

# A Simple Threshold Authenticated Key Exchange from Short Secrets

Michel Abdalla<sup>1</sup>, Olivier Chevassut<sup>2</sup>, Pierre-Alain Fouque<sup>1</sup>, and David Pointcheval<sup>1</sup>

<sup>1</sup> Département d’Informatique  
École normale supérieure  
45 Rue d’Ulm, 75230 Paris Cedex 05, France  
{Michel.Abdalla,Pierre-Alain.Fouque,David.Pointcheval}@ens.fr  
<http://www.di.ens.fr/~{mabdalla,fouque,pointche}>

<sup>2</sup> Lawrence Berkeley National Laboratory  
Berkeley, CA 94720, USA  
OChevassut@lbl.gov  
<http://www.dsd.lbl.gov/~chevassu>

**Abstract.** This paper brings the password-based authenticated key exchange (PAKE) problem closer to practice. It takes into account the presence of firewalls when clients communicate with authentication servers. An authentication server can indeed be seen as two distinct entities, namely a gateway (which is the direct interlocutor of the client) and a back-end server (which is the only one able to check the identity of the client). The goal in this setting is to achieve both transparency and security for the client. And to achieve these goals, the most appropriate choices seem to be to keep the client’s password private—even from the back-end server—and to use threshold-based cryptography. In this paper, we present the Threshold Password-based Authenticated Key Exchange (GTPAKE) system: GTPAKE uses a pair of public/private keys and, unlike traditional threshold-based constructions, shares only the private key among the servers. The system does not require any certification—except during the registration and update of clients’ passwords—since clients do not use the public-key to authenticate to the gateway. Clients only need to have their password in hand. In addition to client security, this paper also presents highly-desirable security properties such as *server password protection* against dishonest gateways and *key privacy* against curious authentication servers.

**Keywords.** Threshold Protocols, Password-based Authentication.

## 1 Introduction

**Problem Description.** Consider the familiar scenario where you are at the airport waiting for your flight. You have checked-in and have now half an hour to kill. What do you do? Turn on your laptop, switch on your wireless card, and pick up the airport wireless LAN. You are prompted for a password to authenticate yourself and upon successful authentication a port is opened for you to browse

the Internet and/or read your e-mails. Now you may wonder what happens under the hood. We have indeed talked to an airport gateway, often termed hotspot, that has in turn talked to your mobile-Internet provider. T-Mobile is an example of such a provider in the United States. The gateway has passed—in an encrypted form—your password to the provider for authentication and gets back a yes/no depending whether or not the authentication was successful.

Although this model is very attractive in practice, existing security solutions implementing it have major drawbacks since the gateway gains some amount of information about your password. The ideal solution is a cryptographic algorithm allowing the client to securely exchange a session key with the gateway, but the gateway does not gain any information about the password and the authentication server does not gain any information about the session key. Additional problems also occur if too many people, from the same provider, try to connect to various gateways at the same time. The authentication check from the provider would become a bottleneck. Various authentication points are very desirable. Nevertheless, the password of the client cannot be stored at several places, otherwise the job of hackers would be made much easier.

**Scenario.** To formally define a model for the above scenario, we propose a model in which one can design a protocol among three parties (the client, the gateway and the authentication server) which protects both the session keys and the passwords. We indeed require the three following security notions which capture dishonest behaviors of the client, the authentication server and the gateway respectively: the *semantic security* of the sessions keys, which we model by a Real-Or-Random (ROR) game [2] (it has been proved strictly stronger than the more classical Find-Then-Guess (FTG) model [4]); the *key-privacy* notion [2] which entails that the session key exchanged by two parties with the help of an authentication server is unknown to the authentication server (and also to any other party, granted the semantic security); the *server password protection* which basically means that the gateway cannot learn any information about clients' passwords from the authentication server.

The ultimate goal of this paper is to achieve the highest level of security in the PAKE setting. With the above security notions, breaking into the gateway would not help to gain any information about the passwords, however the authentication server is a security hole. Breaking into the latter would leak the authentication information. Furthermore, according to the above motivation example, a unique authentication point may be a bottleneck. When data information is crucial, a usual solution to protect it is to distribute it among several servers such that a majority of them is needed to recover the initial data. Moreover, when we want to protect a cryptographic service we can split the private information into several parts, each known by one server, so that a majority of them is required to maintain the service without reconstructing the secret key in a single place. Threshold cryptography is the field that provides such solutions and allows to take into account adversaries that can break into any minority of parties. It furthermore solves the bottleneck problem.

**Contribution.** Our contribution in this paper is a provably-secure protocol satisfying the previously mentioned requirements. We have constructed it by defining a simple protocol, called the gateway PAKE (GPAKE) protocol, among three parties (the client, the gateway and the authentication server). GPAKE protects the session keys and the passwords according to a formal security model which we specify in this paper. It provides the additional property to be a variant of AuthA [5] perfectly *transparent* to the client. Transparency means that a client does not (need to) know whether he is talking to a server directly or whether the server is implemented as a gateway, an authentication server, or even an application server. The gateway does not also (need to) know whether the authentication server is distributed. A non-transparent protocol indeed raises real concerns on the utilization of the protocol in practice since clients need to first update their cryptographic stack in order to take advantage of the threshold PAKE feature. A transparent protocol on the other hand lets only domain administrators worry about deploying the threshold PAKE feature to their users. We have developed a threshold version (called gateway threshold PAKE (GTPAKE)) that does not break the transparency property of GPAKE, and defined its execution in our security model. Clients already running the two-party AuthA protocol (e.g. OMDHKE [8]) will not have to upgrade their stack when administrators add an extra layer of protection with GTPAKE!

**Related Work.** Several password-based key exchange protocols in the threshold setting have been proposed in the past by MacKenzie et al. [14] and by Di Raimondo and Gennaro [10], to name a few. In particular, the protocol by Di Raimondo and Gennaro in Eurocrypt 2003 [10], which is a threshold version of the 2-party KOY protocol [12], is proven secure in the standard model. In that paper, they have introduced the notion of *transparent* protocols, where the initial protocol and its threshold version are the same from the point of view of the client. Unfortunately, their solution is not very efficient from a practical point of view since it requires several rounds of communication between the client and the server and among the servers themselves. Moreover, like previous proposals of threshold password-based key exchange protocols, their protocol requires the password to be shared among all the servers.

The solution we propose in this paper does not require passwords to be shared across different servers. Instead, we only share the secret decryption key for a public-key encryption scheme under which all passwords are encrypted. This provides an additional feature: it is quite easy for a client to modify his password. He just needs to send the new one encrypted under the authentication servers' public key.

Moreover, contrary to the hybrid model of Halevi and Krawczyk [11], where the server has a public/secret key pair and the client only knows a password, the client is not required to check the authenticity of the public key during the execution of the protocol. Integrity is required only during the registration or when the user wants to update his password.

**Organization of the Paper.** In Section 2, we present the formalization used to define the execution of the GPAKE and GTPAKE protocols. Our formalization extends that of Abdalla et al. to the threshold setting [2]. In Section 3, we present the intractability assumptions used throughout the paper. In Section 4, we describe the GPAKE system and show that it achieves *semantic security* and *key privacy* in a provable secure way. In Section 5, we describe the GPAKE system’s threshold version and show that it is secure —via a reduction from the security of GTPAKE to the intractability assumptions— against dictionary attacks.

## 2 Security Model

In this section, we present the security model we will use in the rest of the paper to define the execution of our protocol for threshold password-authenticated key exchange.

### 2.1 Overview

**GATEWAY-ORIENTED PASSWORD-BASED KEY EXCHANGE.** A gateway-oriented password-based key exchange is a three-party protocol among a client, a gateway, and an authentication server. The goal of protocol is to establish an implicitly authenticated session key between the client and the gateway with the help of the authentication server, where the authentication is done by means of a short password. In our model, the password is known to both the client and the authentication server, but not to the gateway. In fact, no long-term secrets are stored in the gateway. The authentication server, on the other hand, is assumed to know the password. While the communication channel between the gateway and the authentication server is assumed to be authenticated and private, the channel connecting the client to the gateway may be insecure and perhaps under the control of an adversary.

The security goals of our gateway-oriented password-based key exchange model are also somewhat different from those of previous models for password-based schemes. In particular, we ask that the session key shared between the gateway and the client should remain private to the authentication server (see Section 2.2 for more details). Moreover, we also ask that the chances of the gateway learning some information on the password after multiple interactions with the server, perhaps concurrently, should be negligible.

**PROTOCOL PARTICIPANTS.** The participants in a gateway-oriented password-based key exchange are the client  $C \in \mathcal{C}$ , the gateway  $G \in \mathcal{G}$ , and the authentication server  $S \in \mathcal{S}$ . We denote by  $\mathcal{U}$  the set of all participants (i.e.,  $\mathcal{U} = \mathcal{C} \cup \mathcal{G} \cup \mathcal{S}$ ) and by  $U$  a non-specific participant in  $\mathcal{U}$ . Each client  $C \in \mathcal{C}$  holds a password  $\text{pw}_C$ . Each server  $S \in \mathcal{S}$  holds a vector of passwords  $\text{PW}_S = \langle \text{pw}_C \rangle_{C \in \mathcal{C}}$  with an entry for each client.

## 2.2 Security Model

Since we assume an authenticated and private channel between the gateway and the server, the communication model is similar to previous one for 2-party authenticated key exchange. In particular, we adopt the Real-Or-Random (ROR) security model of Abdalla *et al.* [2] for password-based authenticated key exchange protocol, which in turn implies that of Bellare *et al.* [4]. As in the standard model, all the interactions between an adversary  $\mathcal{A}$  and the protocol participants in the ROR model are done via oracle queries. Let  $U^i$  denote the instance  $i$  of a participant  $U$ . The list of oracles available to the adversary are as follows:

- $\text{Execute}(C^i, G^j)$ : This query models passive attacks in which the attacker eavesdrops on honest executions among a client instance  $C^i$  and a gateway instance  $G^j$ . The output of this query consists of the messages that were exchanged during the honest execution of the protocol.
- $\text{Send}(U^i, m)$ : This query models an active attack against the client or gateway instance  $U^i$ , in which the adversary may intercept a message and then modify it, create a new one, or simply forward it to the intended recipient. The output of this query is the message that the participant instance  $U^i$  would generate upon receipt of message  $m$ .

**The Real-Or-Random Model [2].** In the ROR model, in addition to the above-mentioned oracles, an attacker is also given access to a less restrictive **Test** oracle. Let  $b$  be a bit chosen uniformly at random at the beginning of the experiment defining the semantic security of session keys. The **Test** oracle in the ROR model is defined as follows:

- $\text{Test}(U^i)$ : If no session key for instance  $U^i$  is defined, then return the undefined symbol  $\perp$ . Otherwise, return the session key for instance  $U^i$  if  $b = 1$  or a random key of the same size if  $b = 0$ .

As in standard models, the **Test** oracle in the ROR model also tries to capture the adversary’s ability (or inability) to tell apart a real session key from a random one. The main difference is that it does so not only for a single session but for all sessions. More precisely, the adversary in the ROR model is not restricted to ask a single **Test** query, but it can in fact ask multiple ones. All **Test** queries in this case will be answered using the same value for the hidden bit  $b$  that was chosen at the beginning of the experiment defining the semantic security of the session keys. That is, the keys returned by the **Test** oracle are either all real or all random. However, in the random case, the same random key value is returned for **Test** queries that are asked to two instances that belong to the same session (see notion of partnering below). The goal of the adversary in the ROR model is still the same: to guess the value of the hidden bit  $b$  used to answer **Test** queries. The adversary is considered successful if it guesses  $b$  correctly.

**PARTNERING.** As in [2], we use the notion of partnering based on session identifications (*sid*), which says that two instances are partnered if they hold the

same non-null *sid*. More specifically, a client instance  $C^i$  and a gateway instance  $G^j$  are said to be partners if the following conditions are met: (1) Both  $C^i$  and  $G^j$  accept; (2) Both  $C^i$  and  $G^j$  share the same session identifications; (3) The partner identification for  $C^i$  is  $G^j$  and vice-versa; and (4) No instance other than  $C^i$  and  $G^j$  accepts with a partner identification equal to  $C^i$  or  $G^j$ . In practice, the *sid* is taken to be the partial transcript of the conversation among the client and the gateway instances before the acceptance.

**FRESHNESS.** Differently from [2], we opt not to embed the notion of freshness inside the definition of the oracles. Instead, we take the more standard approach of explicitly defining the notion of freshness and mandating the adversary to only perform tests on *fresh* instances. The two approaches are however equivalent. In particular, we say that a *instance* of a client or gateway is *fresh* if it has accepted.

**FORMAL DEFINITION.** Let **SUCC** denote the event in which the adversary is successful. The **ake – ror-advantage** of an adversary  $\mathcal{A}$  in violating the semantic security of the protocol  $P$  in the ROR sense and the **advantage function** of the protocol  $P$ , when passwords are drawn from a dictionary **Dict**, are respectively

$$\begin{aligned} \mathbf{Adv}_{P,\mathbf{Dict}}^{\text{ake-ror}}(\mathcal{A}) &= 2 \cdot \Pr[\mathbf{SUCC}] - 1 \text{ and} \\ \mathbf{Adv}_{P,\mathbf{Dict}}^{\text{ake-ror}}(t, R) &= \max_{\mathcal{A}} \{ \mathbf{Adv}_{P,\mathbf{Dict}}^{\text{ake-ror}}(\mathcal{A}) \}, \end{aligned}$$

where the maximum is over all  $\mathcal{A}$  with time-complexity at most  $t$  and using resources at most  $R$  (such as the number of queries to its oracles). The definition of time-complexity that we use henceforth is the usual one, which includes the maximum of all execution times in the experiments defining the security plus the code size.

Please note that, as proven in [2], any scheme proven secure in the ROR model is also secure in the model of Bellare *et al.* [4]. The converse, however, is not necessarily true (see [2] for more details).

**AUTHENTICATION.** The notion of semantic security does not guarantee the existence of a partner, but only the secrecy of the session key (implicit authentication). In order to address this problem, one usually adds mechanisms for explicit authentication of client and gateway instances. In this paper, we only consider *unilateral authentication* of the gateway, by which a client instance can be ensured that it has in fact established a key with the gateway instance it intended to. As in [7], we denote by  $\mathbf{Succ}_{\mathcal{A}}^{\text{G-auth}}$  the probability that adversary  $\mathcal{A}$  successfully impersonates the gateway in an execution of the protocol. This is the probability with which a client instance accepts without having a gateway partner. The advantage function of the protocol can be defined as in previous cases.

**Key Privacy.** The notion of key privacy was introduced in [2] to capture the idea that the session key computed by two parties with the aid of an authentication server should only be known to those two parties and not to the server. The goal in this case is to reduce the amount of trust one puts into the server.

In order to meet this goal, one has to consider an adversary with access to all the secret information stored in the server and then show that such adversary cannot distinguish actual session keys from random ones if we restrict this adversary to test sessions in which the keys are shared between two oracles. The latter restriction is important since an adversary with access to the secrets of the authentication server could always establish a key with a client by playing the roles of the gateway and authentication server. Since one of our main goals is to show that the key shared between the client and the gateway is not known to the authentication server, we also use the notion of key privacy.

To capture the above intuition more formally, [2] considers an adversary which has access to all the secrets held by the authentication server and to the oracles used in the experiment defining semantic security. They then introduce a new type of oracle, called `TestPair`, whose goal is to capture the adversary’s ability to distinguish the real session key shared between any two oracle instances from a random one. The inputs to the `TestPair` oracle are the specific oracle instances whose shared session key the adversary thinks it can tell apart from a random key.

**FORMAL DEFINITION.** Consider an execution of the key exchange protocol  $P$  by an adversary  $\mathcal{A}$  with access to all the secret held by the authentication server as well as to the `Execute`, `Send`, and `TestPair` oracles. Let `SUCC` denote the event in which the adversary is successful in guessing the hidden bit used by `TestPair` oracle when only asking `TestPair` queries to instances pairs that have accepted. The **kp-advantage** of an adversary  $\mathcal{A}$  in violating the key privacy of the protocol  $P$  in the ROR sense ( $\text{Adv}_{P, \text{Dict}}^{\text{ake-kp}}(\mathcal{A})$ ) and the **advantage function** of  $P$  ( $\text{Adv}_{P, \text{Dict}}^{\text{ake-kp}}(t, R)$ ), when passwords are drawn from a dictionary `Dict`, are then defined as in previous definitions.

**Server Password Protection.** As we mentioned earlier, one of the goals of our protocol is to guarantee that the gateway is not capable of learning the user’s password that is stored in the server. Clearly, as in the case of semantic security, one cannot hope for much since, in each interaction, the adversary may be able to eliminate one candidate password from the list of possible passwords. However, we ask that the adversary should not be able to do much better than. That is, if the adversary interacts  $q$  times with the server, then the probability that it can distinguish the true password from a random one in the dictionary should be only negligibly larger than  $O(q/N)$ , where  $N$  is the size of the dictionary. The hidden constant in this case should be as small as possible (preferably 1). Note that, in this definition, the dictionary is assumed to be uniformly distributed.

### 2.3 Threshold Security Model

**THRESHOLD GATEWAY-ORIENTED PASSWORD-BASED KEY EXCHANGE.** A  $(t, k, n)$ -threshold gateway-oriented password-based key exchange is an extension of the

basic gateway-oriented password-based key exchange in which the authentication server is a distributed entity. More specifically, the clients' passwords are no longer known to any single server. Instead, each server in the set of  $n$  authentication servers is assumed to hold a share of the secret key of a public-key encryption scheme, under which clients' passwords are encrypted. The authentication of any client will require the cooperation of some size- $k$  subset of honest servers. In addition, any adversary who learns  $t$  or fewer shares of the secret key should not learn any information about the clients' passwords.

**PARTICIPANTS.** The participants in a threshold protocol are the client  $C \in \mathcal{C}$ , the gateway  $G \in \mathcal{G}$ , the set of authentication servers  $\{S_1, \dots, S_n\}$  with  $S_i \in \mathcal{S}$ , and a trusted dealer.

**Semantic Security.** The definition of semantic security of threshold protocols follows the one given above for gateway-oriented protocols. At the beginning of the semantic security experiment, the adversary selects a subset of at most  $t = k - 1$  servers to corrupt. We say that the adversary is *static* when it chooses the set of servers to corrupt in advance, before seeing anything. A special server, called the combiner, will be used to perform some tasks that do not require any secret. The combiner is also in charge of all communications between the gateway and the other servers.

The dealer generates a public key  $\text{pk}$  and a secret key  $\text{sk}$  for an encryption scheme. Then, he performs the secret sharing of  $\text{sk}$  and sends the part  $\text{sk}_i$  to  $P_i$  along with a verification key  $\text{vk}_i$ . The adversary obtains the secret key shares of the corrupted servers, along with the public key and the verification keys.

After this phase, the adversary is given access to the same set of oracles used in the standard security model for gateway-oriented password-based authenticated key exchange protocols.

**Robustness.** We say that a threshold scheme is *robust* when it takes into account *malicious* adversaries whose behavior can be different from the protocol. To force the servers to correctly perform their job, we use proofs of validity in our protocol. This also enables the combiner to correctly decrypt.

### 3 Diffie-Hellman Assumptions

In this section, we recall the definitions of standard Diffie-Hellman assumptions and introduce some new variants, which we use in the security proof of our protocol. We also present some relations between these assumptions.

#### 3.1 Classical Assumptions

Henceforth, we assume a finite cyclic group  $G$  of prime order  $p$  generated by an element  $g$ . We also call the tuple  $\mathbb{G} = (G, g, p)$  a represented group.



**Computational Diffie-Hellman Assumption: CDH.** The CDH assumption, in a represented group  $\mathbb{G}$ , with respect to the basis  $X$ , states that given two elements  $X' = X^u$  and  $Y = X^v$ , where  $u$  and  $v$  were drawn at random from  $\mathbb{Z}_p$ , it is hard to compute  $Y' = Y^u = X^{uv}$ . This can be defined more precisely by considering an experiment  $\mathbf{Exp}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X)$ , in which we select an exponent  $u$  in  $\mathbb{Z}_p$ , an element  $Y$  in  $G$ , compute  $X' = X^u$ , and then give both  $X'$  and  $Y$  to  $\mathcal{A}$ . Let  $Y'$  be the output of  $\mathcal{A}$ . Then, the experiment  $\mathbf{Exp}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X)$  outputs 1 if  $Y' = Y^u$  and 0 otherwise. We define the *advantage* of  $\mathcal{A}$  in violating the CDH assumption with respect to  $X$  as  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X) = \Pr[\mathbf{Exp}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X) = 1]$ , the *advantage* of  $\mathcal{A}$  in violating the CDH assumption (with a random basis) as  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}) = \mathbf{E}_X [\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X)]$ , and the *advantage functions*,  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(t, X)$  and  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(t)$ , as the maximum values of  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A}, X)$  and  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(\mathcal{A})$  over all  $\mathcal{A}$  with time-complexity at most  $t$ .

It is also often assumed that, *independently of what the value of  $X$  is (as long as it is a generator, of order  $p$ )*, the CDH problem with respect to the basis  $X$  is hard: the maximal value of  $\mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(t, X)$  over all generators  $X$  is small for any reasonable  $t$ .

**Decisional Diffie-Hellman Assumption: DDH.** Roughly, the DDH assumption, with respect to the basis  $X$ , states that the distributions  $(X, X' = X^u, Y, Y' = Y^u)$  and  $(X, X' = X^u, Y, Z = Y^v)$  are computationally indistinguishable when  $Y$  is drawn at random from  $G$ , and  $u$  and  $v$  are drawn at random from  $\mathbb{Z}_p$ . As before, we can define the DDH assumption more formally by defining two experiments,  $\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-real}}(\mathcal{A}, X)$  and  $\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-rand}}(\mathcal{A}, X)$ . In both experiments, we compute two random values  $X' = X^u$  and  $Y$  as before. But in addition to that, we also provide a third input, which is  $Y^u$  in  $\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-real}}(\mathcal{A}, X)$  and  $Y^v$ , for a random  $v$ , in  $\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-rand}}(\mathcal{A}, X)$ . The goal of the adversary is to guess a bit indicating the experiment he thinks he is in. We define the *advantage* of  $\mathcal{A}$  in violating the DDH assumption, with respect to the basis  $X$ ,  $\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(\mathcal{A}, X)$ , as  $\Pr[\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-real}}(\mathcal{A}, X) = 1] - \Pr[\mathbf{Exp}_{\mathbb{G}}^{\text{ddh-rand}}(\mathcal{A}, X) = 1]$ , and *advantage* of  $\mathcal{A}$  in violating the DDH assumption (with random basis) as  $\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(\mathcal{A}) = \mathbf{E}_X [\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(\mathcal{A}, X)]$ . The *advantage functions*  $\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(t, X)$  and  $\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(t)$  are then defined in a similar manner as above.

Again, it is also often assumed that, *independently of what  $X$  is (as long as it is a generator, of order  $p$ )*, the DDH problem with respect to the basis  $X$  is hard: the maximal value of  $\mathbf{Adv}_{\mathbb{G}}^{\text{ddh}}(t, X)$  over all generators  $X$  is small for any reasonable  $t$ .

### 3.2 Password-Based Chosen-Basis Diffie-Hellman Assumptions

The actual proofs of security of our protocol use password-related versions of the above Diffie-Hellman assumptions, in which the adversary has some control over the basis, hence the name *password-based chosen-basis decisional/computational Diffie-Hellman* assumptions. They make use of a dictionary  $\mathcal{D} = \{U_1, \dots, U_N\}$

of size  $N$ . Then, we assume that when the adversary has not correctly predicted the password (1 chance over  $N$ ), he has no significant advantage. Hence, its overall advantage cannot be significantly larger than  $1/N$ .

We start by presenting the password-based chosen-basis computational Diffie-Hellman assumption.

**Definition 1 (PCCDH).** *Let  $\mathbb{G} = (G, g, p)$  be a represented group and let  $\mathcal{A}$  be an adversary. Consider the following experiment, where  $\mathcal{D}$  is a dictionary of  $N$  elements in  $G$ .*

**Experiment  $\text{Exp}_{\mathbb{G}}^{\text{pccdh}}(\mathcal{A}, \mathcal{D})$**   
 $(X, s) \leftarrow \mathcal{A}(\text{find}, \mathcal{D})$   
 $\Pi \xleftarrow{R} \mathcal{D}; Y \xleftarrow{R} G$   
 $K \leftarrow \mathcal{A}(\text{guess}, s, Y, \Pi)$   
**return** 1 if  $K = \text{CDH}(X/\Pi, Y)$

We define the advantage of  $\mathcal{A}$  in violating the PCCDH assumption with respect to the dictionary  $\mathcal{D}$ ,  $\text{Adv}_{\mathbb{G}, N}^{\text{pccdh}}(\mathcal{A}, \mathcal{D})$ , the advantage of  $\mathcal{A}$   $\text{Adv}_{\mathbb{G}, N}^{\text{pccdh}}(\mathcal{A})$ , and the advantage functions,  $\text{Adv}_{\mathbb{G}, N}^{\text{pccdh}}(t, \mathcal{D})$  and  $\text{Adv}_{\mathbb{G}, N}^{\text{pccdh}}(t)$ , as above.  $\diamond$

In our security proofs, we actually need a simple variation of the above problem, in which the adversary, in the second stage, outputs a set of  $s$  candidates for the CDH value. The adversary wins if the set indeed contains  $K$ . This problem is thus named *Set Password-based Chosen-basis Computational Diffie-Hellman Problem (SPCCDH)*.

We define the advantage  $\text{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(\mathcal{A}, \mathcal{D}, s)$  of  $\mathcal{A}$  in violating the SPCCDH assumption with respect to the dictionary  $\mathcal{D}$ , the advantage  $\text{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(\mathcal{A}, s)$  of  $\mathcal{A}$ , and the advantage functions  $\text{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(t, \mathcal{D}, s)$  and  $\text{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(t, s)$  as in previous definitions.

Fortunately, the two new assumptions are not so strong, since one can prove that the SPCCDH problem is equivalent to the CDH problem as proven in Appendix B. The general result proven in Appendix B can be simplified in the particular case of not so large dictionaries:

**Lemma 2.**  $\text{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(t, s) \leq \frac{1}{N} + N^2 s^2 \times \text{Adv}_{\mathbb{G}}^{\text{cdh}}(2t + \tau_e)$ .

In addition to the computational assumptions above, we also make use of the following decisional assumption in our security proofs.

**Definition 3 (PCDDH).** *Let  $\mathbb{G} = (G, g, p)$  be a represented group and let  $\mathcal{A}$  be an adversary. Consider the following experiment, defined for  $b = 0, 1$ , where  $\mathcal{D}$  is the dictionary of  $N$  elements in  $G$ .*

**Experiment  $\text{Exp}_b^{\text{pcddh}}(\mathcal{A}, \mathcal{D})$**   
 $(X, Y, s) \leftarrow \mathcal{A}(\text{find}, \mathcal{D})$   
 $\Pi \xleftarrow{R} \mathcal{D}; s_0, s_1 \xleftarrow{R} \mathbb{Z}_p$   
 $Y' \leftarrow Y^{s_0}; X' \leftarrow (X/\Pi)^{s_b}$   
 $b' \leftarrow \mathcal{A}(\text{guess}, s, X', Y', \Pi)$   
**return**  $b'$

We define the advantage of  $\mathcal{A}$  in violating the PCDDH assumption with respect to the dictionary  $\mathcal{D}$ ,  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(\mathcal{A}, \mathcal{D})$ , the advantage of  $\mathcal{A}$ ,  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(\mathcal{A})$ , and the respective advantage functions of  $\mathbb{G}$  for a given value  $N$ ,  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(t, \mathcal{D})$  and  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(t)$ , as above.  $\diamond$

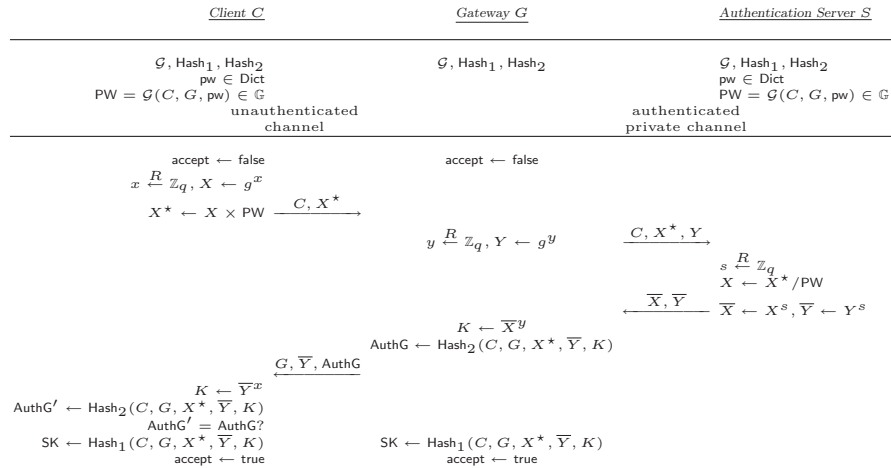
Fortunately again, this problem is not new. It has already appeared in [3] under the name PCDDH2. In that paper, the authors have also shown that  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(t, \mathcal{D})$  and  $\text{Adv}_{\mathbb{G},N}^{\text{pcddh}}(t)$  cannot be significantly larger than  $2/N$ .

## 4 The Gateway PAKE System

In this section, we describe GPAKE, the underlying gateway-oriented password-based protocol used in the construction of our threshold gateway-oriented password-based protocol.

### 4.1 Description

Our gateway-oriented password-based protocol, GPAKE, builds upon previous password-based key exchange protocols in [5, 8, 13], which in turn are based on the encrypted key exchange of Bellare and Merritt [6]. Its description is given in Figure 1, where  $\mathbb{G} = (G, g, q)$  is a represented group;  $\ell$  is a security parameter; and  $\mathcal{G} : \mathcal{U}^2 \times \text{Dict} \rightarrow \mathbb{G}$ ,  $\text{Hash}_1 : \mathcal{U}^2 \times \mathbb{G} \times \mathbb{G} \rightarrow \{0, 1\}^\ell$ , and  $\text{Hash}_2 : \mathcal{U}^2 \times \mathbb{G} \times \mathbb{G} \rightarrow \{0, 1\}^\ell$ , are random oracles.



**Fig. 1.** GPAKE: A gateway-oriented password-based authenticated key exchange protocol.

The protocol consists of four message exchanges, two between the client and the gateway and two between the gateway and the authentication server. Since the channel connecting the gateway to the server is assumed to be authenticated and private, from the client perspective, the protocol resembles almost exactly the 2-party protocol OMDHKE in [8]. The only difference is in the key derivation function, which does not include the password.

The protocol starts with the client choosing a random element  $x \in \mathbb{Z}_p$  and computing  $X = g^x$ , and encrypting it using  $\mathcal{G}(C, G, \text{pw})$  as a mask to obtain  $X^*$ . The client then sends to the gateway both  $X^*$  and its identity string  $C$ . Upon receiving a message from the client, the gateway chooses a random element  $y \in \mathbb{Z}_p$  and computes  $Y = g^y$ , and forwards to the server both  $Y$  and the value  $X^*$  that it has received from the client. Upon receiving the values  $(X^*, Y)$  from the gateway, the server computes  $X = X^*/\mathcal{G}(C, G, \text{pw})$ , chooses a random element  $s \in \mathbb{Z}_p$ , computes  $(\bar{X} = X^s, \bar{Y} = Y^s)$ , and sends it back to the gateway. Upon receiving  $(\bar{X}, \bar{Y})$ , the gateway computes the Diffie-Hellman key  $K = \bar{X}^y$ , the authenticator  $\text{AuthG} = \text{Hash}_2(C, G, X^*, \bar{Y}, K)$  and the session key  $\text{SK} = \text{Hash}_1(C, G, X^*, \bar{Y}, K)$ , and sends  $(G, \bar{Y}, \text{AuthG})$  to the client. Upon receiving  $(G, \bar{Y}, \text{AuthG})$ , the client computes the Diffie-Hellman key  $K = \bar{Y}^x$ , checks whether  $\text{AuthG} = \text{Hash}_2(C, G, X^*, \bar{Y}, K)$ , and sets the session key to  $\text{SK} = \text{Hash}_1(C, G, X^*, \bar{Y}, K)$  if the test passes. The session identification is defined to be the transcript  $(C, G, X^*, \bar{Y})$  of the conversation between the client and the gateway.

## 4.2 Security

**Semantic security.** As the following theorem states, GPAKE is a secure gateway-oriented password-based key exchange protocol as long as the SPCCDH problem is hard in  $\mathbb{G}$ . As shown in Section 3, this is equivalent to assuming that CDH problem is hard in  $\mathbb{G}$ . The proof can be found in the full version of this paper [1]. Nevertheless, we note here that the proof of security assumes Dict to be a uniformly distributed dictionary and of size smaller than  $2^\ell$ .

**Theorem 4.** *Let  $\mathbb{G} = (G, g, q)$  be a represent group of prime order  $q$  and let Dict be a uniformly distributed dictionary of size  $N = |\text{Dict}|$ . Let GPAKE describe the gateway-oriented protocol associated with these primitives as defined in Figure 1. Then,*

$$\begin{aligned} \mathbf{Adv}_{\text{GPAKE}, \mathbb{G}, \text{Dict}}^{\text{ake-ror}}(t, q_{\text{exe}}, q_{\text{fake-C}}, q_{\text{fake-G}}, q_{\text{active}}, q_{\text{test}}, q_{\text{Hash}_1}, q_{\text{Hash}_2}, q_G) \leq \\ \frac{q_{\text{active}}^2 + q_G^2}{q} + \frac{q_{\text{exe}}^2}{q^2} + 2 \frac{q_{\text{Hash}_1}^2 + q_{\text{Hash}_2}^2}{2^\ell} + \\ 2 (q_{\text{Hash}_1} + q_{\text{Hash}_2}) \cdot \mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(t + (4q_{\text{exe}} + 4)\tau_e) + \\ 2 \cdot q_{\text{active}} \cdot (q_{\text{Hash}_1} + q_{\text{Hash}_2}) \cdot \mathbf{Adv}_{\mathbb{G}}^{\text{cdh}}(t + 2\tau_e) + 4 \cdot q_{\text{fake-C}}/N + \\ 2 \cdot q_{\text{fake-G}} \cdot \mathbf{Adv}_{\mathbb{G}, N}^{\text{spccdh}}(t, q_{\text{Hash}_1} + q_{\text{Hash}_2}), \end{aligned}$$

where  $q_G$ ,  $q_{\text{Hash}_1}$ , and  $q_{\text{Hash}_2}$  represent the number of queries to the  $\mathcal{G}$ ,  $\text{Hash}_1$  and  $\text{Hash}_2$  oracles, respectively;  $q_{\text{exe}}$  represents the number of queries to the  $\text{Execute}$  oracle;  $q_{\text{fake-C}}$  and  $q_{\text{fake-G}}$  represent the number of attempts of the adversary to fake the client and the gateway, respectively;  $q_{\text{active}}$  represents the total number of queries to the  $\text{Send}$  oracle;  $q_{\text{test}}$  represents the total number of queries to the  $\text{Test}$  oracle; and  $\tau_e$  denotes the exponentiation computational time in  $\mathbb{G}$ .

*Remark 5.* In the security model presented in Section 2, the adversary is not allowed to corrupt gateway instances. Consequently, the proof of Theorem 4 does not guarantee the security of GPAKE in that scenario. Even though GPAKE appears to be secure in the presence of such adversaries, a new proof of security would be required in this case.

**Key Privacy.** As the following shows, GPAKE achieves the goal of key privacy as long as the DDH problem is hard in  $\mathbb{G}$ .

**Theorem 6.** *Let  $\mathbb{G} = (G, g, q)$  be a represent group of prime order  $q$  and let  $\text{Dict}$  be a uniformly distributed dictionary of size  $N = |\text{Dict}|$ . Let GPAKE describe the gateway-oriented protocol associated with these primitives as defined in Figure 1. Then,*

$$\text{Adv}_{\text{GPAKE}, \mathbb{G}, \text{Dict}}^{\text{ake-kp}}(t, q_{\text{exe}}, q_{\text{test}}) \leq 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{ddh}}(t + (4q_{\text{exe}} + 4)\tau_e),$$

where  $q_{\text{exe}}$  and  $q_{\text{test}}$  represent the total number of queries to the  $\text{Execute}$  and  $\text{TestPair}$  oracles; and  $\tau_e$  denotes the exponentiation computational time in  $\mathbb{G}$ .

The proof of Theorem 6 is in [1]. It is worth mentioning that, when proving the security of GPAKE, we do not give the server access to a  $\text{Send}$  oracle since we assume the server to be honest but curious. We do so because, in the actual implementation of GPAKE, the server is distributed and we assume the majority of them to be honest (see Section 5). We also note that, in order to prove key privacy in scenarios where the majority of servers is corrupted, additional modifications would need to be made to GPAKE. These modifications would include the addition of an authenticated Diffie-Hellman protocol between  $C$  and  $G$  as done in [2] and a proof that the pair  $(\bar{X}, \bar{Y})$  is well formed.

**Server Password Protection.** The server password protection of GPAKE follows directly from the password-based chosen-basis decisional Diffie-Hellman assumption (PCDDH) introduced in [3] and recalled in Section 3. More specifically, it is easy to see that the interaction between the gateway and the server corresponds exactly to the security experiment for PCDDH. Since the security of the latter was shown in [3] to be only negligibly larger than  $2/N$ , where  $N$  is the size of the dictionary, it follows (via a standard hybrid argument as in [3]) that, in each interaction with the server, an adversarial gateway cannot do much better than eliminating two passwords from the list of possible candidates with each interaction. As a result, after  $q$  interactions with the server, the advantage of a malicious gateway would be only negligibly larger than  $2q/N$ .

## 5 The Gateway Threshold PAKE System

In this section, our goal is to distribute the authentication server in the previous gateway-oriented password-based protocol to prevent malicious adversaries that can corrupt up to  $k$  out of  $n$  servers. The solution is robust against static adversaries. The threshold version is *transparent* from the point of view of the client since it communicates only with the gateway. The threshold version is also *transparent* from the point of view of the gateway since a special authentication server, called the *combiner*, is the only server with which the gateway communicates. We also assume that the channel between the gateway and the combiner is authenticated and private. We can use signature and encryption schemes in order to fulfill this requirement using threshold signature and encryption.

**Description.** Let  $\mathbb{G}$  is a cyclic subgroup of prime order  $q$ . We assume that the authentication servers  $\{S_1, \dots, S_n\}$  share a secret ElGamal encryption key  $\text{sk} = x$  using Shamir scheme with threshold  $k$ . Server  $i$  knows  $\text{sk}_i = x_i$  and has a verification key  $\text{vk}_i = v^{x_i}$ , where  $v$  is a generator of  $\mathbb{G}$ . The encryption of the password  $c = \mathbf{E}_{\text{pk}}(\text{PW}) = (e_c, f_c)$  under the public ElGamal encryption key  $\text{pk} = (g, y)$  is authenticated in a public file. At the beginning of the protocol, the combiner,  $S_1$  wlog, receives  $X^* = X \times \text{PW}$ , and  $Y = g^y$ .

**First Stage** The combiner encrypts  $X^*$  into  $(e, f)$  and proves the validity of the encryption using a zero-knowledge proof of equality of discrete-log  $\text{EDLog}_{y,g}(e/X^*, f)$ . Given  $X^*$  and  $(e, f)$ , where  $e = X^* \times y^r$  and  $f = g^r$ , then we have to show that  $\log_y(e/X^*) = \log_g(f)$ . Then, he broadcasts to all servers  $Y, X^*, (e, f)$  and the proof of validity  $\text{EDLog}_{y,g}(e/X^*, f)$ .

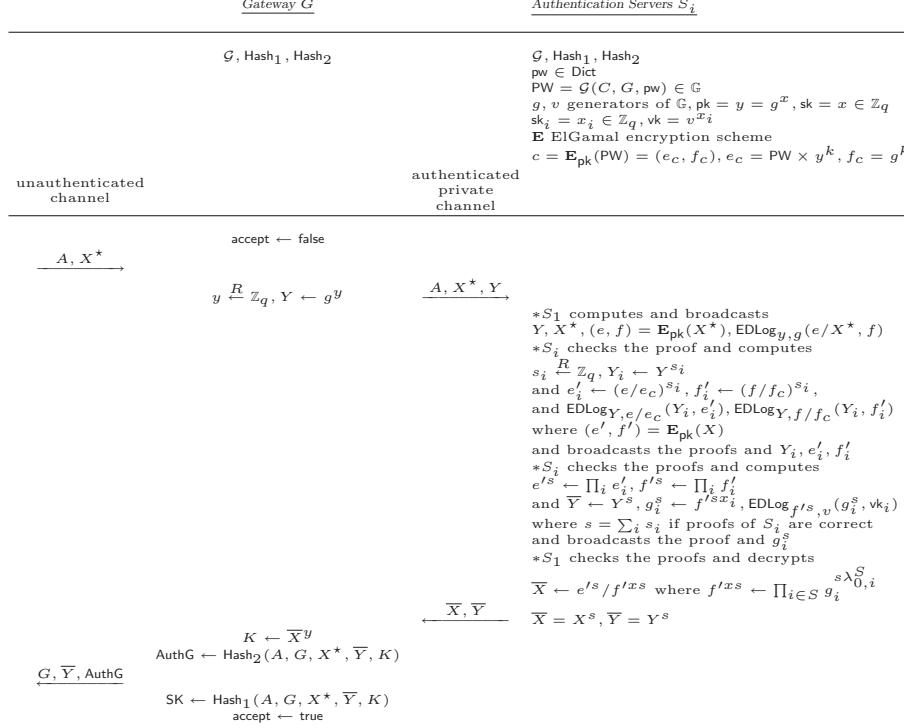
**Second Stage** Each authentication server  $S_i$  checks the proofs and chooses a random value  $s_i \in \mathbb{Z}_q$ . Then, he computes  $Y_i = Y^{s_i}$ ,  $e'_i = (e/e_c)^{s_i}$  and  $f'_i = (f/f_c)^{s_i}$  along with a proof of validity that  $Y_i$  and the two parts of  $(e'_i, f'_i)$  have been raised to the same power  $s_i$ :  $\text{EDLog}_{Y,e/e_c}(Y_i, e'_i)$  and  $\text{EDLog}_{Y,f/f_c}(Y_i, f'_i)$ . All these values are broadcast to all the servers.

**Third Stage** They check all proofs and they compute the values  $e'^s = \prod_i e'_i = e'^{\sum_i s_i}$ ,  $f'^s = \prod_i f'_i = f'^{\sum_i s_i}$ , and  $Y^s = \prod_i Y_i$ . Then, they perform a threshold decryption of  $(e'^s, f'^s)$  by computing  $g_i^s = (f'^s)^{x_i}$ . Next, they prove the validity of the decryption by computing  $\text{EDLog}_{f'^s, v}(g_i^s, \text{vk}_i)$ . Finally, they broadcast the proof along with their decryption share  $g_i^s$ .

**Forth Stage** The combiner, without any secret, can compute  $f'^{xs} = \prod_{i \in S} g_i^{s \lambda_{0,i}^s}$  by using Lagrange interpolation formula and  $\bar{X} = X^s = e'^s / f'^{xs}$  using  $k$  valid decryption shares. He sends to the gateway  $\bar{X}$  and  $\bar{Y} = Y^s$  using the authentication and private channel.

For a pictorial description of our threshold protocol, please refer to Figure 2.

**Security.** We now analyze the security of the threshold variant we just presented. To this end, we show a reduction between an adversary  $\mathcal{A}$  against the



**Fig. 2.** Threshold version of the gateway-oriented password-based authenticated key exchange protocol. Since no change is required on the client side with respect to the non-threshold protocol in Figure 1, the client side has been omitted in the diagram.

threshold scheme and an attacker  $\mathcal{B}$  against the underlying provably-secure GPAKE protocol described in the previous section.

**Theorem 7.** *If the underlying GPAKE protocol is semantically secure and the DDH assumption holds in  $\mathbb{G}$ , then GTPAKE is a semantically secure threshold protocol against a static adversary.*

*Proof.* Our goal is to reduce an adversary  $\mathcal{A}$  against the GTPAKE protocol to an attacker  $\mathcal{B}$  against the GPAKE protocol. We need to simulate the threshold environment for  $\mathcal{A}$  that can corrupt any subset of size at most  $k - 1$  servers from the environment of the attacker  $\mathcal{B}$ . We essentially need to simulate the communications among the servers, *i.e.* the decryption parts at the end of the protocol and the proofs of validity.

Let  $i_1, \dots, i_{k-1}$  be the set of corrupted servers. Recall  $x_i = F(i) \bmod q$  for all  $1 \leq i \leq n$ , and  $x = F(0) \bmod q$ .

To simulate the adversary's view, we simply choose the  $x_{i_j}$  belonging to the set of corrupted servers at random from the set  $\mathbb{Z}_q$ . This allows us to simulate all

the messages, proofs of validity and decryption shares coming from the corrupted servers.

Once these values are chosen, the values  $x_i$  for the uncorrupted servers are also completely determined modulo  $q$ , since we have  $k$  points (the  $k - 1$  points of the corrupted servers and the point 0). The value at point 0 is the decryption value  $f^{sx}$ . However, they cannot be easily computed since  $F(0)$  is secret and corresponds to the secret key of the ElGamal scheme. The ElGamal secret key cannot be chosen by the attacker  $\mathcal{B}$ . Indeed, as all the passwords are encrypted using the ElGamal public key in a public file,  $\mathcal{B}$  cannot know it, unless he recovers the passwords and can easily break the semantic security of the GPAKE scheme. However, we can easily compute the decryption parts  $g_i^s = f^{sx_i}$  of the uncorrupted server by using the value  $f^{xs}$ :

$$g_i^s = f^{sx\lambda_{i,0}^S} \times \prod_{j \in S \setminus \{0\}} g_j^{s\lambda_{i,j}^S}$$

where  $S = \{0, i_1, \dots, i_{k-1}\}$ .

For the “proofs of validity”, one can invoke the random oracle model for the hash function  $H$  to get soundness and perfect zero-knowledge. The soundness is similar to that in Appendix A.

Moreover, the interactive proof system is zero-knowledge against an *honest* verifier since the adversary’s view can be simulated without knowing the values  $x_i$ . This view includes the values of the random oracle at those points where the adversary has queried the oracle, so the simulator is in complete charge of the random oracle. Whenever, the adversary makes a query to the random oracle, if the oracle has not been previously defined at the given point, the simulator defines it to be a random value, and returns the value to the adversary. When we have to perform a fake proof for  $(u_i, \bar{u}_i)$ , since the simulator does not know  $x_i$ , he chooses at random  $c \in \mathbb{Z}_q$  and  $z \in \mathbb{Z}_q$  and defines the values of the random oracle at  $(p, q, g, \bar{g}, u_i, \bar{u}_i, g^z/u_i^c, \bar{g}^z/\bar{u}_i^c)$  to be  $c$ . With all but negligible probability, the simulator has not defined the random oracle at this point before, and so it is free to do so. It is easy to verify that the distribution produced by this simulator is perfect.

Finally, as we need the semantic security of ElGamal encryption scheme, the security is based on the DDH assumption.  $\square$

## Acknowledgements

The authors would like to thank Raymond Choo for pointing out a flaw in a previous version of our schemes that made them susceptible to attacks by corrupted gateways. Such attacks are not currently taken into account by our security model. The first, third, and fourth authors have been supported in part by France Telecom R&D as part of the contract CIDRE, between France Telecom R&D and Ecole Normale Supérieure. The second author is supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, Mathematical



Information and Computing Sciences Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. This document is report LBNL-58175. Disclaimer available at <http://www-library.lbl.gov/disclaimer>.

## References

1. Michel Abdalla, Olivier Chevassut, Pierre-Alain Fouque, and David Pointcheval. A simple threshold authenticated key exchange from short secrets. Full version of current paper. Available from authors' web pages.
2. Michel Abdalla, Pierre-Alain Fouque, and David Pointcheval. Password-based authenticated key exchange in the three-party setting. In Serge Vaudenay, editor, *PKC 2005: 8th International Workshop on Theory and Practice in Public Key Cryptography*, volume 3386 of *Lecture Notes in Computer Science*, pages 65–84, Les Diablerets, Switzerland, January 23–26, 2005. Springer-Verlag, Berlin, Germany.
3. Michel Abdalla and David Pointcheval. Interactive Diffie-Hellman assumptions with applications to password-based authentication. In Andrew Patrick and Moti Yung, editors, *Financial Cryptography 2005*, volume 3570 of *Lecture Notes in Computer Science*, pages 341–356, Roseau, The Commonwealth Of Dominica, February 28 – March 3, 2005. Springer-Verlag, Berlin, Germany.
4. Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 139–155, Bruges, Belgium, May 14–18, 2000. Springer-Verlag, Berlin, Germany.
5. Mihir Bellare and Phillip Rogaway. The AuthA protocol for password-based authenticated key exchange. Contributions to IEEE P1363, March 2000.
6. Steven M. Bellovin and Michael Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *1992 IEEE Symposium on Security and Privacy*, pages 72–84, Oakland, CA, May 1992. IEEE Computer Society Press.
7. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Security proofs for an efficient password-based key exchange. In *ACM CCS 03: 10th Conference on Computer and Communications Security*, pages 241–250, Washington D.C., USA, October 27–30, 2003. ACM Press.
8. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. New security results on encrypted key exchange. In Feng Bao, Robert Deng, and Jianying Zhou, editors, *PKC 2004: 7th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 145–158, Singapore, March 1–4, 2004. Springer-Verlag, Berlin, Germany.
9. David Chaum and Torben P. Pedersen. Wallet databases with observers. In Ernest F. Brickell, editor, *Advances in Cryptology – CRYPTO'92*, volume 740 of *Lecture Notes in Computer Science*, pages 89–105, Santa Barbara, CA, USA, August 16–20, 1992. Springer-Verlag, Berlin, Germany.
10. Mario Di Raimondo and Rosario Gennaro. Provably secure threshold password-authenticated key exchange. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 507–523, Warsaw, Poland, May 4–8, 2003. Springer-Verlag, Berlin, Germany.
11. Shai Halevi and Hugo Krawczyk. Public-key cryptography and password protocols. *ACM Transactions on Information and System Security*, 2(3):230–268, August 1999.

12. Jonathan Katz, Rafail Ostrovsky, and Moti Yung. Efficient password-authenticated key exchange using human-memorable passwords. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 475–494, Innsbruck, Austria, May 6–10, 2001. Springer-Verlag, Berlin, Germany.
13. Philip D. MacKenzie. The PAK suite: Protocols for password-authenticated key exchange. Contributions to IEEE P1363.2, 2002.
14. Philip D. MacKenzie, Thomas Shrimpton, and Markus Jakobsson. Threshold password-authenticated key exchange. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 385–400, Santa Barbara, CA, USA, August 18–22, 2002. Springer-Verlag, Berlin, Germany.
15. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.

## A Basic Tools

### A.1 Threshold Secret Sharing

Let  $q$  be a prime and  $1 \leq k \leq n < q$ . Shamir secret sharing over  $\mathbb{Z}_q$  is a  $k$ -out-of- $n$  sharing where any subset of at least  $k$  parties can recover the secret, but any subset of strictly less than  $k$  parties cannot gain any information about the secret. It is defined as follows: the dealer knows a secret  $x \in \mathbb{Z}_q$ , chooses  $k - 1$  random points  $f_1, \dots, f_{k-1} \in \mathbb{Z}_q$ , sets  $f_0 = x$  and define the polynomial  $F(X) = \sum_{j=0}^{k-1} f_j X^j$ . For  $1 \leq i \leq n$ , let  $x_i = F(i) \in \mathbb{Z}_q$  be the  $i$ -th share of  $x$ . Thanks to the Lagrange equality, we can show that if  $k$  shares are revealed,  $x$  is completely and can be determined by interpolation. For  $S \subset \mathbb{Z}_q$  of cardinality  $k$ , any  $i \in \mathbb{Z}_q$  and any  $j \in S$ , there exists an element  $\lambda_{i,j}^S = \prod_{j' \in S \setminus \{j\}} (i - j') / \prod_{j' \in S \setminus \{j\}} (j - j') \in \mathbb{Z}_q$  such that

$$F(i) = \sum_{j \in S} \lambda_{i,j}^S F(j) \pmod{q}$$

### A.2 Zero-Knowledge Proof of Equality of Discrete Logarithm.

Let  $\mathbb{G}$  be a group of prime order  $q$  with generators  $g$  and  $\bar{g}$ . Let  $\text{EDLog}_{g,\bar{g}}$  be the language of pairs  $(u, \bar{u}) \in \mathbb{G}^2$  such that  $\log_g u = \log_{\bar{g}} \bar{u}$ . We will use a zero-knowledge proof of membership for the language  $\text{EDLog}_{g,\bar{g}}$ . It is not a proof of knowledge since we only want to prove the correctness of the values computed by the authentication servers. We use the zero-knowledge proof system due to Chaum and Pedersen[9]. Although it happens to also be a proof of knowledge, we will not use this property. We describe here the non-interactive version using the Fiat-Shamir heuristic in the random oracle model.

Let  $(u, \bar{u}) \in \text{EDLog}_{g,\bar{g}}$  be given, so there exists  $r \in \mathbb{Z}_q$  such that  $u = g^r$  and  $\bar{u} = \bar{g}^r$ .

- The prover chooses  $t \in \mathbb{Z}_q$  at random, computes  $w = g^t$  and  $\bar{w} = \bar{g}^t$ . He computes  $c = H(p, q, g, \bar{g}, u, \bar{u}, w, \bar{w})$  where  $H$  is modeled as a random oracle. Finally, he computes  $z = t + rc \pmod q$  and sends to the verifier,  $(c, z)$
- The verifier checks whether the following equality holds

$$c = H(p, q, g, \bar{g}, u, \bar{u}, g^z/u^c, \bar{g}^z/\bar{u}^c)$$

It is well-known that the interactive version of this proof is sound since a cheating prover can be accepted only with probability at most  $1/q$ . Assume that  $(u, \bar{u}) \notin \text{EDLog}_{g, \bar{g}}$ , then  $u = g^r$  and  $\bar{u} = \bar{g}^{\bar{r}}$  where  $r \neq \bar{r}$ . If a proof is correct, then there exists a unique  $z$  such that  $w = g^z/u^c$  and  $\bar{w} = \bar{g}^z/\bar{u}^c$ , then  $z - rc = t$  and  $z - \bar{r}c = \bar{t}$ . Therefore, we get  $t - \bar{t} = (\bar{r} - r)c \pmod q$  and since  $\bar{r} - r \neq 0 \pmod q$ , there is a unique value for  $t - \bar{t}$  and so with probability  $1/q$ , the cheating prover is detected.

Finally, the interactive proof system is zero-knowledge against an *honest* verifier

## B Proof of Lemma 2

In this section, we show that the Set Password-based Computational Diffie-Hellman SPCCDH problem is equivalent to the (basic) computational Diffie-Hellman problem CDH: For proving this relation, one simply applies the splitting lemma [15]:

**Lemma 8 (Splitting Lemma).** *Let  $S \subset A \times B$  such that  $\Pr[(a, b) \in S] \geq \alpha$ . For any  $\beta < \alpha$ , define*

$$T = \left\{ (a, b) \in A \times B \mid \Pr_{b' \in B} [(a, b') \in S] \geq \alpha - \beta \right\}$$

*Then (i)  $\Pr[T] \geq \beta$  (ii)  $\forall (a, b) \in T, \Pr_{b' \in B} [(a, b') \in S] \geq \alpha - \beta$ .*

Let  $\mathcal{A}$  be an adversary against the SPCCDH problem, with success probability  $\alpha = 1/N + \varepsilon$ . Then, we can use the splitting lemma, with  $\beta = \varepsilon/2$ , on

$$A = \{(\omega, X, \mathcal{D})\} \text{ and } B = \{1, \dots, N\} \approx \mathcal{D}.$$

Our adversary  $\mathcal{B}$  receives as input a random CDH instance  $(U, X)$ . It chooses a random tape  $\omega$  for  $\mathcal{A}$ , as well as  $N$  random distinct exponents  $u_i \in \mathbb{Z}_p$ . It defines  $U_i = U^{u_i}$ , which specifies the dictionary  $\mathcal{D}$ : with probability greater than  $\varepsilon/2$ , the success probability is greater than  $1/N + \varepsilon/2$ , over the probability space  $B = \{1, \dots, N\}$ . It is thus a multiple of  $1/N$ , not smaller than  $1/N + \nu$ , where  $\nu$  is the maximum in  $\{1/N, \varepsilon/2\}$ . One first simply runs  $\mathcal{A}$  with a random  $k$ , and with probability greater than  $1/N + \nu$ , one gets a first set  $\mathcal{S}_1$  with  $K = \text{CDH}(X/U_k, Y) = \text{CDH}(X/U^{u_k}, Y)$ . One runs  $\mathcal{A}$  again, with another random  $k' \neq k$ , and with probability greater than  $\nu$ , one gets a second set  $\mathcal{S}_2$  with  $K' = \text{CDH}(X/U_{k'}, Y) = \text{CDH}(X/U^{u_{k'}}, Y)$ . Then,  $\text{CDH}(X, U) = (K/K')^{1/(u_k - u_{k'})}$ . By choosing two elements at random in  $\mathcal{S}_1$  and  $\mathcal{S}_2$ , one gets  $\text{CDH}(X, U)$  with probability  $1/s^2$ .