Public-key cryptography and password protocols

Shai Halevi* Hugo Krawczyk†

February 2, 1999

Abstract

We study protocols for strong authentication and key exchange in asymmetric scenarios where the authentication server possesses a pair of private and public keys while the client has only a weak human-memorizable password as its authentication key. We present and analyze several simple password protocols in this scenario, and show that the security of these protocols can be formally proven based on standard cryptographic assumptions. Remarkably, our analysis shows *optimal* resistance to off-line password guessing attacks under the choice of suitable public key encryption functions. In addition to user authentication, we enhance our protocols to provide two-way authentication, authenticated key exchange, defense against server's compromise, and user anonymity. We complement these results with a proof that public key techniques are unavoidable for password protocols that resist off-line guessing attacks.

As a further contribution, we introduce the notion of *public passwords* that enables the use of the above protocols in situations where the client's machine does not have the means to validate the server's public key. Public passwords serve as "hand-held certificates" that the user can carry without the need for special computing devices.

* IBM T.J.Watson Research Center, PO Box 704, Yorktown Heights, New York 10598, USA Email: shaih@watson.ibm.com.

† Department of Electrical Engineering, Technion, Haifa 32000, Israel, and IBM T.J. Watson Research Center, New York, USA. Email: hugo@ee.technion.ac.il.

1 Introduction

In this paper we study the use of human passwords for strong authentication and key exchange in *asymmetric* scenarios where the authentication server can store a strong secret (such as the private key for public-key encryption) while the client uses a weak human-memorizable password as its only authentication key. This asymmetry arises naturally in applications, such as remote user authentication, where the user does not carry any computational device (e.g., a laptop or smart-card) capable of storing a long secret. It also arises in applications of protocols such as SSL, IPSEC and SET in the cases where the client end does not possess a public key[1].

A common problem with password-based methods is the low entropy available in user-chosen passwords, which may be used by an attacker to mount password-guessing attacks. Such attacks may proceed by simply guessing passwords and verifying the guessed value using publicly available information, such as the transcript of a legitimate authentication session between the user and server. This attack is very powerful, since it can be performed off-line, so the attacker does not need to interact with the legitimate parties, and can use a lot of computing power. (In contrast, on-line attacks where the attacker actively tries different passwords against the server are easier to detect and limit.)

The first to deal with the use of public key techniques in conjunction with password authentication were Gong, Lomas, Needham and Saltzer [12]. They suggested that by providing the authentication server with a pair of private/public keys, one could protect weak human passwords against strong attacks via the use of public key encryption. The emphasis of that work was in protecting passwords against off-line password-guessing attacks. Our work extends the public key based approach of [12] in several ways with special emphasis on the analysis and provability of the proposed protocols.

- We consider several simple and intuitive protocols for password authentication, and formally analyze their security. We show that the security of these protocols strongly depends on the choice of the public key encryption function, and demonstrate how some natural and "seemingly secure" realizations of the protocols can be broken. On the other hand, by strengthening the notion of encryption (to resist some form of chosen ciphertext attack) we can formally prove the security of the protocols. In particular, we show optimal resistance to off-line password-guessing attacks.
- We enhance our basic protocols to provide important features such as two-way authentication, authenticated key exchange, resistance to server's compromise, and user anonymity.
- We introduce the notion of *public passwords*, as an enabler for the use of our protocols in situations where the certified server's public key is not available to the client's machine.
- Finally, we prove a general theoretic result, showing that the use of public key techniques is unavoidable in password protocols that provide defense against off line guessing attacks.

Security of password protocols. The main difficulty in designing secure password mechanisms arises from the fact that the space of passwords is usually small and much easier to attack than random cryptographic keys. In particular, exhaustive search attacks as the off-line guessing attacks mentioned above become practical. Moreover, using a low-entropy password as a key to a cryptographic function, can transform an otherwise strong function into a weak one. Namely, when using passwords as cryptographic keys, one makes the assumption that these functions remain secure

[1] SET defines a "certless" option to support these cases. SSL and IPSEC currently do not define such password-based modes of authentication but the need for them has been repeatedly pointed out.

even when the keys are chosen from a very small set. These assumptions are so unusual that, to the best of our knowledge, no one has been able to formally define the requirements from these cryptographic functions under which existing protocols can be proved secure.[2]

Our work avoids these problems by providing mechanisms that do not use the password as a key to cryptographic functions and which we formally prove secure based on standard cryptographic assumptions. Our assurance of security is very strong: we prove that the attacker cannot do better than just trying its luck in active (on-line) impersonation attempts (e.g., by trying to authenticate to the server using a guessed password). Note that if the attacker performs v such attempts and the password is taken from a dictionary of size d then the attack can succeed with probability v/d. We show that additional off-line work by the attacker does not increase this probability (as long as the attacker cannot break the encryption function).

Security of public key encryption. In this work we use public key encryption to design password authentication protocols. Although the protocols themselves are very simple and intuitive, proving their security is not straightforward. In particular, it turns out that the basic notion of security for encryption algorithms (i.e., secrecy against eavesdroppers) is not sufficient to ensure the security of these protocols, and that a stronger notion must be used. We briefly discuss these notions here.

The basic notion of security for public key encryption, due to Goldwasser and Micali [11], essentially means that it is infeasible to derive any partial information on the encrypted plaintext given its ciphertext. In particular, given x_1, x_2, c it should be infeasible to determine whether or not c is an encryption of x_1 or x_2. (This implies that such a secure public-key encryption algorithm must be randomized.) Although this notion provides very high assurance of secrecy protection, it is still not enough to ensure the security of our protocols. Indeed, in Section 3.5 we show how the use of a particular encryption function (which satisfies the Goldwasser-Micali notion) leads to an *insecure* implementation of the protocols. Hence a stronger notion of security of the encryption algorithm is needed. This stronger notion of security, known as resistance to *chosen ciphertext attacks*, was introduced by Rackoff and Simon [22], and is also the subject of [8, 2]. We prove that when using encryption that satisfies this stronger property (e.g., [3, 6]) the protocols that we describe are indeed secure.

Public passwords. Our protocols enjoy several attractive properties and are suited for implementation in cases where the authentication server possesses a public key. However, this requires the client machine to know the *correct* value of this public key. Under certain circumstances, this is possible via a certification of the server's public key by a trusted party or via some other form of trusted distribution to the client machine. However, if a user needs to authenticate from a remote machine that does not have a way to validate the correct public key of the authentication server, then the security of these protocols is in danger. In these cases, we propose to provide the users with a *digest* of the server's public key. This digest, typically of length 60-80 bits, does not need to be memorized by the user. It can be safely written on paper, a plastic card, etc. In Section 4 we discuss several ways to implement these digests; we suggest that the user does not even need to type such a digest, but just recognize it when displayed. We call this digest a *public password*. This notion may be of significant practical value beyond our applications in order to bootstrap trust in public keys before a public key infrastructure is in place (and possibly even after that).

[2] It may be possible to define these requirements in terms of idealized assumptions like the random oracle model; see [19]. However, when we replace these ideal functions with actual cryptographic functions the security of the resultant scheme is unknown.

Applications. The protocols described and analyzed in this paper can have extensive application in scenarios where users authenticate themselves using human-memorizable passwords. For example, they can be used to improve the user authentication techniques in remote applications such as telnet, ftp, etc., and for exchanging a key used to protect the actual data transmitted under these applications. Moreover, our protocols can be used in conjunction with general security protocols, such as IPSEC or SSL, to create "secure sessions" between server and terminal. This is done by first exchanging a key between the terminal and the server using our protocols, namely, where the authentication of the terminal side is done using the user's password. Then this key is used with the IPSEC or SSL mechanisms for the protection of the data itself.

We note that if the terminal does have its own public key, and the server has the means to verify this key, then one can directly use SSL or IPSEC (with their own public-key-based key-exchange mechanisms) to establish a secure session over which the application's data, including user's password, is transmitted. The advantage of this method is that the authentication of the terminal can be based on a strong cryptographic key. It is interesting to note that our analysis of Section 5 shows that the transmission of the user's password for authentication is secure provided that the encrypted channel is secure against chosen ciphertext attacks. Although our analysis there focuses on public key encryption we note that the same analysis is valid for symmetric-key encryption. In the later case, a combination of semantically secure encryption and message authentication provide for the required resistance to chosen ciphertext attacks.

Necessity of public key techniques. All the published solutions for password authentication protocols that resist guessing attacks use public key techniques; we show that this is no accident. Specifically, we prove that every authentication protocol which resists off-line password-guessing attack requires the use of some form of public key techniques. More precisely, we show that given any password protocol resistant to off-line guessing attacks one can build a secure key-exchange protocol (i.e., a protocol that allows parties that do not share any initial secret to exchange a fresh secret via a public and authenticated conversation, as in the case of the Diffie-Hellman protocol). From the outset, this result implies that a secure password protocol is "at least as hard to devise" as a secure key-exchange protocol. Moreover, using a result from [14], our proof implies that the task of building a secure password authentication protocol using only simple symmetric cryptographic primitives, such as hash functions, block ciphers or pseudorandom functions, is of paramount difficulty if not impossible. This provides yet another indication for the inherent complexity of the design of secure password protocols.

Related work. We have already mentioned the work of Gong et al. [12] which was also the first to deal with the problem of guessing attacks against password protocols. Another very influential work, by Bellovin and Merrit [4], introduced Encrypted Key Exchange (EKE) which became the basis for many of the subsequent works in this area, e.g., [5, 15, 23, 19, 21, 24]. The security of these solutions has not been proven and is based on heuristic arguments. Other forms of password authentication, such as one-time passwords, are discussed in Section 2. For a survey of works and techniques related to password authentication see [20, 16].

Organization. In Section 2 we briefly discuss basic password mechanisms and the security requirements for password-based authentication and key exchange. In Section 3 we present our protocols and their extensions, and in Section 4 we expand on the notion of public passwords and discuss some issues regarding their implementation. The formal definitions and a formal proof of security are presented in Section 5. Finally, in Section 6 we prove the necessity of public key techniques for designing password authentication resistant to password-guessing attacks.

2 Password Mechanisms and Their Security

In this section we briefly discuss several mechanisms for password authentication, and describe the notion of security for these mechanisms, which we use in this paper. A reader who is familiar with password security issues can skip most of this section and go right to the description of our notion of security in Subsection 2.3.

2.1 A brief survey

Password transmission. The simplest password mechanism is the transmission of a password in the clear from the user to the server. To validate the password, the server stores a file containing either the plain passwords (attached to the user name) or an image of the passwords under a one-way function. The latter is the classic method of the `Unix` system and is used in remote authentication for functions like `ftp` and `telnet`. In the case of remote authentication the drawback of this mechanism is clear as the password can be easily read by an eavesdropper from the network.

Challenge response. A more secure form of password authentication uses the so called challenge response mechanisms. In this case the password is never transmitted in the clear but is used to compute a secret function on a challenge which is selected by the authentication server with each new authentication instance. This provides freshness for the authentication but leaves the password open to *password-guessing attacks* which work as follows. The attacker is assumed to have access to a relatively small dictionary, containing many common passwords. It first records an authentication session including the challenge and the corresponding response from the user. Later, the attacker tries a set of possible passwords on the challenge to see whether the same response is obtained. If so, the password was found (with high probability). Unfortunately, in reality many passwords are indeed found in such dictionaries, thus the above attack is highly effective.

One-time passwords. One variant of challenge-response mechanisms is the so called "one-time password" authentication (e.g., see [18, 13]), in which the user uses a different password every time it tries to authenticate itself. If these one-time passwords are derived from a human password, the latter is still vulnerable to password-guessing attacks. This can be avoided by providing the user with a list of one-time passwords written on paper. This has the advantage that such a password cannot be re-used (even by the local terminal or by somebody breaking into the authentication server). However it entails the inconvenience for the user of carrying a long list of passwords, the requirement to keep this list safe and secret, and the need to type relatively complex strings into the terminal. In addition, this mechanism is vulnerable to several attacks, ranging from stolen passwords (e.g., copied from the person's paper) to man-in-the-middle attacks, and does not support the important extensions discussed next.

Beyond simple authentication. Password mechanisms can provide additional functionalities on top of one-way authentication of user to server. In particular, we often need them to have the following features.

- *Mutual authentication.* Not only the user authenticates itself to the server, but the server authenticates to the user as well. This is important to avoid man-in-the-middle and server impersonation attacks. The importance of such two-way authentication increases with the need to authenticate remote users over completely untrusted networks like the Internet.
- *Authenticated key-exchange.* At the end of the protocol, the user and the server share a secret session key. This session key is then used to authenticate or encrypt subsequent communica-

tion in the current session. This prevents hijacking of sessions by an intruder, data forgery and data exposure.

- *User identity protection.* An eavesdropper to the authentication protocol does not learn the identity of the user. (This is particularly important with remote authentication of mobile users.)

In all these cases the strength of the authentication provided by the password mechanism is usually the security bottleneck for the added functionalities.

2.2 Security of password authentication

Below we informally describe the notion of security for password authentication which we use in this work. Refer to Subsection 2.3 for a semi-formal definition and to Section 5 for a completely formal one. We begin by presenting a list of basic attacks that a password-based protocol needs to guard against.

- *Eavesdropping.* The attacker listens on the line and tries to learn some useful information from the on-going communication.
- *Replay.* The attacker records messages which were sent in past communications and re-sends them at a later time.
- *Man-in-the-middle.* The attacker intercepts the messages sent between the parties and replaces them with its own messages. It plays the role of the user in the messages which it sends to the server, and at the same time plays the role of the server in the messages that it sends to the user.
- *Password-Guessing attacks.* The attacker is assumed to have access to a relatively small *dictionary* containing common choices of passwords. There are primarily two ways in which the attacker can use the dictionary:
  - *Off-line attack.* The adversary records past communication, and then goes over the dictionary and looks for a password which is consistent with the recorded communication. If such a password is found, the attacker concludes that this is the password of the user.
  - *On-line attack.* The attacker repeatedly picks a password from the dictionary and tries to use it in order to impersonate as the user. If the impersonation fails, the attacker eliminates this password from the dictionary and tries again, using a different password. The standard ways of preventing such on-line attacks in practice are to either limit the number of failed runs that a user is allowed to have before the password is expired, or reduce the rate in which the user is allowed to make login attempts.
- *Insider-assisted attacks.* Quite often, an attacker may recruit the help of insiders to mount attacks. In fact, it is often the case that the attacker is just another user of the system. Therefore, we should take into account the possibility that the attacker may have its own account (or accounts), together with a valid password for this account. We should make sure that even in this case, the protocol prevents an attacker from using a "legitimate" account in order to attack other accounts.
- *Exposure of secrets.* An attacker may occasionally also get access to sensitive data which is supposed to be kept secret at the participating parties (i.e., if the server or the user are compromised). In this case the goal is to minimize the effect that the compromise of any single key or file has on the entire system. (An example for this is the notion of perfect forward secrecy [7].) In particular, in the context of passwords mechanism, one needs to

consider the effect of a compromised password on the derived session-keys (and vice-versa), and the effects of compromising the password file or the secret-key of the server. In principle, the compromise of any of these secrets can potentially influence any of the other secrets, so one needs to consider the following matrix of threats and vulnerabilities:

vulnerability of →	session keys	password
to compromise of →	other session keys	session keys
	passwords	other passwords
	server's password file	server's password file
	server's cryptographic keys	server's cryptographic keys

2.3 A semi-formal definition

Here we sketch our definition for the basic notion of security for a password-based one-way authentication protocol. We present the notion of security through the description of the attacker that we consider. The attacker, which we call an *intruder* I, is allowed to watch regular runs of the protocol between the *user* U and the *server* S, and can also actively communicate with the user and the server in replay, impersonation and man-in-the middle attacks. The intruder I can prompt the parties to initiate new authentication sessions. In each session, I can see all the messages sent between the U and S, and can intercept these messages, change them to any value of its choice, or drop them altogether. It can also send arbitrary messages to the parties (including replay of previous messages). Finally, I also gets to see whether S accepts the authentication or not. On top of these abilities, we also allow the intruder to establish other "accounts" with the server, using its own passwords, run arbitrarily many authentication sessions using these accounts, and use the data so gathered for its attack. (This provision allows the intruder full control over arbitrarily many users.)

For a protocol to be secure in the presence of such a intruder, we require that whenever the server S accepts an authentication session with the user U, it is the case that this user indeed participated in this authentication session. Namely, we require that at the end of each session between U and S, the server outputs (U, sid) and the user outputs (S, sid), where sid is a *session identifier* which is unique to that session (and is the same for the user and the server). We say that the intruder *breaks the authentication protocol* if S either outputs the same pair (U, sid) twice, or if it outputs a pair (U, sid) but user U does not output a matching pair (S, sid).

It is clear that if the user picks its password from a dictionary $\mathcal{D}$, then an intruder that attempts m active impersonation attacks with the server can break the protocol with probability of at least $m/|\mathcal{D}|$ (just by trying in each impersonation attempt a different password from $\mathcal{D}$). A protocol is considered secure if I cannot do significantly better than this trivial bound.

Informal Definition 1 *We say that a password authentication protocol ensures one-way authentication up to ϵ (for some $\epsilon > 0$), if no feasible[3] attacker can break the protocol with probability higher than $\epsilon + \frac{m}{|\mathcal{D}|}$, using only m impersonation attempts with the server.*

We note that keeping the value $m/|\mathcal{D}|$ small should be enforced by auditing mechanisms that either expire a password after a certain number of failed authentication attempts, or limit the rate

[3]The computational "feasibility" of attackers can be formalized as polynomial-time or, concretely, using some specific computational bounds. In particular, we want to consider realistic attackers that can perform an exhaustive search over a password dictionary but cannot break the cryptographic primitives used in the protocol. For example, it is reasonable to assume that an attacker can carry out 2^{40} computational steps, that may be enough to exhaust a password dictionary, but not 2^{80} as may be needed to break some encryption scheme.

in which such attempts can be made.

3 Password Authentication Using the Server's Public Key

Below we present several protocols in the asymmetric scenario where the authentication server is assumed to have a pair of private and public keys while the client authenticates on the basis of a user's password. It is a requirement for the security of these protocols that the server's public key used by the client for encryption be a valid public key for that server. As discussed in the Introduction, when the client machine does not have a way to validate this public key we suggest to provide the user with a *public password* that acts as a "hand-held certificate" for that key (see also Section 4). For the sake of illustration and concreteness, we assume in the description of our protocols the availability of such a public password.

For clarity of presentation and analysis we start by presenting only the basic protocols for user authentication. Later (Section 3.4) we augment these protocols to provide mutual authentication and key exchange. We first introduce some terminology and tools common to our different protocols.

Terminology and tools: The user name is denoted by U and the server's name is denoted by S. A flow of a protocol is denoted with arrows. For example, $S \leftarrow m \leftarrow U$ means message m sent from U to S. The secret password memorized by the user U is denoted by spwd (typically, the value of spwd is computed as a hash value of the user's typed secret password). The public key of the server S is denoted pk_S. The public password of the user is denoted by ppwd. In all the protocols of this section we have $\mathsf{ppwd} = MD(\mathsf{pk}_S)$, where MD is a collision-resistant (or second pre-image resistant) hash function, e.g., SHA-1.

In the protocols below, the symbol n stands for a "nonce", namely, a non-repeating string freshly chosen by S with each protocol execution. This can be implemented as a counter, or as a random string which is long enough as to have only negligible probability of ever repeating. The symbol $\textsc{Enc}$ stands for a randomized encryption scheme, which we assume is resistant to some forms of chosen ciphertext attack (see below). We denote encryption using $\textsc{Enc}$ under the public key of S by $\textsc{Enc}_{\mathsf{pk}_S}$.

In the protocols of Section 3.4 we use families of pseudorandom functions [10], which we denote by $\textsc{prf}$. An individual function in the family is indexed by its key, e.g., $\textsc{prf}_k$. We use pseudorandom functions for key derivation as well as message authentication codes. Typical implementations of pseudorandom functions are based on block ciphers and keyed cryptographic hash functions.

In our description of protocols we omit an initialization flow in which U communicates to S the fact that it wants to talk to S. Also, for simplicity, we omit explicit specification of some of the obvious (yet essential) verification steps by the parties.

3.1 Encrypted password transmission

We start by presenting an extremely simple protocol where the password is encrypted with the server's public key, and then sent to the server for verification. We note that this protocol is reminiscent of the Identification Protocol of [12] (but is even simpler). Later we augment this protocol to provide stronger security properties.

Encrypted Password Transmission

	Set-up: $\mathsf{ppwd} := MD(\mathsf{pk}_S)$	
S		U
1. Pick a nonce n	$\rightarrow \mathsf{n}, \mathsf{pk}_S \rightarrow$	Check $\mathsf{ppwd} = MD(\mathsf{pk}_S)$
2. Verify password	$\leftarrow U, \mathsf{n}, \textsc{Enc}_{\mathsf{pk}_S}(\mathsf{spwd}, \mathsf{U}, \mathsf{S}, \mathsf{n}) \leftarrow$	

Notice that the public password ppwd identifies the public key as being the authentic server's public key, and thus the user can safely use it to encrypt its password. We remark that this simple protocol does not rely on the password as a cryptographic key by itself, thus it avoids the weakness of choosing a cryptographic key from a too small space as discussed earlier in this paper. We also stress again that the encryption scheme is *randomized*, and thus an attacker cannot cannot simply guess spwd and verify the guess by re-encrypting $\textsc{Enc}_{\mathsf{pk}_S}(\mathsf{spwd}, \mathsf{U}, \mathsf{S}, \mathsf{n})$.

An important aspect of this protocol is the use of the *nonce* n sent from S to U. This acts as a proof of freshness without which the authentication could be trivially broken by replaying the ciphertext. However, this protocol uses the encryption function $\textsc{Enc}$ not only to hide the password but also to *bind* the nonce to the password. Thus, it requires that the encryption function will have other properties than simply hiding the encrypted message. For example, it should be infeasible (without knowing spwd) to obtain $\textsc{Enc}_{\mathsf{pk}_S}(\mathsf{n}', \mathsf{spwd})$ from $\textsc{Enc}_{\mathsf{pk}_S}(\mathsf{n}, \mathsf{spwd})$ for some other $\mathsf{n}' \neq \mathsf{n}$. For example, even using a perfect one-time pad encryption of $(\mathsf{n}, \mathsf{spwd})$ would be insecure here since modifying it to an encryption of $(\mathsf{n}', \mathsf{spwd})$ is trivial even without knowing the encryption key. Similarly, ElGamal encryption is vulnerable to such an attack, too. Our analysis below provides a characterization of the additional properties required from the encryption function in order to ensure the security of this protocol.

The above simple protocol is a special case of a broader (and more powerful) family of protocols that we call "encrypted challenge-response" mechanisms and that we discuss next.

3.2 Generic encrypted challenge-response protocol

Here we propose to use the challenge-response approach but to encrypt the user's response under the server's public key as a means to protect against password-guessing attacks. It turns out that this approach, although natural and intuitive, is less straightforward than it may seem at first glance. Indeed, the intuition that merely hiding the response from an attacker should be enough to prevent guessing-attacks is false. To stress this point we present in Section 3.5 a password-guessing attack against a particular "bad implementation" of this protocol, that succeeds even when the public key encryption scheme in use provides provable secrecy protection against eavesdroppers. Fortunately, we are able to show that under a stronger (yet achievable) notion of security for the encryption function, our protocols are provably secure.

Another drawback of regular challenge-response mechanisms is the use of weak human passwords as keys to cryptographic functions; we note that the security of these functions under such a small space of keys is clearly questionable. In contrast, in our basic authentication protocols we use the password under functions that do not have any "cryptographic requirements" related to the password but just require very simple combinatorial properties (e.g., being one-to-one).

We first present a skeleton protocol using a *generic challenge-response* function f that combines the password spwd and challenge n into some response, and analyze the security of this protocol in terms of the structure of f and the security of the encryption function. Armed with this information we then proceed to suggest concrete protocols that achieve secure user authentication.

Generic Encrypted Challenge-Response Protocol

	Set-up: $\mathsf{ppwd} := MD(\mathsf{pk}_S)$		
	S		U
1.	Pick a nonce n	$\rightarrow \mathsf{n}, \mathsf{pk}_S \rightarrow$	Check $\mathsf{ppwd} = MD(\mathsf{pk}_S)$
2.	Decrypt and verify	$\leftarrow U, \mathsf{n}, \textsc{Enc}_{\mathsf{pk}_S}(f(\mathsf{spwd}\ ; U, S, \mathsf{n})) \leftarrow$	

Security of the encryption function ENC. The basic notion of security for public key encryption, called *semantic security*, was introduced by Goldwasser and Micali [11]. In a nutshell, an encryption scheme is said to be semantically secure if, given a public key pk, a ciphertext c and two possible plaintexts x_1, x_2, it is infeasible to determine is c is an encryption of x_1 or an encryption of x_2. (Clearly, such an encryption function must be randomized, so that simply re-encrypting x_1, x_2 and comparing the result to c does not work.) This, in turn, implies other strong properties of the encryption function such as the infeasibility to derive any partial information on the encrypted plaintext given its ciphertext.

As we already said in the introduction, this notion of security is not enough to ensure the security of our protocols. Instead, we require that the encryption function ENC satisfies the stronger notion of resistance to chosen ciphertext attacks, due to Rackoff and Simon [22]. According to this definition, it should be infeasible to determine if c is an encryption of x_1 or x_2, even when you are given some "extra help" in the form of the ability to ask for the decryption of ciphertexts of your choice (but not for the decryption of c itself).

In fact, in this work we use a seemingly weaker form of chosen ciphertext attacks, which we call *ciphertext-verification* attacks. In this form, the "extra help" is limited to the ability to generate pairs (x', c') of plaintext and ciphertext (with $c' \neq c$), and to query whether or not c' is an encryption of x'. Such a query is called a verification query.

Informal Definition 2 *Let* ENC *be an encryption scheme and let $\epsilon > 0$. We say that* ENC *resists ciphertext verification attacks with security ϵ, if no feasible adversary has an advantage of more than ϵ (over a random guess) in deciding if c is an encryption of x_1 or x_2 , even after asking arbitrarily many ciphertext verification queries for pairs (x', c') with $c' \neq c$.*

Stronger notions of security against chosen ciphertext attacks can be found in [22, 8, 3, 2]. In particular, Bellare and Rogaway presented in [3] a simple encoding of data (called OAEP) for use with RSA encryption that provides defense against these strong attacks. Although the analysis of that construction is given on the basis of ideal random functions, it should be considered as a good heuristic and, in particular, advisable for use in our (less demanding) scenario. Also, very recently Cramer and Shoup [6] described a simple encryption scheme which is provable secure against the strongest type of chosen ciphertext attacks without using an "ideal random function".

Structure of the function f. Below we assume that the function $f(\cdot\ ;\cdot)$ has the property that for every fixed strings spwd, x, the induced functions $f(\mathsf{spwd}; \cdot), f(\cdot\ ; x)$ are one-to-one. We say that a function f as above is *one-to-one on its components.* (For example, the concatenation function $f(x; y) = (x, y)$ has this property, as does the XOR function if x and y are of the same length.) We note that in fact, it is sufficient that $f(\mathsf{spwd}; \cdot)$ be collision-resistant and it does not actually have to be one-to-one. Nonetheless, for clarity of presentation we assume below that it is one-to-one.

Informal Theorem 1 *Let ϵ be any positive real number, let $\mathcal{E}$ be an encryption scheme that resists ciphertext verification attacks with security ϵ, and let f be one-to-one on its components. Then the encrypted challenge-response protocol using $\mathcal{E}$ and f ensures one-way authentication up*

to $\epsilon' = \ell \cdot m \cdot \epsilon$, where ℓ is the number of login attempts by the user and m is the number of active impersonation attacks against the protocol.[4]

We present the proof of Theorem 1 in Section 5.3.

Corollary 1 *The encrypted password transmission protocol is secure under the above assumptions on the encryption scheme $\mathcal{E}$.*

Proof. Obviously the concatenation function $f(\mathsf{spwd}\ ; \mathsf{n}, U, S) = (\mathsf{spwd}, \mathsf{n}, U, S)$ is one-to-one on its components. ■

Remark (*user anonymity*). It is possible to derive, from the above generic scheme, protocols in which the user's identity is only sent encrypted in the second flow under the server's public key. In this case, the remote terminal first informs the server of a request for authentication but does not specify the user. In the above protocol, S can send ppwd and the challenge without knowing the identity of the specific user (here we use the fact that all users carry the same value of ppwd). This provides the important *user anonymity* property in cases of remote and mobile authentication.

3.3 Resistance to server compromise

Although the proof of security implies that every implementation of the Encrypted Challenge-Response protocol ensures one-way authentication, different implementations may have very different security properties with respect to the compromise of the secret information stored on the server. (For example, it is clear that the encrypted password transmission protocol of Section 3.1 becomes totally insecure once the private key of the server is compromised).

To protect against compromise of the server, one can use some common heuristics for the definition of the f function. For example, one can set

$$\begin{aligned} p_1 &= H_1(\mathsf{spwd}, U, S) \\ p_2 &= H_2(\mathsf{spwd}, U, S) \\ p_3 &= H_3(p_2, \mathsf{salt}) \\ f(\mathsf{spwd}; \mathsf{n}, U, S) &\stackrel{\text{def}}{=} \langle \textsc{Mac}_{p_1}(\mathsf{n}, U, S), p_2, \mathsf{n} \rangle \end{aligned}$$

and have the server store the values salt, p_3 and p_1 (where H_1, H_2, H_3 are one-way functions, MAC is a message authentication code, and salt is a random string).

The above mechanism defends against compromise of *either* password file or server's private key (but not simultaneously against both)[5]. If the password file is compromised but the server's private key remains secret, then the attacker still needs to mount a password-guessing attack to find p_2. If, on the other hand, the attacker gets the server's private key but does not gain access to the password file, then it still cannot trivially authenticate the user since it needs to be able to compute the value $\textsc{Mac}_{p_1}(\mathsf{n}, U, S)$.

We stress that in this case we are making the heuristic assumption that the attacker cannot break the MAC function in any better way than a password-guessing attack. As said before, this is a non-standard assumption for most MAC functions since they were designed to be keyed over a much larger space. Still, in our case this assumption is not the basis for the authentication security but only a second line of defense in case of server's key compromise.

[4]For example, if ENC has security $\epsilon = 2^{-80}$, the password expires after 100 failed trials, and users change their password before trying 1000 login attempts, then we get $\epsilon' \leq 2^{-80} \cdot 100 \cdot 1000 \approx 2^{-63}$. Thus, no attacker has probability of more than $100/|\mathcal{D}| + 2^{-63}$ for a successful impersonation.

[5]Defense against compromise of both the password file and the server's private key can be achieved by using Lamport's one-time password mechanism [18, 20] instead of the fixed value p_2. This mechanism requires a pre-established limit on the number of password authentications before the password value in the server is to be re-set.

3.4 Mutual authentication and key exchange

Here we add to the above basic authentication protocols the capability of authenticating the server to the user as well as of exchanging an authenticated secret key between the two. This added functionality is needed in many security applications. In particular, our solutions can provide authenticated key exchange for protocols such as IPSEC and SSL where the client's end authenticates via a user's password.

Our extensions for mutual authentication and key exchange follow the general design of SKEME [17]. The basic idea is that the user U adds an encryption of a random key k to the authentication information that it sends to S in the second flow of the protocol. The server S uses this key to authenticate itself, by using k as a key to a MAC function. (It is the sole ability of S to decrypt k which forms the basis for the server's authentication.) A shared key can be derived by applying a pseudorandom function, keyed with the above key k, to the exchanged information. (In actuality, we use the pseudorandom function PRF in this protocol both as a MAC and for key derivation.) Note that here the strength of the authentication from S to U is based on the cryptographic keys of S and hence it is stronger than in other mechanisms that base this authentication on the strength of the user's password. An analysis of this later form of authentication, based on public key encryption resistant to chosen ciphertext attacks, can be found in [1].

Mutual Authentication and Key Exchange

	Set-up: $\mathsf{ppwd} := MD(\mathsf{pk}_S)$		
	S		U
1.	Pick a nonce n	$\rightarrow\ \mathsf{n}, \mathsf{pk}_S\ \rightarrow$	Check $\mathsf{ppwd} = MD(\mathsf{pk}_S)$
2.	Decrypt and verify	$\leftarrow\ U, \mathsf{n}, \textsc{Enc}_{\mathsf{pk}_S}(k, f(\mathsf{spwd}\ ; \mathsf{n}, k, U, S))\ \leftarrow$	Pick random key k
3.	$y := \textsc{prf}_k(\mathsf{n}, S, U)$	$\rightarrow\ y\ \rightarrow$	Check $y = \textsc{prf}_k(\mathsf{n}, S, U)$
4.	Set $k' := \textsc{prf}_k(y)$		Set $k' := \textsc{prf}_k(y)$

Note that the two first flows are the same as in the generic encrypted challenge response protocol of section 3.2 except that the key k is included in the encryption and in the function f.

The above protocol does not provide perfect forward secrecy, since if the servers private key is eventually exposed then the session key k' is revealed. As with any key-exchange protocol, perfect forward secrecy can be added through the use of Diffie-Hellman exchange. The resulting protocol is as follows (below we assume a common prime modulus over which the DH exchange is carried. We omit the mod p notation.)

Mutual Authentication and Diffie-Hellman Key Exchange

	Set-up: $\mathsf{ppwd} := MD(\mathsf{pk}_S)$		
	S		U
1.	Pick n, g^x	$\rightarrow\ \mathsf{n}, g^x, \mathsf{pk}_S\ \rightarrow$	Check $\mathsf{ppwd} = MD(\mathsf{pk}_S)$
			Pick k, g^y
2.	Decrypt and verify	$\leftarrow\ U, \mathsf{n}, g^y, c\ \leftarrow$	$c := \textsc{Enc}_{\mathsf{pk}_S}(k, f(\mathsf{spwd}\ ; \mathsf{n}, g^x, g^y, k, U, S))$
3.	$z := \textsc{prf}_k(c)$	$\rightarrow\ z\ \rightarrow$	Check $z = \textsc{prf}_k(c)$
4.	Set $k' := \textsc{prf}_k(g^{xy})$		Set $k' := \textsc{prf}_k(g^{xy})$

We note that if g^x is chosen at random in every run, then the nonce n is not needed and can be omitted. The derivation of the session key through the application of $\textsc{prf}_k$ to the DH key

g^{xy} is intended to "hash" the DH key into a shorter and stronger key (it also makes the protocol resistant to the breaking of either the Diffie-Hellman exchange or the encryption function, namely, to compute k' an attacker needs to be able to compute g^{xy} and also to find the value k'). Finally, we stress that the information in the second argument of the function f can be hashed under a collision resistant hash function (such as SHA-1) before computing f on it. This preserves the security properties that we prove and shortens the information to fit under the encryption.

As mentioned before, these protocols can provide user anonymity (as required, for example, in IPSEC) by including the user identity under the public key encryption.

3.5 Semantic security is not sufficient

For the proof of Theorem 1 we need to assume that the encryption scheme resists some (weak) form of chosen ciphertext attack. We comment that this requirement is needed for the Encrypted Challenge Response protocol, even if the function f is assumed to have additional cryptographic properties. Specifically, it can be shown that there exists an encryption scheme which preserves secrecy (but is not resilient to chosen ciphertext attack), and for which this protocol is vulnerable to a password-guessing attack, regardless of the choice of the function f (as long as f is a deterministic function).

To see this, let ENC be any encryption that encrypts bit-by-bit (e.g., the encryption scheme in [11], which is proven to be semantically secure). To attack the protocol that uses this scheme, the intruder I records the server's message $x = \langle \mathsf{pk}_S, \mathsf{n} \rangle$ and intercepts the user's response $c = \textsc{Enc}_{\mathsf{pk}_S}(f(\mathsf{spwd}\ ; U, S, \mathsf{n}))$. It then modifies the response to get c' where c' is the same as c, except that the encryption of the last bit of $f(\mathsf{spwd}\ ; U, S, \mathsf{n})$ is replaced by an encryption of the bit '0'. The modified response is sent by the attacker to S. Depending on whether the server accepts or not, the intruder now knows the least significant bit of $f(\mathsf{spwd}\ ; U, S, \mathsf{n})$, and it can use this to eliminate (approximately) half of the passwords in the dictionary. This can be repeated with different challenges, until only a single password remains in the dictionary.

One should note that as opposed to an on-line password-guessing attack, this attack only requires about $\log |\mathcal{D}|$ attempts before the password is revealed (e.g., 20 attempts for a dictionary of one million passwords).

4 Public Passwords

In order to enable the use of our secure protocols in cases where the client's machine cannot verify the authenticity of the server's public key, we suggest to provide the user with a hashed version of this public key. We call this information a "public password". This results in an extension to the usual human-password paradigm, where the user carries not only a secret password, but also a public password. The latter *requires no secrecy* protection but requires integrity. The public password should be short enough so that a human user is able to recognize it if displayed, or even to type it in if requested to do so, but it does not need to be memorized and can be safely written down on a piece of paper, a sticker, a plastic card, etc.

In our applications, the public password serves as "hand-held certificate" for a public key, which the user can conveniently carry with him. Whenever presented with the actual public key (e.g., after being transmitted to the user's terminal) the user can verify the validity of the public key against the hand-held certificate. This enables a human user to participate in protocols that otherwise would be impossible to carry out without a memory device. This notion may be useful in other scenarios as well. This solution is suited, for example, to credit-card applications, where the hand-held certificate can be recorded on the credit-card itself. Moreover, even when public key

infrastructure is available, hand-held certificates may be useful as a supplement to the trust level offered by other mechanisms such as certification authorities (e.g., X.509), distributed directories (e.g., Secure DNS), and others.[6]

As said, we use public passwords as digests of public keys. Hence, it should be infeasible to find a second public key that hashes to the same public password. The length of the public password depends on the amount of trust that we have on the party generating these public keys. If this party (user or server) is trusted not to look for collisions in the hash function during the process of key generation, then the public password needs only to resist "second preimage attacks". That is, it should be infeasible – given a public key pk – to find another pk' such that $H(\mathsf{pk}) = H(\mathsf{pk}')$. In this case, a public password of 60 to 80 bits will suffice. If the generator of the public key is not trusted then H needs to be fully collision resistant and then its output should be in the 120-160 bit range. For the uses in this paper it seems reasonable to assume that the key-generation process is done properly.

4.1 Representation and identification of public passwords

Even though public passwords are short enough to be carried by a human being, they usually represent unstructured strings. Thus, for a user to be able to read, recognize, and type the public password, it is advisable to have a user-readable format for these passwords. A representation for mapping arbitrary binary strings into easy-to-read (and write) words was introduced in the context of one-time passwords [13]. This solution defines a dictionary of 2048 words (mostly English words, 2 to 4 letter long) and a mapping of each 11-bit string to a different word in the dictionary. Thus a 66 bit string is represented by 6 words from that dictionary, e.g., `moss mont sit rear rage pit`.

Such a representation could be used by a public password system as well. Of course, many other representations are possible, e.g., using just alphanumerics (without case distinction) would require about 12 characters to represent 60-bit strings, e.g., `a6et qw29 hzjv`.

Another difference between the public passwords and the secret ones is that in our applications there is no need for the user to type in the public password. Consider again the case of a public password ppwd consisting of a hash of a public key pk. The latter is stored in some remote machine and sent over the network to the local terminal where the user is working. The terminal computes the hash of the public key and then compares the computed value with the public password ppwd. This can be done by having the user enter ppwd, but also by just displaying the computed hash on the screen and asking the user to approve it. Moreover, many users may be able, after some time, to recognize the right value without even carrying it with them.

In this case, however, it is important that a user carefully checks for the validity of the displayed value. Thus, the user-interface should be designed carefully to avoid the tendency of users to answer every question by simply hitting the `Enter`-key. An example of one possible user-interface is to display five strings to the user, one of which is the correct password, and have the user type the number corresponding to the right public password. For instance, if the public password is `moss mont sit rear rage pit`, the user may be presented with:

```
1. eddy weak half net ohio ok
2. moss mont sit rear rage pit
3. ivan laud loy an gal but
```

[6] One can envision applications of this notion where users carry hand-held certificates for public keys of a few entities with which they interact frequently (e.g., the public key of their administrative domain, etc.), thus avoiding the need to rely on public-key infrastructure in every connection with these entities. This is somewhat similar to carrying a short list of often used telephone numbers to avoid the need to refer to the directory for every phone call.

```
4. bloc ave fire grad beef aye
5. vary hone ton limb pry stew
```

to which he is required to answer with '2'. We stress that it is important that the user does not blindly decide by a matching prefix or so, as this would help an attacker in delivering a public key of its choice. Thus, in the above example it may be more effective to present the user with strings that are actually visually related to each other (and only one being the right value). There may be other graphical encodings of bits that will be even more easily recognizable by the user. Some recent work on a similar recognition problem can be found in [9].

5 Formal definitions and proofs

In this section we formally define our notion of security for password protocols and prove that the Encrypted Challenge-Response protocol is secure, provided that the encryption scheme used in it resists "ciphertext-verification attacks".

5.1 One way password authentication

We define a general model of *one-way password authentication*, intended to capture the realistic scenarios where such authentication protocols are run and to establish their security requirements. This definition includes the description of the parties to the protocol, the attacker and its capabilities, and the notion of security. Here we restrict ourselves to one-way authentication and do not formalize the notions of security related to more general tasks such as two-way authentication, defense against compromised servers, and key exchange (see Section 5.4).

Before formally describing our model, we motivate some of its elements. Three important aspects captured by this model are: (i) the existence of many users in the system, (ii) the fact that different runs of the protocol, called *sessions*, can be executed sequentially and/or concurrently by one or more users (e.g., the same or different users can have open authentication sessions with the server at the same time), and (iii) the realistic capabilities of attackers which are powerful but not omnipotent. We assume that the attacker not only controls the information transmitted over the communication lines (i.e., a "man in the middle") but can also corrupt and control some of the users of the system. In the later case the attacker knows, and even chooses, the secrets of these parties. The computational power of the attacker is large but limited. We call this a *feasible* attacker. Our treatment and analysis allows for formalizing the exact notion of feasibility via polynomial-time computation or using concrete time bounds. We do not specify this here; the underlying assumption is that the attacker's power does not suffice to break the cryptographic primitives (e.g., encryption) used in the protocol.

For simplicity, and to strengthen the model, we assume that all the users of the system, except for one, are controlled by the attacker; yet, it should be infeasible for the attacker to impersonate the legitimate user to the server. Thus we consider three parties: the (legitimate) user, the server, and the attacker. Among the capabilities of the latter will be the ability to create and register additional (corrupted) users. For reasons of simplicity, we consider a single server but we stress that extensions to more servers is straightforward. This server is assumed to be secure, i.e. not controlled by the attacker.

We now turn to describe the formal details of our authentication model. Following the common practice in cryptography, we define this model by means of a "probabilistic game" involving the legitimate parties and the attacker. We call this a probabilistic game since the parties are allowed to use randomized algorithms in their execution of the game, i.e., they possess a source of "random

coins". The rules of the game define the capabilities of the attacker, and also define what it means for the attacker to "win" (i.e., to break the scheme). A scheme is deemed secure if feasible attackers can only win with very small probability.

5.1.1 The one-way password authentication game

The players. The players in the probabilistic game are the *user*, the *server*, and an *intruder*. A password-based scheme (U, S) specifies the protocols followed by the user and the server, respectively, whereas the intruder can follow any arbitrary (feasible) protocol (denoted I). All the protocols in this game (user, server and intruder) are *message-driven*. That is, for a given internal state and input message, the protocol specifies the changes in the state and the resulting output message. (Note that we identify the user, server and intruder with the protocols that they run.)

The game. Each game is parameterized by a security parameter k (which controls the strength of the underlying cryptographic functions and keys), and a "dictionary" $\mathcal{D}$, containing a set of possible passwords. We assume that k and $\mathcal{D}$ are known to everyone (including the intruder). The game itself proceeds as follows.

First, there is a set-up phase, in which passwords and cryptographic keys are selected. In this phase, the server picks a name, which is an arbitrary string denoted by S, and publishes it. The protocol for the server can also specify the choice by S of secret and/or public cryptographic keys. If public keys are used then S publishes them too. Then, the user picks an arbitrary user-name, U, and a password spwd from $\mathcal{D}$.[7] The user publishes U and gives spwd to the server while keeping it secret from the intruder. The intruder can also register additional users with the server at any time (before, during, or after the set-up phase), by picking any pair of user-name U' and password spwd' (provided that $\mathsf{U}' \neq \mathsf{U}$), publishing U' and giving spwd' to the server.

After the set-up phase, the intruder I has full control over all the "parties" U' it created, as well as over the communication between U and S. To model the ability of the intruder to control the communication lines, we formally let U and S communicate only through I. Every message that U and S send goes to I, and every message they receive comes from I. The intruder I may choose to forward the messages between U and S unchanged, or it may choose to modify or eliminate some messages or send other messages instead (including replay of previous messages). At any time, I can send special "prompt" messages to the parties, causing them to start new authentication sessions (in particular, several simultaneous sessions by the same or different users are possible). Sessions have unique identifiers (i.e., different sessions have different identifiers). This game is run until the intruder I decides to halt.

To capture our security requirements, we require that the protocols for U and S specify some outputs. These are special outputs, different than the messages exchanged by the parties, and intended to record events related to the security of the authentication. User U outputs a pair $(\mathsf{S}, \mathsf{sid})$ whenever it authenticated itself to server S under session identifier sid. Server S outputs a pair $(\mathsf{U}, \mathsf{sid})$ whenever a successful authentication by user U has been verified during session sid. If an attempt to authenticate by (alleged) user U during session sid fails (i.e. is not verified as correct by S) then S outputs $(\mathsf{U}, \mathsf{sid}, \perp)$. The latter is needed, so that we can count the "number of failed authentication attempts", after which the password must expire.

[7] For simplicity we assume that the user chooses a passwords from $\mathcal{D}$ uniformly at random; however, our analysis can easily accommodate other probability distributions over $\mathcal{D}$.

5.1.2 Security

Based on the above game we define what is considered as a secure one-way password authentication protocol (U, S). We first develop some terminology regarding such protocols and the actions of the intruder I.

- A protocol (U, S) is said to be *syntactically correct*, if whenever all the messages between U and S in a session sid are passed unchanged, then S and U output $(\mathsf{U}, \mathsf{sid})$ and $(\mathsf{S}, \mathsf{sid})$, respectively. (This means that if no attacker is active against the protocol then the authentication should succeed. We will use this as a "non-triviality" condition to prevent calling "secure" protocols where S is specified to reject all authentications.)
- An event in which the server outputs $(\mathsf{U}, \mathsf{sid})$ but the legitimate user U does not output a pair $(\mathsf{S}, \mathsf{sid})$ is called a *successful impersonation*. An event in which the server outputs $(\mathsf{U}, \mathsf{sid}, \perp)$ is called an *authentication failure*. An event in which the server outputs a pair $(\mathsf{U}', \mathsf{sid})$ after already outputting some other pair $(\mathsf{U}'', \mathsf{sid})$ in the past, is called a *successful replay*. (Here $\mathsf{U}', \mathsf{U}''$ are arbitrary users, and sid is the same in both pairs.) We refer to all the above events as *active impersonation attempts*.
- An (ℓ, m)-*run* of the game is a run with at most m active impersonation attempts, and in which the legitimate user outputs at most ℓ pairs $(\mathsf{S}, \mathsf{sid})$. An (ℓ, m)-*win* for the intruder is an (ℓ, m)-run which contains at least one successful impersonation or replay event.

Our definition of security essentially says that the "best" possible strategy for the intruder is to actively try passwords with the server until the right password of the user is found. (As discussed earlier in the paper, such an attack can be thwarted by limiting the number of "authentication failures" allowed to each user.)

Definition 1 *Let $\epsilon(\cdot, \cdot, \cdot)$ be a positive real function, and let (U, S) be a syntactically correct protocol. We say that (U, S) ensures one-way password authentication up to ϵ, if for every feasible intruder I, every finite dictionary $\mathcal{D}$, every value k for the security parameter, and every ℓ, m, we have*

$$\Pr\left[(\ell, m)\text{-}win \text{ for } I\right] \leq \frac{m}{|\mathcal{D}|} + \epsilon(k, \ell, m)$$

where the probability is taken over the random coins of U, S and I in an execution of the above probabilistic game with dictionary $\mathcal{D}$ and security parameter k.

5.2 Security of the encryption function

Here we give a formal definition for our notion of a public key encryption scheme which resists *ciphertext-verification attacks*. Let $\mathcal{E} = (\textsc{Gen}, \textsc{Enc}, \textsc{Dec})$ be an encryption scheme, where $\textsc{Gen}$ is the key-generation algorithm, $\textsc{Enc}$ is the (probabilistic) encryption algorithm and $\textsc{Dec}$ is the decryption algorithm. A ciphertext-verification attack is formally defined via the following game, which involves the three algorithms and an adversary A.

1. The key-generation algorithm is run (with security parameter k) to generate a secret/public key pair, $(\mathsf{sk}, \mathsf{pk})$. The adversary A is given the public key pk.
2. The adversary adaptively generates queries (x_i, c_i). Below we refer to these as *verification queries*. For each verification query (x_i, c_i) the adversary is told whether or not $x_i = \textsc{Dec}_{\mathsf{sk}}(c_i)$.
3. The adversary generates a pair of plaintexts x_1, x_2 of the same length, and asks for an encryption of one of them. Below we call this the *test query* of A. With probability $1/2$, A gets an answer $c = \textsc{Enc}_{\mathsf{pk}}(x_1)$, and with probability $1/2$ it gets $c = \textsc{Enc}_{\mathsf{pk}}(x_2)$.

Server algorithm $S(\mathcal{D}, k)$: (Server name denoted by S)

Setup: pick $(\mathsf{sk}_S, \mathsf{pk}_S)$ with security parameter k
publish pk_S and record sk_S
accept and record pairs $(\mathsf{U}, \mathsf{spwd})$

Operation: when prompted to start a new session,
pick a new session-id sid, record it and send it to the initiator of the session
upon receipt of a message $\langle \mathsf{U}, \mathsf{sid}, a \rangle$,
if you have records for $(\mathsf{U}, \mathsf{spwd})$ and sid
if $\textsc{Dec}_{\mathsf{sk}_S}(a) = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$
output $(\mathsf{U}, \mathsf{sid})$ and delete the record for sid
else output $(\mathsf{U}, \mathsf{sid}, \perp)$
else ignore incoming message

User algorithm $U(\mathcal{D}, k)$: (User name denoted by U)

Setup: pick $\mathsf{spwd} \leftarrow \mathcal{D}$, publish U and secretly give spwd to server
record server's name and public key $(\mathsf{S}, \mathsf{pk}_S)$

Operation: upon receipt of a message sid,
set $a = \textsc{Enc}_{\mathsf{pk}_S}(f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}))$
send $\langle \mathsf{U}, \mathsf{sid}, a \rangle$ to server
output $(\mathsf{S}, \mathsf{sid})$

Figure 1: The Server- and User-algorithms of the Encrypted Challenge-Response protocol.

4. The adversary may adaptively generate more verification queries (x_i, c_i), subject to the constraint that $c_i \neq c$. Again, for each verification query (x_i, c_i), the adversary is told whether or not $x_i = \textsc{Dec}_{\mathsf{sk}}(c_i)$.
5. The adversary A guesses whether c is an encryption of x_1 or x_2.

Definition 2 *An encryption scheme* $\mathcal{E} = (\textsc{Gen}, \textsc{Enc}, \textsc{Dec})$ *is said to* resist ciphertext-verification attacks *with security* $\epsilon = \epsilon(k)$ *if for any feasible adversary* A,

$$\Big| \Pr[A \textit{ guesses "encryption of } x_1\textit{"} \mid \textsc{Dec}_{\mathsf{sk}}(c) = x_1] - \Pr[A \textit{ guesses "encryption of } x_1\textit{"} \mid \textsc{Dec}_{\mathsf{sk}}(c) = x_2] \Big| \leq \epsilon$$

where the probabilities are taken over the execution of $\textsc{Gen}$, *the random coins of* A, *and the randomness used in Step 3 above. The difference between these two probabilities is called the* advantage *of* A.

5.3 Proof of Theorem 1

We start by casting the Encrypted Challenge-Response protocol in the syntax of Definition 1. The server and user protocols are specified in Fig. 1. This description assumes that the user stores the server's public key. As we discussed earlier, this requirement can be avoided using public passwords. Using the syntax of Definitions 1 and 2, we re-state Theorem 1 as follows:

Theorem 1 *Let $\mathcal{E}$ be an encryption scheme that resists ciphertext-verification attacks with security $\epsilon(k)$, and let f be a function which is one-to-one on its components. Then, the Encrypted Challenge-Response protocol (U, S) with encryption $\mathcal{E}$ and function f ensures one-way password authentication up to $\epsilon'(k, \ell, m) = m \cdot \ell \cdot \epsilon(k)$.*

In the following proof of Theorem 1 we formulate several auxiliary Lemmas that we later prove in Subsection 5.3.1.

Proof We begin by setting a few notations. For the rest of the proof, fix the security parameter k and the dictionary $\mathcal{D}$. Now consider a run of the game between U, S and I. We denote the messages that the server sends in response to prompts for a new session by $\mathsf{sid}_1 \ldots \mathsf{sid}_n$, the messages that the intruder sends to U by $\mathsf{sid}'_1 \ldots \mathsf{sid}'_{n'}$, the responses of U to these messages by $y'_1 \ldots y'_{n'}$, and the messages that I sends to S by $y_1 \ldots y_{n''}$.

In what follows, we use the convention that U represents the legitimate user not corrupted by the intruder. Recall that each reply y_i from I to S is of the form $\langle \mathsf{U}', \mathsf{sid}', a' \rangle$, for some user U'. The replies for which $\mathsf{U}' = \mathsf{U}$ play a special role in the proof. Below we call them U-*replies* (i.e., y is a U-reply if $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$ for some values of sid, a).

A U-reply $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$ is considered *successful* for the intruder, if it is different than all the replies of the user U (i.e., $y \neq y'_j$ for all $j = 1 \ldots n'$) and yet it is accepted by the server (namely, the server outputs the pair $(\mathsf{U}, \mathsf{sid})$). The intruder I is said to *fool the server* if it generates at least one successful U-reply. Our first observation is that the intruder must fool the server in order to win the game.

Lemma 2 *If a run of the game between U, S and I is a win for the intruder I, then in this run the intruder I fools the server.*

Our next observation is that without loss of generality, we can assume that the intruder I never forwards an unmodified U-reply to the server (since by definition this U-reply cannot be successful), and also never replays the same message twice to the sever (since the server always reject replays). Formally, we say that an intruder I is *restrained* if it never forwards an unmodified U-reply to S and never replays the same message twice to S.

Lemma 3 *Let ℓ, m be two integers and let p be a probability $(0 \leq p \leq 1)$. If there exists an intruder I such that $\Pr[I$ fools the server in an (ℓ, m)-game$] = p$, then there also exists a restrained intruder I' such that $\Pr[I'$ fools the server in an (ℓ, m)-game$] = p$.*

For the rest of the proof, we consider only restrained intruder algorithms. We note that for such algorithms, any U-reply that is accepted by the server is necessarily a successful one. We continue by way of contradiction. Namely, we show that if there exists a restrained intruder I that fools the server with probability of more than $v/|\mathcal{D}| + \epsilon'$, then there also exists an adversary A which has advantage of more than $\epsilon = \epsilon'/\ell m$ in guessing the correct answer in a ciphertext-verification attack on $\mathcal{E}$.

The proof proceeds with two lemmas. In the first lemma, we show that if it is possible to fool the server in an (ℓ, m)-run with probability p, then it is also possible to fool it with probability p/m in an $(\ell, 1)$-run. That is, instead of sending to the server m U-replies, the intruder sends only a single U-reply, and the success probability of the intruder then decreases from p to p/m.

Lemma 4 *Let ℓ, m be two integers and let p be a probability $(0 \leq p \leq 1)$. If there exists a restrained intruder I, such that $\Pr[I$ fools the server in an (ℓ, m)-run$] = p$, then there also exists a restrained intruder I' such that $\Pr[I'$ fools the server in an $(\ell, 1)$-game$] \geq p/m$.*

In particular, Lemma 4 implies that if it is possible to fool the server in an (ℓ, m)-run with probability $m/|\mathcal{D}| + \epsilon'$, then it is also possible to fool it with probability $1/|\mathcal{D}| + \epsilon'/m$ in an $(\ell, 1)$-run. The next and final lemma shows that if the latter is possible, then it is also possible to achieve advantage $\epsilon'/\ell m$ in a ciphertext-verification attack against the encryption scheme $\mathcal{E}$.

Lemma 5 *Let ℓ be an integer and let δ be a real number ($0 \leq \delta \leq 1$). If there exists a restrained intruder algorithm I such that $\Pr[I$ fools the server in an $(\ell, 1)$-run$] = \delta + 1/|\mathcal{D}|$, then there also exists an adversary A, which has advantage of at least δ/ℓ in a ciphertext-verification attack against the encryption scheme $\mathcal{E}$.*

The combination of Lemmas 2 through 5 yield Theorem 1 as follows: Assume that there is an intruder that wins in an (ℓ, m)-run with probability more than $\epsilon' + m/|\mathcal{D}|$. By Lemma 2, it means that I also fools the server with the same probability. Lemma 3 then implies that there exists a restrained intruder I' that fools the server with the same probability, and by Lemma 4, there also exists a restrained intruder algorithm I'' that fools the server with probability more than $\epsilon'/m + 1/|\mathcal{D}|$ in an $(\ell, 1)$-run. Finally, Lemma 5 implies that there exists an adversary A which has an advantage of more than $\epsilon'/\ell m$ in guessing the right answer in a ciphertext-verification attack against $\mathcal{E}$, contradicting our assumption that $\mathcal{E}$ resists ciphertext-verification attacks with security $\epsilon = \epsilon'/\ell m$. ■

5.3.1 Proofs of the lemmas

Proof of Lemma 2 Consider a run of the game in which I *does not* fool the server, and we prove that I does not win in this run of the game. First, since all the session-id's that S chooses are unique, and since S discards session-id sid once it accepts it, then it never outputs the same value sid twice. It is left to show that whenever the server outputs a pair $(\mathsf{U}, \mathsf{sid})$, the user outputs a matching pair $(\mathsf{S}, \mathsf{sid})$. To see that, notice that the server only outputs the pair $(\mathsf{U}, \mathsf{sid})$ after it receives a U-reply $\langle \mathsf{U}, \mathsf{sid}, a \rangle$. If I does not fool the server, then all these U-replies that S accepted indeed came from U and were not modified. But the user U always output the pair $(\mathsf{S}, \mathsf{sid})$ after sending $\langle \mathsf{U}, \mathsf{sid}, a \rangle$. ■

Proof of Lemma 3 Consider a run τ in which I fools the server, and let $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$ be a successful U-reply in this run. Clearly, y was not replayed, since the server always rejects (or ignores) replayed messages. Also, y were not forwarded unmodified from U (by the definition of a successful reply).

Consider now the run τ' which results by omitting all the replays to the server as well as all the U-replies that were forwarded unmodified from U to S. This run would also include the U-reply $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$, and moreover, this U-reply would still be successful for the intruder. To see that, notice the following points:

- The server in τ' is prompted for the same number of sessions as in τ, and so it produces the same set of session-id's. In particular, if the server in τ has a record for session-id sid, then so does the server in τ'.
- The communication between the intruder and the user U in τ' is the same as in τ. Therefore, if y is not equal to any of the user replies in τ, then it still is not equal to any of the user replies in τ'.
- If y is successful in τ then it means that $\textsc{Dec}_{\mathsf{sk}_S}(a) = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$, and this still holds true in τ'. Therefore, the server in τ' also accepts y and outputs the pair $(\mathsf{U}, \mathsf{sid})$.

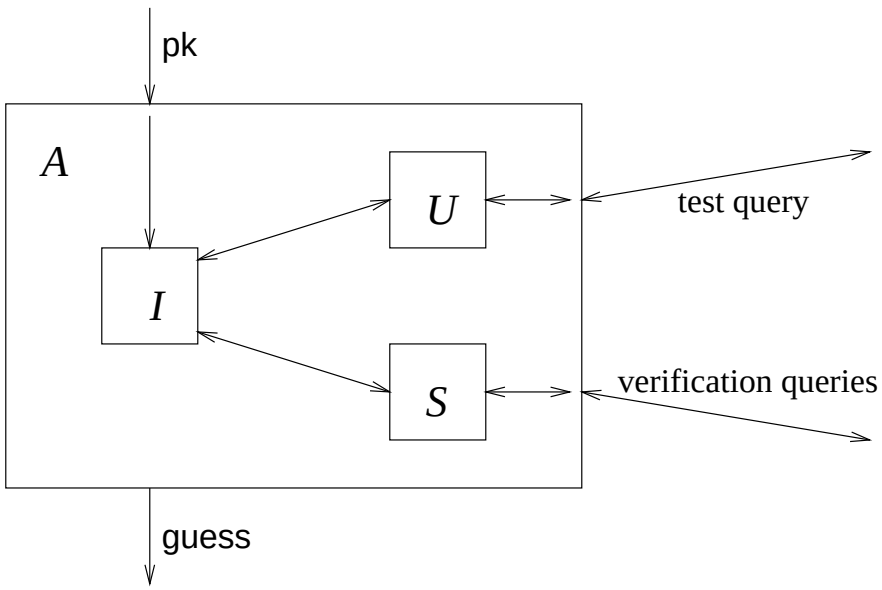

Figure 2: The reduction of Lemma 5.

Notice also that if τ is an (ℓ, m)-game, then so is τ' (since it has at most the same number of outputs for the user and server as τ).

Consider, therefore, an intruder I' that behaves exactly like I, except that it does not send any of the replayed messages and unmodified U-replies that I does. Note that for every I-message that I' does not send, I' knows what would have been the server's response to that message had it been sent. Namely, it would be rejected if it is a replay, and it would be accepted if it is an unmodified U-reply that is not replayed.

If follows that if we have $\Pr[I$ fools the server in an (ℓ, m)-game$] = p$, then the same still holds for I'. ■

Proof of Lemma 4 The intruder algorithm I' is almost the same as I, except for the way it handles U-replies. At startup time, I' picks a random index $i \in \{1, \ldots, m\}$. Then, it behaves exactly as I does, except that I' only sends to the server the i'th U-reply (among all the U-replies that I sends). For all the other U-replies, I' does not send them, and instead behaves as if they were sent and rejected by the server.

Let τ be an (ℓ, m)-run in which I fools the server, and let j be the the index of the first successful U-reply that I sends to the server in this run, and denote this reply by y. Consider now what happens in the corresponding run using I', assuming that I' happens to choose the index $i = j$. Recall that I' behaves just like I except for the U-replies that I sends to the server. Notice also that the first $j-1$ such U-replies were rejected in τ (since the j'th U-reply if the first successful one), and that since $i = j$ then I' behaves as if these U-replies were indeed rejected. Hence, the actions of I' exactly mirror those of I, and in particular I' produces the same U-reply y. Using the same arguments as in the proof of Lemma 3, this U-reply would still be successful.

We conclude that whenever I fools the server in a run τ, and I' chooses the index of the first successful U-reply in that run, then I' would also fool the server. Therefore, if I has probability p of fooling the server with m U-replies, then I' has probability at least p/m of fooling the server with a single U-reply. ■

Proof of Lemma 5 We start by describing the algorithm of the adversary A, and then we analyze it to show that it satisfies the assertion of the lemma. The adversary A gets a public key pk as an initial input, and then it needs to mount a ciphertext-verification attack against the encryption scheme $\mathcal{E}$. To mount that attack, it uses the intruder algorithm I from the premise of the lemma.

On a high level, A has two interfaces: an "external interface" where it issues the verification queries and test query of the ciphertext verification attack, and an "internal interface" where it runs the intruder algorithm I, and interacts with it, playing the roles of both the server S and the user U. See a pictorial demonstration in Fig. 2.

After A receives its input pk, it picks names S, U for the server and user, respectively, picks a random password $\mathsf{spwd} \in \mathcal{D}$, and also picks a random index $t \in \{1, \ldots, \ell\}$. It records all of these values, and then then runs I, giving it the public key pk and the names S, U as input, and responding to I's actions as described below:

1. Whenever I "opens a new account" by giving the server the pair $(\mathsf{U}', \mathsf{spwd}')$, A records this pair. When I prompts the server for a new session, the adversary A picks a new session-id sid_i, records it, and sends it back to I.

2. When I sends a message sid'_j to the user, A responds with a triple $y'_j = \left\langle \mathsf{U}, \mathsf{sid}'_j, a_j \right\rangle$ where a_j is computed as follows.

   If $j < t$ then A sets $w_j = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}'_j)$ and $a_j = \textsc{Enc}_{pk}(w_j)$.

   If $j = t$ then A sets $w_t = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}'_t)$ and $\bar{w}_t = f(0;0)$, and asks its *test query* of the ciphertext-verification attack, by setting $x_1 = w_t, x_2 = \bar{w}_t$, asking for an encryption of one of them, and receiving a ciphertext c which is an encryption of either x_1 or x_2. Then, A sets $a_t = c$.

   If $j > t$ then A sets $\bar{w}_j = f(0\ ; 0)$ and $a_j = \textsc{Enc}_{pk}(\bar{w}_j)$.

3. When I sends to the server a message $y_i = \langle \mathsf{U}', \mathsf{sid}', a' \rangle$ with $\mathsf{U}' \neq \mathsf{U}$, the adversary A responds as follows: If it does not have a record for user U' or for session-id sid', then A ignores the message y_i (as the server would do in such a case). If $a' = a_t$ (i.e., a' is the same ciphertext that was returned in the test query), then A rejects the reply y_i.[8] Otherwise, A finds the password spwd' associated with U', computes $z_i = f(\mathsf{spwd}'\ ; \mathsf{U}', \mathsf{S}, \mathsf{sid}')$, and issues a verification query (of the ciphertext-verification attack) to find out whether $\textsc{Dec}_{\mathsf{sk}}(a') = z_i$. If it equals then A accepts the reply y_i, and if not it rejects it.

4. Finally, when I sends its (single) U-reply to the server (denoted $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$), A does the following: It verifies that $a \neq a_t$ (i.e., a is not the same ciphertext that was returned in the test query), computes $z = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$, and uses one last verification query, asking whether or not $\textsc{Dec}_{\mathsf{sk}}(a) = z$. If the answer is yes, A guesses that c is an encryption of x_1. Otherwise, A guesses that c is an encryption of x_2.

   If $a = a_t$ then A guesses that c is an encryption of x_1. This is completely arbitrary and does not effect the analysis below. To see why, recall that I is a restrained intruder, and so it never forwards an unmodified U-reply. It follows that if $a = a_t$ then the session-id sid in y is not the one that was used to compute w_t in Step 2. Hence, this U-reply will be rejected, so it cannot be a successful one.

   Once this is done, A halts.

[8] Notice that the server indeed would have rejected this reply. If a_t is an encryption of $f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$ then it will be rejected since $\mathsf{U}' \neq \mathsf{U}$, and if a_t is an encryption of $f(0;0)$ then it will be rejected since $f(0;0)$ has the wrong format.

Analysis of A. We first need to show that the queries made by A on its "external interface" adhere to the conditions set for the ciphertext-verification attack. Specifically, we need to show that A never asks a verification query with the same ciphertext it got for its test query, but this is guaranteed since A always checks this condition before making any verification query (See in Steps 3 and 4 above).

It is left to analyze the success probability of A. To that end, we define the following $\ell + 1$ "mental experiments", which are all variations on the probabilistic game between U, S and I. For $j = 0, 1, \ldots \ell$, the j'th experiment (which we denote E_j) consists of running the usual game up to the j'th message that I sends to U. Namely, for the j first messages $\mathsf{sid}'_1 \ldots \mathsf{sid}'_j$, the user U follows the protocol and sets $a_i = \textsc{Enc}_{\mathsf{pk}}(f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}'_i))$, and $y'_i = \langle \mathsf{U}, \mathsf{sid}'_i, a_i \rangle$. For all the other messages $\mathsf{sid}'_{j+1}, \mathsf{sid}'_{j+2}, \ldots$, the user instead sets $a_i = \textsc{Enc}_{\mathsf{pk}}(f(0;0))$, and $y'_i = \langle \mathsf{U}, \mathsf{sid}'_i, a_i \rangle$.

As in the usual game, we say that I fools the server in E_j if the server accepted any of the U-replies that I sends to it. (Recall that I is a restrained intruder, so the U-replies that it sends to the server are all different than what it receives from U.) We denote by p_j the probability that I fools the server in an $(\ell, 1)$-run in E_j. We now note that p_ℓ is exactly the probability that I fools the server in an $(\ell, 1)$-run of the original game, and by the premise of the lemma, we have $p_\ell \geq \delta + 1/|\mathcal{D}|$. On the other hand, we have

Proposition 6 $p_0 \leq 1/|\mathcal{D}|$

Before proving Proposition 6, we show how it is used to complete the proof of Lemma 5. Let t be the index that A picks at the beginning of its operations. Note that if the test query of A is answered with an encryption of $x = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}'_t)$, then the game that I sees is exactly the t'th mental experiment, so its success probability is p_t. On the other hand, if the test query of A is answered with an encryption of $f(0;0)$, then the game that I sees is exactly the $(t-1)$'th mental experiment, so its success probability is p_{t-1}. Thus we get

$$\begin{aligned}
&\Pr[A \text{ guesses "encryption of } x_1\text{"} \mid \textsc{Dec}_{\mathsf{pk}}(c) = x_1] \\
&\qquad - \Pr[A \text{ guesses "encryption of } x_1\text{"} \mid \textsc{Dec}_{\mathsf{pk}}(c) = x_2] \\
&= \sum_{j=1}^{\ell} \Pr[t = j] \cdot (p_j - p_{j-1}) \\
&= \frac{1}{\ell} \sum_{i=1}^{\ell} (p_j - p_{j-1}) \quad = \quad \frac{1}{\ell}(p_\ell - p_0) \quad \geq \quad \delta/\ell
\end{aligned}$$

■

Proof of Proposition 6. Let $y = \langle \mathsf{U}, \mathsf{sid}, a \rangle$ be the first U-reply that I sends in the experiment E_0, and recall that this reply is successful if and only if $\textsc{Dec}_{\mathsf{sk}}(a) = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$.

Since f is one-to-one on its components, it follows that for every $\mathsf{spwd}' \neq \mathsf{spwd}$, $f(\mathsf{spwd}'\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}) \neq f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid})$. Therefore, for every key pair $(\mathsf{pk}, \mathsf{sk})$, every ciphertext a and every session-id sid, it holds that

$$\Pr_{\mathsf{spwd} \in \mathcal{D}} \left[\textsc{Dec}_{\mathsf{sk}}(a) = f(\mathsf{spwd}\ ; \mathsf{U}, \mathsf{S}, \mathsf{sid}) \right] \leq \frac{1}{|\mathcal{D}|}$$

Since the messages that the adversary sees in experiment E_0 up until it sends y are all independent of the password spwd, then y itself must also be independent of spwd, and therefore the probability that it is successful (taken over the choice of the password spwd) is at most $1/|\mathcal{D}|$. ■

5.4 Mutual authentication and key exchange

In this paper we do not formalize or prove the security of the extensions to our basic one-way authentication protocol as presented in Section 3.4 for providing mutual authentication and key exchange. Formalizing authentication and key-exchange, in general, is beyond the scope of this paper. However, we point out that a suitable framework for the analysis of such protocols has been recently developed by Bellare, Canetti and Krawczyk [1]. That paper analyzes key-exchange protocols which are very similar to the ones presented here. The significant difference is that they do not analyze password-based protocols. Fortunately, it can be shown that our security formalization and proof of the Encrypted Challenge-Response protocol for one-way authentication provides a *password-based authenticator* in the language of [1], and so the analysis methodology from that work applies to our protocols too.

(Authenticators are protocols that guarantee "secure delivery" of messages in the presence of active attackers. It is shown in [1] that when an authenticator is applied on top of a key-exchange protocol which is secure against eavesdroppers-only (such as the Diffie-Hellman protocol) then the whole key exchange is secure against active attackers too. This is the way that we build the key-exchange protocols presented in this paper.)

In this regard, it is worth noting that while [1] only present protocols that use the same form of authentication in the two directions of the protocol, here we use a password-based authentication from user to server, and an encryption-based authentication from server to user. Yet, this framework is extendible to this "a-symmetry" in the authentication methods.

6 Do secure password protocols require public-key tools?

It is interesting to note that although in a password setting, the user and server have a *shared secret* (albeit, a weak one), all the strong password mechanisms proposed in the literature, including ours, employ public key techniques (e.g., [12, 4]). An interesting question, therefore, is whether password authentication mechanisms resistant to off-line guessing can be built based on symmetric key techniques only, e.g. solely based on the existence of a secure block cipher. Below we provide very strong evidence against such possibility, by proving that a public key primitive such as key exchange is inherently necessary in the construction of password protocols which resist password-guessing attacks.

What we specifically prove is that given a secure password protocol resistant to guessing attacks one can use it (without further cryptographic functions) to implement a key-exchange protocol. The latter is a protocol that allows parties that do not share any initial secret to exchange a fresh secret via a public and authenticated conversation, as in the case of the Diffie-Hellman protocol. It follows that secure password authentication protocol *must* use whatever machinery is required for building secure cryptographic key exchange protocols. Or, equivalently, that these cryptographic techniques are essential for constructing strong password authentication. Combined with a result by Impagliazzo and Rudich [14], our result implies that constructing secure password protocols using only symmetric key techniques such as block ciphers and hash functions is extremely unlikely (barring some major breakthrough in cryptography and complexity theory, see discussion below).

6.1 Secure password authentication implies key-exchange

For simplicity, in the presentation below we focus on protocols for exchanging a single bit. Clearly, such protocols can be used to implement exchange of longer keys (simply by repeating the same protocol many times). Below we formally define such protocols and their security against passive

eavesdroppers; we then prove that such key-exchange protocols can be implemented using any secure password protocol.

Definition 3 *A two-party protocol (A, B) is a key exchange protocol for one bit, if at the end of the protocol, both A, B output the same bit.*[9] *(The input to both A, B is a security parameter k.) Let $\epsilon(\cdot)$ be a function, and let (A, B) be a key exchange protocol for one bit. We say that (A, B) is secure up to ϵ if no feasible eavesdropper E can guess the bit that A, B output with probability better than $1/2 + \epsilon(k)$, where k is the security parameter of (A, B).*

Theorem 2 *Any protocol that ensures one-way password authentication up to $\epsilon(k, \ell, m)$ can be transformed into a key exchange protocol for one bit, which is secure up to $\epsilon'(k) = \epsilon(k, 1, 1)$.*

Proof overview We start with a somewhat informal overview of the proof. Assume that we have a password protocol that resists off-line guessing attacks. In particular, this means for any dictionary $\mathcal{D}$ that is used by the parties, a passive eavesdropping adversary (which only listen on the lines) cannot guess the user's password with probability significantly larger than $1/|\mathcal{D}|$. In particular, if the parties are using a dictionary of size 2, then the adversary cannot guess the user's password with probability significantly larger than a half.

In order to exchange a secret bit, the two parties use the password protocol with a dictionary of size 2. To exchange a bit, each of the parties chooses at random a password from the dictionary, and then they execute the password protocol with one party playing the server and the other playing the user. It follows from the security of the password protocol that this execution succeeds if and only if they both choose the same password. If the execution succeeds, then their secret bit is set to '0' in the case that the password that they chose was the first password in the dictionary, or to '1' otherwise. If the execution fails (i.e., they chose different passwords) then they choose new passwords and try again. Since the protocol resists password-guessing attacks, then an eavesdropping adversary cannot guess the password that was used in a successful execution, and so the exchanged bit is indeed secret. After expected 2 trials, the parties will be able to exchange the secret bit. A more formal description follows.

Formal proof Let $\epsilon(k, \ell, m)$ be a real function and let (U, S) be any (syntactically correct) password protocol that ensures one-way authentication up to ϵ. The key exchange protocol (A, B) works as follows: The parties use the protocol (U, S) with A playing the role of U and B playing the role of S. On security parameter k, A and B first execute the setup phase of the protocol (U, S) with a dictionary $\mathcal{D}$ of size two and the same security parameter, *but without exchanging the password.* Instead, each of A, B picks a random password from $\mathcal{D}$ and behaves as if this is the password that was exchanged. Then, A, B execute an authentication session, and at the end of this session B (who plays the role of the server S) sends to A a message, telling it if the session was successful or not. If the session does not succeed, then A, B repeat the whole process (including the setup phase) again, until the session is successful. Once the session succeeds, each party outputs a '0' if it chose the first password in the dictionary for this session, or '1' if it was the second password. We note that the protocol (A, B) ends after expected two executions of the protocol (U, S).

We start by proving that when the parties choose different passwords, then the authentication session fails (except with probability of at most $2\epsilon(k, 1, 1)$). To see that, consider an intruder I that simply guesses a password from $\mathcal{D}$ at random and tries an authentication session with the server

[9] In fact, we may allow the protocol to have a small probability of error, when the bits output by the two parties are different. Of course, to be of any use, the probability that the parties agree on the bit must be strictly larger than the probability that an eavesdropper guesses this bit.

using that password. On one hand, since (U, S) is syntactically correct, we know that this intruder wins if it chooses the correct password (which happens with probability 1/2). On the other hand, since (U, S) is secure then we also know that the intruder cannot win with probability more than $1/2 + \epsilon$. Hence we have

$$\begin{aligned} 1/2 + \epsilon &\geq \Pr[I \text{ wins}] \\ &= 1/2 \cdot \Pr[I \text{ wins} \mid \text{password is correct}] + 1/2 \cdot \Pr[I \text{ wins} \mid \text{password is incorrect}] \\ &= 1/2 + 1/2 \cdot \Pr[I \text{ wins} \mid \text{password is incorrect}] \end{aligned}$$

and therefore $\Pr[I \text{ wins} \mid \text{password is incorrect}] \leq 2\epsilon$. It follows that the probability that the bit output by the two parties (A, B) is not the same, is at most $2\epsilon/(1 + 2\epsilon) < 2\epsilon$.

We now show that (A, B) is indeed secure up to $\epsilon'(k) = \epsilon(k, 1, 1)$. Formally, we reduce the security of (A, B) to that of (U, S), by showing that if there is an eavesdropper E with advantage of more than ϵ' against (A, B), then there also exists an intruder I with advantage of more than ϵ against (U, S). In this reduction we let the intruder I use the algorithm E as a subroutine in its attack against the protocol (U, S).

The intruder I is given the security parameter k, a dictionary $\mathcal{D}$ of two passwords, and the names (and possibly public keys) of the user and server. It prompts the parties to start an authentication session and records that session, passing the messages back and forth without modifying anything. Denote the transcript of that session by τ. Then, I repeatedly picks pairs of passwords from $\mathcal{D}$ $\langle a_i, b_i \rangle$. As long as $a_i \neq b_i$, the intruder I generates a session of the protocol (A, B) in which A picks the password a_i and B picks the password b_i, and records the transcript of this session τ_i. In the first instance where $a_i = b_i$, the intruder I invokes the eavesdropper E, giving it as input the sequence of transcripts $\langle \tau_1, \ldots, \tau_{i-1}, \tau \rangle$. The eavesdropper then guesses a bit σ, and then I attempts an on-line impersonation attack using the first passwords from $\mathcal{D}$ if $\sigma = 0$, and the second if $\sigma = 1$.

It is clear that the view of the eavesdropper E in this execution is identical to the view when it interacts with the real protocol (A, B), and therefore, by our assumption, E guesses the right bit (and I uses the right password) with probability of more than $1/2 + \epsilon'(k)$. Hence, we have

$$\Pr[I \text{ wins in a (1,1)-game}] > 1/2 + \epsilon'(k) = 1/2 + \epsilon(k, 1, 1)$$

which contradicts the security assumption on the protocol (U, S). ■

Remark: In the above proof we use the fact that strong password-based authentication protocols need to resist off-line guessing attacks for any size dictionary $\mathcal{D}$. In particular, even for $|\mathcal{D}| = 2$. That is, given two candidate passwords only it shouldn't be feasible for the attacker to guess the right password with probability significantly better than 1/2. This is guaranteed for our protocols by Theorem 1, and is also commonly conjectured (but not proved) for the EKE protocols of [4] and some of its variants. The assumption in these protocols is that given a single candidate password, there is no efficient way to check off-line whether this password is the one used by the user; this implies the above condition on dictionaries of size 2. We stress that if a given protocol is guaranteed to resist guessing attacks with dictionaries of at least t entries, for some number t, then the proof still works but the complexity of exchanging a bit is t expected trials instead of 2.

6.2 Discussion

Since key-exchange protocols are themselves "public key tools", then Theorem 2 can be viewed as suggesting a positive answer to the question in the title of this section. Certainly, it means

that whatever tools are needed for key exchange, are also needed for secure password protocols. But is it possible that both password protocols and key exchange protocols can be implemented using only primitives such as symmetric ciphers and hash functions? The answer to this question is "probably not", and it follows from the work of Impagliazzo and Rudich [14]: In that work, they devised a mathematical model in which "only secret-key tools are available", and analyzed which protocols can or cannot be obtained in this model. There are two features that distinguish the Impagliazzo-Rudich model from the standard model of computation.

Powerful adversary. To eliminate all "public key tools" from the model, the adversary is given the power to "solve any NP computational problem". Specifically, the adversary is given an oracle access to an NP-complete problem. (In particular, this means that given a "public key", the adversary can find the corresponding "secret key" in polynomial time.)

Access to a random oracle. To enable "non public-key tools" in the presence of this powerful adversary, all the parties in this model are given access to a random function f, mapping arbitrary-length strings into strings of length k (where k is a security parameter). Given such function, it is easy to implement primitives such as collision-intractable hashing, symmetric encryption, etc. For example, collision-intractable hashing is achieved simply by applying the function f to an arbitrary-length string. Also, when two parties share a secret key α, they can use f to generate a random one-time pad by computing the i'th block of the pad as $f(\alpha, i)$, and then use this pad for encryption.

The main result in [14] is that secure key exchange protocols are not possible in this model. Namely

Theorem 3 ([14]) *There is no key-exchange protocol in the Impagliazzo-Rudich model which is secure up to ϵ, for any $\epsilon < 1/2$.*

Using Theorem 2, we therefore have

Corollary 7 *There is no password protocol in the Impagliazzo-Rudich model that ensures one way authentication up to ϵ, for any $\epsilon < 1/2$.*

Implications in the standard model. The above corollary suggests that "non public-key tools" are not sufficient for secure password protocols, as there is a model where the former exist and the latter do not. But in fact, this result has also implications for the standard model of computations. Specifically, it means that exhibiting a password protocol which can be proven secure based only on the security of "generic secret-key tools" (such as symmetric encryption, MAC or collision-intractable hashing) is at least as hard as proving that $P \neq NP$. To see that, assume that we have such a password protocol (U, S). Since the security of this protocol is proven based only on the security of the "generic secret-key tools" that it uses, then the protocol remains secure also in the model where all parties have access to a random function, and these "secret-key tools" are implemented using that function. However, the Impagliazzo-Rudich result says that in this model, being able to solve any NP-complete problem is sufficient to break the protocol. Since we assume that the scheme is proven secure, then we know that no efficient adversary can break it, which implies that no efficient adversary can solve NP-complete problems. Hence this would prove that $P \neq NP$.

References

[1] M. Bellare, R. Canetti and H. Krawczyk, "A Modular Approach to the Design and Analysis of Authentication and Key Exchange Protocols", *Proceedings of the Thirtieth ACM Symposium on the Theory of Computation (STOC)*, 1998, pp. 419–428.

[2] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway, "Relations Among Notions of Security for Public-Key Encryption Schemes", *Advances in Cryptology - CRYPTO'98 Proceedings*, Lecture Notes in Computer Science Vol. 1462, H. Krawczyk, ed., Springer-Verlag, 1998, pp. 26–45.

[3] M. Bellare, and P. Rogaway, "Optimal Asymmetric Encryption – How to encrypt with RSA", *Advances in Cryptology - EUROCRYPT'94 Proceedings*, Lecture Notes in Computer Science Vol. 950, A. De Santis ed, Springer-Verlag, 1995.

[4] S. M. Bellovin and M. Merritt, "Encrypted Key Exchange: Password- Based Protocols Secure Against Dictionary Attacks", *Proceedings of the IEEE. Symposium on Research in Security and Privacy*, Oakland, May 1992.

[5] S. M. Bellovin and M. Merritt, "Augmented Encrypted Key Exchange: a Password-Based Protocol Secure Against Dictionary Attacks and Password File Compromise", *Proceedings of the First ACM Conference on Computer and Communications Security*, 1993, pp. 244-250.

[6] R. Cramer and V. Shoup, "A Practical Public Key Cryptosystem Provably Secure against Adaptive Chosen Ciphertext Attack", *Advances in Cryptology - CRYPTO'98 Proceedings*, Lecture Notes in Computer Science Vol. 1462, H. Krawczyk, ed., Springer-Verlag, 1998, pp. 13–25.

[7] W. Diffie, P. C. Van-Oorschot, and M. J. Weiner. Authentication and authenticated key exchanges. *Designs, Codes and Cryptography*, 2:107–125, 1992.

[8] D. Dolev, C. Dwork, and M. Naor. "Non-malleable cryptography". *Proceedings of the Twenty Third Annual ACM Symposium on Theory of Computing*, pages 542-552, 1991.

[9] I. Goldberg, H. Finney and R. Levien. "Visual Fingerprints" and "Snowflakes". `http://www.cs.berkeley.edu/~iang/visprint.html`

[10] O. Goldreich, S. Goldwasser and S. Micali. "How to Construct Random Functions", *Journal of the ACM*, Vol. 33, no. 4, 1986, pp. 792-807

[11] S. Goldwasser, and S. Micali. "Probabilistic Encryption", *Journal of Computer and System Sciences*, Vol. 28, 1984, pp. 270-299.

[12] L. Gong, M. Lomas, R. Needham, and J. Saltzer, "Protecting Poorly Chosen Secrets from Guessing Attacks", *I.E.E.E. Journal on Selected Areas in Communications*, Vol. 11, No. 5, June 1993, pp. 648-656.

[13] N. Haller, "The S/KEY One-Time Password System", RFC 1760, Feb. 1995.

[14] R. Impagliazzo and S. Rudich. "Limits on the provable consequences of one-way permutations". In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, 1989, pages 44-61.

[15] D. Jablon, "Strong Password-Only Authenticated Key Exchange". *Computer Communication Review*, ACM SIGCOMM, vol. 26, no. 5, pp. 5-26, October 1996.

[16] C. Kaufman, R. Perlman, and M. Speciner, "Network Security," Prentice Hall, 1997.

[17] H. Krawczyk, "SKEME: A Versatile Secure Key Exchange Mechanism for Internet,", *Proceedings of the 1996 Internet Society Symposium on Network and Distributed System Security,* Feb. 1996, pp. 114-127.

[18] L. Lamport, "Password authentication with insecure communication," *Comm. of the ACM,* Vol. 24 Number 11, Nov 1981, pp. 770-772.

[19] S. Lucks, "Open Key Exchange: How to Defeat Dictionary Attacks Without Encrypting Public Keys", *The Security Protocol Workshop '97*, Ecole Normale Superieure, April 7-9, 1997.

[20] A. Menezes, P. Van Oorschot and S. Vanstone, "Handbook of Applied Cryptography," CRC Press, 1997.

[21] S. Patel, "Number Theoretic Attacks On Secure Password Schemes" *IEEE Symposium on Security and Privacy*, Oakland, California, May 5-7, 1997.

[22] C. Rackoff and D. Simon, "Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack", *Advances in Cryptology - CRYPTO'91 Proceedings*, Lecture Notes in Computer Science Vol. 576, J. Feigenbaum ed, Springer-Verlag, 1991.

[23] M. Steiner, G. Tsudik, and M. Waidner, "Refinement and Extension of Encrypted Key Exchange", *Operating Systems Review*, vol. 29, Iss. 3, pp. 22-30 (July 1995).

[24] T. Wu, The Secure Remote Password Protocol, in *Proceedings of the 1998 Internet Society Network and Distributed System Security Symposium*, San Diego, CA, Mar 1998, pp. 97-111.