branding. Our prototype allows the server to specify some limited “branding” customizations to the login dialog box, including displaying a logo and explanatory text, as can be seen in Figure 3(b). A common objection to server-specified branding is that the protocol becomes insecure due to phishing attacks: while it is true that an attacker could put in a different logo or text in Figure 3(b), the attacker *gains nothing* in doing so: the protocol cryptographically protects the password, even when the user’s browser runs the protocol with attacker’s server. At best, the attacker can interrupt communication, but will gain no information. Our limited branding does not give the attacker enough power to completely spoof the user interface and trick the user into using the attacker’s own dialog box, due to the trusted UI path via the browser notification bar.

6.2 Performance

In Table 1, we report timings for our implementation. Timings reported are an average of 10 timings, with standard deviation. The total runtime of the protocol includes the network latency for the communication within a corporate network where the client and server machines were located. The average ping time on the network was 48.55 ms (stdev. 37.76).

We report two different sets of timings: “cross-platform” timings, using pure Javascript on the client side and PHP with built-in GMP libraries on the server side; and “native” timings, using calls to OpenSSL for cryptographic operations on both the client side and the server side.

The average total runtime from when the user clicked “Login” after entering their password until the protocol completes was 487.72 ms (stdev. 49.93) using cross-platform code, and 109.04 ms (stdev. 47.96) using native code. In comparison, the average total runtime of password authentication based on an HTML form over TLS in our setting takes 66.16 ms (stdev. 27.80)). The difference of about 40ms with our native code implementation corresponds to one additional round trip and is unlikely to be perceptible by the users. Our native cryptographic code is further comparable to Dacosta et al.’s reported performance of DVCert on laptops [23]. Our protocol implementation includes a variety of

Table 1. Average runtime in ms (\pm standard deviation) of extension using cross-platform Javascript cryptographic code, native C (OpenSSL) cryptographic code, compared with standard passwords submitted using an HTML form over SSL

Operation	Pure Javascript	Native C	HTML form + SSL
Client cryptographic computations	354.06 \pm 5.12	5.32 \pm 0.23	—
Server cryptographic computations	36.27 \pm 2.20	6.34 \pm 0.48	—
Total runtime	487.72 \pm 49.93	109.04 \pm 47.96	66.16 \pm 27.80

Software: Mozilla Firefox 21.0, Apache 2.2.22, PHP 5.4.14, GMP 5.1.1, MySQL 5.5.28, OpenSSL 1.0.1e, Mac OS X 10.8.3.

Hardware: 2.6 GHz Intel Core i7 (3720QM), 16 GB of RAM.

Network: Corporate network, ping time 48.55 \pm 37.76ms

operations beyond cryptographic computations, so the total runtime is greater than the sum of cryptographic runtime and communication time.

7 Discussion

7.1 TLS Channel Bindings

The `tls-unique` channel binding works with all TLS ciphersuites, whereas the `tls-server-end-point` binding only works with TLS ciphersuites that use certificate-based server authentication, though these are most widely used in practice.

`tls-server-end-point` may be easier to deploy on the server side since the server certificate is often fixed for long periods, and thus more suitable for multi-server architectures where for example an SSL accelerator handles the TLS connection and then passes the plaintext onto one of potentially many layers of application servers. `tls-server-end-point` is also easily deployable on the client side: for example, the Firefox extension API already makes the server certificate, but not the `Finished` message, available.

In some sense `tls-unique` is a stronger channel binding string: with it, we can achieve security of our generic construction using only CCE security of the TLS channel, whereas with `tls-server-end-point` we rely on the stronger ACCE security notion of TLS. In the end, both allow us to achieve our goal.

TLS keying material exporters [39] are another option for binding to the TLS channel, as they allow an application to obtain keying material derived from the master secret key for a given label. However, TLS channel bindings appear to be the preferred mechanism, and so we focus on them.

7.2 Challenges with PAKE

Although PAKE protocols have been known in the literature since their invention 1992, they have seen almost no deployed adoption for user authentication in real-world protocols and implementations, with the exception of the use of the socialist millionaires' protocol in the Off-the-Record Messaging (OTR) protocol for private instant messaging [26]. Engler et al. [12] recently identified several challenges—divided into two classes, user interface and deployment challenges—to adopting cryptographic protocols for password authentication in the web. It has also been noted that the myriad patents related to PAKE have had a negative impact on adoption [40].

Deployment challenges. This modular architecture may address certain deployment challenges. Engler et al. ask “What is the appropriate layer in the networking stack to integrate PAKE protocols?” They compare two proposed options: TLS-SRP [13] and an earlier draft of the HTTPS-PAKE approach of Oiwa et al. [20]. Adding SRP as a TLS ciphersuite has benefits in that, once implemented, allows multiple applications to use the same TLS implementation. But many drawbacks are identified by Engler et al., including: (i) the need to integrate the application layer with the TLS layer on both the client side (necessitating a complex API between the TLS library and the web browser, for example) and on the server side (which could negatively affect the ability of HTTPS load balancers to terminate TLS connections and then hand them off to web application servers); and (ii) the difficulty of supporting multiple authentication realms within the same domain. HTTPS-PAKE, running as an HTTP authentication mechanism at the application layer, avoids both of these problems. The version of HTTPS-PAKE reviewed by Engler et al. did not have cryptographically strong binding between the two protocols and thus could not prevent man-in-the-middle attacks, but later revisions addressed that issue. Our modular approach avoids the problems that Engler et al. identify for TLS-SRP.

Our approach also better handles the transition from unauthenticated encrypted browsing to authenticated encrypted browsing: a user may browse an HTTPS site for a while before logging in; with TLS-SRP, a new TLS connection is required (and the mechanism for triggering a new TLS connection is unclear); it is much easier to trigger the authentication at the application layer when it is required.

User interface challenges. Engler et al. [12] identify several user interface challenges. We do not aim to fully solve all these challenges in our prototype, as demonstrating a convincing solution to these challenges requires critical examination by usability experts and appropriate user studies. Nonetheless, we have endeavoured to follow some best practices that may at least partially address the identified UI challenges.

It is essential for the security of PAKE protocols that the user always enter their password into a secure dialog box. If the entry mechanism can be spoofed by an attacking website, then the user could be tricked into entering their password directly into a textfield controlled by the attacker. Thus, there must be a *trusted path* to the dialog box in the UI, usually achieved by placing the password entry visibly in the browser chrome (i.e., the parts of the window that make up the browser UI, such as the location bar, rather than the page content). In our prototype, we follow this practice by using Firefox’s notification bar. It has been suggested that permitting users to customize notification bars helps to reduce spoofing attacks [38].

The second and third of Engler et al.’s UI challenges are about how to train users to use the system in the first place, and how to communicate failures to users in a way that they do not fall back on insecure methods. Both of these remain a challenge for usability designers, though again, delivering failure notifications via Firefox’s trusted path for notifications may provide some benefit. Providing forgotten password resets securely remains an open challenge both in practice and in theory and is outside the scope of our goals.

The final challenge noted by Engler et al. is on how to allow website designers to customize and brand the login dialog without compromising security; it has been suggested that lack of customization and branding was a contributing factor to the lack of adoption of HTTP basic and digest authentication. Our prototype allows the server to provide a few customizations to the login dialog box, including a logo, some explanatory text, as shown in Figure 3(b). While an attacker could use stolen images, the benefit to the attacker is minimal since the password entry will be cryptographically protected.

Adoption challenges. A final challenge for any new security technology is facilitating widespread adoption. Such protocols see a “network effect”: it is only useful for a client Alice to use the technology if there are many Bobs who support it, and vice versa. In the end, any secure password authentication technology will be most successful once built in to all major web browsers and web application frameworks. In the meantime, the modular approach in this paper is suitable for gradual deployment. For example, an organization can internally standardize on the use of this approach by deploying an extension to all of its users’ browsers without needing to wait for the browser vendor to support the protocol. The more adoption via extension, the more evidence for interest in the technology, and the greater incentive for vendors to provide a native implementation.

Acknowledgements. Mark Manulis was supported by the German Research Foundation (DFG), project PRIMAKE (MA 4957). Douglas Stebila acknowledges funding from Australian Research Council (ARC) Discovery Project DP130104304.

References

1. Schechter, S.E., Dhamija, R., Ozment, A., Fischer, I.: The emperor's new security indicators. In: 2007 IEEE Symposium on Security and Privacy, pp. 51–65. IEEE Computer Society Press (2007)
2. Sunshine, J., Egelman, S., Almuhimedi, H., Atri, N., Cranor, L.F.: Crying wolf: An empirical study of SSL warning effectiveness. In: USENIX Security 2009 (2009)
3. Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., Stewart, L.: HTTP Authentication: Basic and Digest Access Authentication. RFC 2617 (Draft Standard), Updated by RFC 7235 (1999)
4. Bellare, S.M., Merritt, M.: Encrypted key exchange: Password-based protocols secure against dictionary attacks. In: 1992 IEEE Symposium on Security and Privacy, pp. 72–84. IEEE Computer Society Press (1992)
5. Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated key exchange secure against dictionary attacks. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 139–155. Springer, Heidelberg (2000)
6. Canetti, R., Halevi, S., Katz, J., Lindell, Y., MacKenzie, P.: Universally composable password-based key exchange. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 404–421. Springer, Heidelberg (2005)
7. Abdalla, M., Catalano, D., Chevalier, C., Pointcheval, D.: Efficient two-party password-based key exchange protocols in the UC framework. In: Malkin, T. (ed.) CT-RSA 2008. LNCS, vol. 4964, pp. 335–351. Springer, Heidelberg (2008)
8. International Organization for Standardization (ISO): ISO/IEC 11770-4: Information technology — security techniques — key management — part 4: Mechanisms based on weak secrets (2006)
9. ITU-T X.1035: Password-authenticated key exchange (PAK) protocol (2007)
10. IEEE P1363.2: Standard specifications for password-based public-key cryptographic techniques (2008)
11. Wu, T.D.: The secure remote password protocol. In: NDSS 1998. The Internet Society (1998)
12. Engler, J., Karlof, C., Shi, E., Song, D.: Is it too late for PAKE? In: Web 2.0 Security and Privacy (W2SP) 2009 (2009)
13. Taylor, D., Wu, T., Mavrogiannopoulos, N., Perrin, T.: Using the Secure Remote Password (SRP) Protocol for TLS Authentication. RFC 5054, Informational (2007)
14. Abdalla, M., Bresson, E., Chevassut, O., Möller, B., Pointcheval, D.: Provably secure password-based authentication in TLS. In: Lin, F.C., Lee, D.T., Lin, B.S., Shieh, S., Jajodia, S. (eds.) ASIACCS 2006, pp. 35–45. ACM Press (2006)
15. Abdalla, M., Pointcheval, D.: Simple password-based encrypted key exchange protocols. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 191–208. Springer, Heidelberg (2005)
16. Hao, F., Ryan, P.Y.A.: Password authenticated key exchange by juggling. In: Christianson, B., Malcolm, J.A., Matyas, V., Roe, M. (eds.) Security Protocols 2008. LNCS, vol. 6615, pp. 159–171. Springer, Heidelberg (2011)
17. Altman, J., Williams, N., Zhu, L.: Channel Bindings for TLS. RFC 5929 (Proposed Standard) (2010)
18. Oiwa, Y., Takagi, H., Watanabe, H., Suzuki, H.: PAKE-based mutual HTTP authentication for preventing phishing attacks. In: Maarek, Y., Nejdl, W. (eds.) Proc. 18th International World Wide Web Conference (WWW 2009), pp. 1143–1144. ACM (2009)

19. Oiwa, Y., Watanabe, H., Takagi, H.: PAKE-based mutual HTTP authentication for preventing phishing attacks (2009), <http://arxiv.org/abs/0911.5230>
20. Oiwa, Y., Watanabe, H., Takagi, H., Ioku, Y., Hayashi, T.: Mutual authentication protocol for HTTP (2012), Internet-Draft, <http://tools.ietf.org/html/draft-oiwa-http-mutualauth-12>
21. AIST Research Center for Information Security: (Mutual authentication protocol for HTTP), <https://www.rcis.aist.go.jp/special/MutualAuth>
22. Kwon, T.: Authentication and key agreement via memorable passwords. In: NDSS 2001. The Internet Society (2001)
23. 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)
24. Boyko, V., MacKenzie, P.D., Patel, S.: Provably secure password-authenticated key exchange using diffie-hellman. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 156–171. Springer, Heidelberg (2000)
25. Borisov, N., Goldberg, I., Brewer, E.A.: Off-the-record communication, or, why not to use PGP. In: ACM Workshop on Privacy in Electronic Society (WPES 2004), pp. 77–84. ACM Press (2004)
26. Alexander, C., Goldberg, I.: Improved user authentication in Off-The-Record messaging. In: Yu, T. (ed.) ACM Workshop on Privacy in Electronic Society (WPES 2007), pp. 41–47. ACM Press (2007)
27. 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)
28. Bergsma, F., Dowling, B., Kohlar, F., Schwenk, J., Stebila, D.: Multi-ciphersuite security of the Secure Shell (SSH) protocol. In: Yung, M., Li, N. (eds.) ACM CCS 2014. ACM Press (2014)
29. Brzuska, C., Smart, N.P., Warinschi, B., Watson, G.J.: An analysis of the EMV channel establishment protocol. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013, pp. 373–386. ACM Press (2013)
30. 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)
31. Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DH and TLS-RSA in the standard model. Cryptology ePrint Archive, Report 2013/367 (2013), <http://eprint.iacr.org/2013/367>
32. Giesen, F., Kohlar, F., Stebila, D.: On the security of TLS renegotiation. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013, pp. 387–398. ACM Press (2013)
33. 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)
34. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: Generic compilers for authenticated key exchange. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 232–249. Springer, Heidelberg (2010)
35. Fleischhacker, N., Manulis, M., Azodi, A.: A Modular Framework for Multi-Factor Authentication and Key Exchange. Cryptology ePrint Archive, Report 2012/181 (2012), <http://eprint.iacr.org/2012/181>

36. Manulis, M., Stebila, D., Denham, N.: Secure modular password authentication for the web using channel bindings (full version). Cryptology ePrint Archive, Report 2014/731 (2014), <http://eprint.iacr.org/2014/731>
37. National Institute of Standards and Technology: Recommended elliptic curves for federal government use (1999), <http://csrc.nist.gov/groups/ST/toolkit/documents/dss/NISTReCur.pdf>
38. Dhamija, R., Tygar, J.D.: The battle against phishing: Dynamic security skins. In: Cranor, L.F., Zurko, M.E. (eds.) Symposium on Usable Privacy and Security (SOUPS 2005), pp. 77–88. ACM Press (2005)
39. Rescorla, E.: Keying Material Exporters for Transport Layer Security (TLS). RFC 5705 (Proposed Standard) (2010)
40. Abdalla, M., Bresson, E., Chevassut, O., Möller, B., Pointcheval, D.: Strong password-based authentication in TLS using the three-party group Diffie–Hellman protocol. International Journal of Security and Networks 2, 284–296 (2007)
41. Certicom Research: SEC 1: Elliptic curve cryptography, Version 2.0 (2009)

A Confidential Channel Establishment (CCE)

The model of CCE protocols can be described using the setting of PACCE protocols from Section 2, with a few differences. The first difference is that there are no passwords or identities involved; hence no **Corrupt** oracle is needed, nor is the U' parameter required in the initialization in the **Send^{pre}** query. Further, the security condition is adjusted so that only sessions where the adversary was passive in the handshake stage are considered. The oracles **RevealSK**, **Encrypt** and **Decrypt** remain unchanged. The following definition of CCE security is obtained from Definition 1 by considering the above mentioned modifications.

Definition 2 (CCE security). *An adversary \mathcal{A} is said to (t, ϵ) -break a CCE protocol if \mathcal{A} runs in time t and, when \mathcal{A} terminates and outputs a triple (U, s, b') such that*

(a) $\Pi_U^s.\alpha = \text{accept}$, and

(b) *there exists an instance $\Pi_{U'}^{s'}$ that is partnered to Π_U^s , and*

(c) *\mathcal{A} did not issue **RevealSK**(Π_U^s) or **RevealSK**($\Pi_{U'}^{s'}$) for any $\Pi_{U'}^{s'}$ that is partnered to Π_U^s ,*

then $|\Pr[b' = \Pi_U^s.b] - \frac{1}{2}| \geq \epsilon$.

B Tag-Based Password Authentication (tPAAuth)

The model of tPAAuth can be described using the setting of PACCE protocols from Section 2. A tPAAuth session is executed between a client instance Π_C^s and a server instance $\Pi_S^{s'}$ on input the corresponding password $\text{pw}_{C,S}$ from the dictionary \mathcal{D} and some tag $\tau \in \{0, 1\}^*$. A tPAAuth session is successful if both instances use the same password $\text{pw}_{C,S}$ and tag τ as their input. The requirement

on tag equality leads to the extended definition of partnering: two instances Π_C^s and $\Pi_S^{s'}$ are *partnered* if $\Pi_C^s.\text{pid} = S$, $\Pi_S^{s'}.\text{pid} = C$, $\Pi_C^s.T \approx \Pi_S^{s'}.T$ (matching transcripts), and $\Pi_C^s.\tau = \Pi_S^{s'}.\tau$ (equal tags).

A tPAAuth adversary \mathcal{A} is active and interacts with instances of $U \in \{C, S\}$ using the following oracles:

- $\text{Send}(\Pi_U^s, m)$: This query is identical to Send^{PTE} from the PACCE model except for one important difference — when \mathcal{A} initializes some instance Π_U^s using the special messages $m = (\text{init}, U', \tau)$ or $m = (\text{resp}, U', \tau)$ then it additionally provides as input a tag τ which will be used by the instance in the tPAAuth session. This essentially gives \mathcal{A} full control over the tags that are used in the protocol.
- $\text{Corrupt}(C, S)$: This query reveals the corresponding password $\text{pw}_{C,S}$.

The security of tPAAuth protocols, defined in the following, extends the traditional password authentication requirement that accounts for online dictionary attacks with the requirement of tag equality, which is implied by condition 3 due to the extended definition of partnering.

Definition 3 (tPAAuth security). *An adversary \mathcal{A} is said to (t, ϵ) -break a tPAAuth protocol if after the termination of \mathcal{A} that runs in time t with probability at least $\epsilon + O(n/|\mathcal{D}|)$ where n is the number of initialized tPAAuth instances there exists an instance Π_U^s such that*

1. $\Pi_U^s.\alpha = \text{accept}$, and
2. \mathcal{A} did not issue $\text{Corrupt}(\text{init}(\Pi_U^s), \text{resp}(\Pi_U^s))$ before Π_U^s accepted, and
3. there is no unique instance $\Pi_{U'}^{s'}$ that is partnered to Π_U^s .

A tPAAuth protocol is (t, ϵ) -secure if there is no \mathcal{A} that (t, ϵ) -breaks it; it is secure if it is (t, ϵ) -secure for all polynomial t and negligible ϵ in security parameter κ .

Our tagged variant tSOKE of the Simple Open Key Exchange (SOKE) protocol [14] is shown in Figure 4, and is a secure tPAAuth protocol (see full version [36]).

System parameters

Elliptic curve group nistp192 with generator G of order n

Second generator G' constructed verifiably at random with $\langle G \rangle = \langle G' \rangle$

$G' = (0x8da36f68628a18107650b306f22b41448cb60fe5712dd57a,$

$0x1f64a649852124528a09455de6aad151b4c0a9a8c2e8269c)$

(constructed verifiably at random [41, §3.1.3.2] with seed string “This is the seed string for the web passwords protocol.”)

Registration stage (takes place over a secure channel)

Client A	Server B
1. Enter username id_A .	
2. Enter password pw_{AB} .	
3. $\text{salt} \leftarrow \{0, 1\}^{128}$.	
4. Choose iteration counter $c \in \mathbb{N}^+$.	
5. $h \leftarrow \text{PBKDF2}(\text{SHA-256}, \text{pw}_{AB}, \text{salt}, c, 256)$	
6.	$\xrightarrow{\text{id}_A, \text{salt}, c, h}$
7.	Store salt, c, h for id_A .

Login stage

Client A	Server B
Input: tag τ	Input: tag τ
1. Enter username id_A .	
2. Enter password pw_{AB} .	
3. $x \leftarrow_R \{2, \dots, n-1\}$	
4. $X \leftarrow xG$	
5.	$\xrightarrow{\text{id}_A, X}$
6.	Look up salt, c, h for id_A .
7.	$y \leftarrow_R \{2, \dots, n-1\}$
8.	$Y \leftarrow yG$
9.	$Y^* \leftarrow Y + (h \bmod n)G'$
10.	$\xleftarrow{\text{salt}, c, Y^*}$
11. $h \leftarrow \text{PBKDF2}(\text{SHA-256}, \text{pw}_{AB}, \text{salt}, c, 256)$	
12. $Y \leftarrow Y^* - (h \bmod n)G'$	
13. $Z \leftarrow xY$	$Z' \leftarrow yX$
14. $\text{pms} \leftarrow \text{SHA-256}(\text{id}_A, h, \tau, X, Y^*, Z)$	$K \leftarrow \text{SHA-256}(\text{id}_A, h, \tau, X, Y^*, Z)$
15. $A_1 \leftarrow \text{SHA-256}(\text{pms}, \text{“auth1”})$	
16.	$\xrightarrow{A_1}$
17.	Abort if $A_1 \neq \text{SHA-256}(\text{pms}, \text{“auth1”})$
18.	$A_2 \leftarrow \text{SHA-256}(\text{pms}, \text{“auth2”})$
19.	$\xleftarrow{A_2}$
20. Abort if $A_2 \neq \text{SHA-256}(\text{pms}, \text{“auth2”})$	
21. Accept	Accept

Fig. 4. tSOKE protocol registration and login stages