

Security Issues in the Diffie-Hellman Key Agreement Protocol

Jean-François Raymond and Anton Stiglic

Zero-Knowledge Systems Inc.
{jfr, anton}@zeroknowledge.com

Abstract Diffie-Hellman key agreement protocol [20] implementations have been plagued by serious security flaws. The attacks can be very subtle and, more often than not, have not been taken into account by protocol designers. In this summary we discuss both theoretical attacks against the Diffie-Hellman key agreement protocol and attacks based on implementation details. It is hoped that computer security practitioners will obtain enough information to build and design secure and efficient versions of this classic key agreement protocol.

1 Introduction

In their landmark 1976 paper “New Directions in Cryptography”[20], Diffie and Hellman present a secure key agreement protocol that can be carried out over public communication channels. Their protocol is still widely used to this day.

Even though the protocol seems quite simple, it *can* be vulnerable to certain attacks. As with many cryptographic protocols, the Diffie-Hellman key agreement protocol (DH protocol) has subtle problems that cryptographers have taken many years to discover. This vulnerability is compounded by the fact that programmers often do not have a proper understanding of the security issues. In fact, bad implementations of cryptographic protocols are, unfortunately, common [2].

In this work, we attempt to give a comprehensive listing of attacks on the DH protocol. This listing will, in turn, allow us to motivate protocol design decisions. Note that throughout this presentation emphasis is placed on practice. After reading this paper, one might not have an extremely detailed understanding of previous work and theoretical problems, but should have a very good idea about how to securely implement the DH protocol in different settings.

1.1 Related Work

As mentioned previously, flaws in cryptographic protocols are not uncommon. Hence, the problem has received some attention from the cryptography community; here are the most important approaches that have been proposed:

1. The use of verification logics such as BAN [15] to prove protocol properties.
2. Very high level programming languages in which security properties can be proved mechanically (i.e. by computers) [1].

3. Complete proofs of security [7, 6].
4. The use of robustness principles, i.e. rules of thumb, protocol design principles [3].

The biggest problem with the first approach is that encryption primitives are dissociated from the verification logics which implies that they do not provide complete proofs of security [6]. As an example of this problem one just needs to look at the problem of encryption and signature ordering: most verification logics do not complain when messages are encrypted before being signed which possibly results in a security vulnerability [3].

The second approach seems promising however the best known proof mechanization techniques are not efficient enough and only a few cryptographic primitives have been included in the model.

The third suggestion is the most powerful. The main problem is that the proofs are somewhat involved and proving the correctness of complex protocols seems quite difficult. Note also that it is not entirely obvious that the claims that are proved are adequate.

The robustness principles are useful in that they can help in preventing common errors. However, it is hard to exhaustively list all important robustness principles, and so using these principles does not give us peace of mind as there are no security guarantees. Furthermore, the protocol designer must be comfortable and competent in verifying security properties. For example Principle 3 of [3], which states

Be careful when signing or decrypting data that you never let yourself be used as an oracle

might not be understood by individuals that do not have a background in cryptography.

The most important problem with all of the above approaches is that low level implementation issues are not spelled out. Hence, unless one has a solid grasp of all of the technical details, it is very easy to make low-level implementation errors and difficult to debug code. Also notice that none of these approaches deal with problems specific to the cryptographic primitives used.

Many standards have been developed for the DH protocol (see Appendix A), unfortunately none describe the issues and attacks in detail. More importantly, none motivate the design decisions. In [8], work has been done to characterize the security of the DH protocols introduced in various standards but not much is discussed about known attacks and implementation details are ignored. Our work can be considered as complementary to that of [8], describing and studying the most important theoretical and practical issues and considering implementation details.

1.2 Overview

Section 2 presents a mathematical background of the basics needed to understand the DH protocol and the types of attacks it is vulnerable to. In section 3 we give attacks which are based on mathematical tricks. Authentication is discussed in section 4. In section 5 we discuss attacks on DH that exploit implementation details. In section 6, we expose some subtleties that appear when using the DH shared secret to obtain a key which can be used in other cryptographic operations. The information acquired in

sections 3, 4, 5 and 6 is used to present implementation guidelines in section 7. The conclusion can be found in section 8.

2 The Diffie-Hellman Key Agreement Protocol

The DH key agreement protocol allows two users, referred to as Alice (\mathcal{A}) and Bob (\mathcal{B}), to obtain a shared secret key over a public communication channel. An attacker, *eavesdropping* at the messages sent by both Alice and Bob will not be able to determine what the shared secret key is. This is an extremely useful primitive because the shared secret can be used to generate a secret session key that can be used with symmetric crypto-systems¹ (e.g. DES) or message authentication codes (MAC). We now give some basic notions from mathematics that are needed to understand the protocol.

2.1 Mathematical Background

The computations required in the DH protocol are carried out in a group.

Groups A *group* (G, \star) consists of a set G and a binary operation \star that takes elements of G as inputs. \star has the following properties:

1. (associativity) $a \star (b \star c) = (a \star b) \star c$, for all $a, b, c \in G$.
2. (identity element) There is an element $1 \in G$, called the identity, that has the property that $1 \star a = a \star 1 = a$, for all $a \in G$.
3. (inverse element) For each $a \in G$, there exists a value denoted by a^{-1} such that $a \star a^{-1} = a^{-1} \star a = 1$.

An *Abelian group* is a group having the following additional property:

4. (commutativity) $a \star b = b \star a$ for all $a, b \in G$.

For *finite groups* (G finite), the order of a group is defined as the cardinality (size) of G . The *order* of an element a of a finite group G is defined to be the smallest value t such that $a^t := \underbrace{a \star a \star \dots \star a}_t = 1$.

Cyclic Groups A cyclic group is a group that has the property that there exists an element g such that all elements in G can be expressed as g^i (for different i s). If g generates all elements of the group (G, \star) , g is a generator and we say it generates (G, \star) . Note that the order of a generator g equals the order of the group it generates.

Subgroups We say that G' is a subgroup of G if (G', \star) forms a group and $(G' \subseteq G)$. If G is a finite group, then the order of a subgroup G' will always divide the order of G (*Lagrange's theorem*, see for example [28]).

¹ for an introduction to cryptography see [50].

Examples of Groups Groups typically used for DH protocols are the set \mathbb{Z}_p^* with multiplication modulo p where p is prime, the multiplicative group of the field² \mathbb{F}_{2^m} and the additive group formed by a collection of points defined by an elliptic curve over a finite field. These groups all have the property that exponentiating is computationally inexpensive and that computing discrete logs is/seems hard (i.e. computationally intractable).

In the remainder of this work, we will take the group to be the set $\mathbb{Z}_p^* = \{1, 2, \dots, p-1\}$ with multiplication modulo p (p prime) and all operations will be taken over this group (g^y will stand for $g^y \bmod p$, for example). Small variations on many of the attacks of the following sections can be easily mounted on DH implementations using other groups. Note that we will abuse the notation a bit by using \mathbb{Z}_p^* when referring to the group composed of the set \mathbb{Z}_p^* with multiplication modulo p .

2.2 The Core DH Protocol

Alice (\mathcal{A}) and Bob (\mathcal{B}) first agree on a large prime number p and an element g ($2 \leq g \leq p-2$) that generates a (cyclic) *subgroup* of large order. These values are usually determined a-priori, and are used for many protocol runs (e.g. they could be public parameters that everybody uses). The rest of the protocol goes as follows:

1. \mathcal{A} chooses a number, x , at random from the set $\{1, \dots, p-2\}$. And \mathcal{B} chooses y randomly from the same set.
2. \mathcal{A} sends g^x to \mathcal{B} and \mathcal{B} sends g^y to \mathcal{A} .
3. The shared secret key is $K = g^{xy}$. \mathcal{A} , knowing x and g^y , can easily calculate $(g^y)^x = g^{xy}$. \mathcal{B} can determine the secret key in a similar manner by computing $(g^x)^y$.

x and y are referred to as the private keys, g^x and g^y are referred to as the public keys and g^{xy} is called the shared (DH) secret key. When the secret keys are used only once we call this an ephemeral DH secret key agreement. We assume that an eavesdropper having access to the public values cannot calculate the shared secret key, this is called the Diffie-Hellman assumption. The Diffie-Hellman assumption is somewhat related to the Discrete Log assumption which states that given a generator g of \mathbb{Z}_p^* and an element β of \mathbb{Z}_p^* , it is infeasible to compute x such that $g^x \equiv \beta$ in \mathbb{Z}_p^* . The relation stems from the fact that if we can compute discrete logs efficiently, we can efficiently compute g^{xy} given only g , g^x and g^y and thus invalidate the Diffie-Hellman assumption, but the converse is not known to be true.

2.3 Half-Certified Diffie-Hellman (or Elgamal Key agreement protocol)

This is a very important and useful variant on the Diffie-Hellman protocol discussed above. First introduced in [22], the protocol is almost exactly the same as the basic one except that a user (Bob) publishes his public key (g^y). The public key (g^y) remains constant for large periods of time and is used by everyone wishing to set up a shared

² see for example [28] for a definition of *field* as well as an overall introduction to algebra.

secret key with Bob. Note that the public key should be authenticated in some way (e.g. by Bob's signature). This mechanism is especially useful for secure anonymous client connections³.

2.4 Attacks

Attacks against the *DH protocol* come in a few flavors:

- **Denial of service Attacks:** Here, the attacker will try to stop Alice and Bob from successfully carrying out the protocol. The attacker can accomplish this in many ways, for example by deleting the messages that Alice and Bob send to each other, or by overwhelming the parties with unnecessary computation or communication.
- **Outsider Attacks:** The attacker tries to disrupt the protocol (by for example adding, removing, replaying messages) so that he gets some interesting knowledge (i.e. information he could not have gotten by just looking at the public values).
- **Insider Attacks:** It is possible that one of the participants in a DH protocol creates a breakable protocol run on purpose in order to try to gain knowledge about the secret key of his peer. This is an important attack if one of the participants holds a static secret key that is used in many key agreement protocol runs. Note that malicious software could be very successful in mounting this attack.

The plausibility of these attacks depends on what assumptions we make about the adversary. For example, if the adversary can remove and replace any message from the public communication channel, the denial of service attack is impossible to prevent. Fortunately, it seems that complete breaks (outsider attacks in which the attacker obtains the shared secret key) and insider attacks can be prevented in many settings⁴.

2.5 Man in the Middle Attacks

An active attacker (Oscar), capable of removing and adding messages, can easily break the core DH protocol presented above. By intercepting g^x and g^y and replacing them with $g^{x'}$ and $g^{y'}$ respectively, Oscar (\mathcal{O}) can fool Alice and Bob into thinking that they share a secret key. In fact, Alice will think that the secret key is $g^{xy'}$ and Bob will believe that it is $g^{x'y}$. This is a *man in the middle* attack [47].

As an example of what can be done with such an attack, consider the case where Alice and Bob use a shared secret key obtained in a DH protocol for symmetric encryption. Suppose Alice sends a message m to Bob and that $ENC_K(x)$ represents the symmetric encryption (e.g. DES) of x using the secret key K .

1. \mathcal{A} sends $ENC_{g^{xy'}}(m)$.
2. \mathcal{O} intercepts $ENC_{g^{xy'}}(m)$ and decrypts it (which he can do since he knows $g^{xy'}$).
3. \mathcal{O} replaces this message with $ENC_{g^{x'y}}(m')$ which he sends to \mathcal{B} . Note that m' can be set to any message.

The encryption scheme is thus clearly compromised as message privacy is violated. In the next section, we study attacks that can be mounted by a less powerful adversary.

³ this scheme is sometimes called Half Static Diffie-Hellman (because one secret, y , is static).

⁴ in practice it is much easier to insert packets than it is to delete them. In any case, we consider all attacks in order to derive a DH protocol that is secure in all practical settings.

3 Attacks Based on Number Theory

The previous man in the middle attack, although it completely breaks the protocol, requires Oscar to be very powerful. For example, if the secret keys are used in conjunction with MACs, Oscar needs to intercept and modify each authenticated message in order to prevent Alice and Bob from detecting that their keys are not identical. In some of the following subsections, Alice and Bob have the same secret key (which Oscar knows). Thus, Oscar only needs to be active during the DH protocol, afterwards he can break the protocols using the shared secret key whenever he wants.

3.1 Degenerate Message Attacks

There are degenerate cases in which the protocol does not work (i.e. it can be broken). For example when g^x or g^y equals one, the shared secret key becomes 1. Since the communication channel is public anybody can detect this anomaly. Fortunately, this situation is impossible in a properly carried out protocol run because both x and y are chosen from $\{1, \dots, p-2\}$ ⁵. However, an insider attack is possible and so DH protocol participants should make sure that their key agreement peer does not send $g^z = 1$.

Simple Exponents If one of x and y can be easily determined, the protocol can be broken. For example, if x equals 1 then $g^x = g$ which any observant attacker will be able to detect. It is very hard to determine where to draw the line here, that is, determining for which values of g^i , i is hard to determine, since this depends entirely on the strategy of the attacker. Any set of i values could be vulnerable, depending on which values of g^i are precomputed, where the search starts, and how it proceeds. In any case, it seems very reasonable to insist that x and y not equal 1.

Simple Substitution Attacks The following attack is very interesting, as it is extremely easy to mount and normally would not come up in theoretical proofs of security. The attacker can force the secret key to be an “impossible” value. If the DH protocol would only be executed by sentient beings this would not be interesting as the anomalies would be easily detected. However in practice DH protocols are carried out by computers and careless implementations might not spot the following attack.

1. \mathcal{O} intercepts g^x and g^y and replaces them with 1.
2. Both \mathcal{A} and \mathcal{B} compute the same shared secret key which equals one.

If the computer program does not realize that g^x , g^y and g^{xy} cannot equal 1, the protocol is vulnerable. Note that the same argument holds for values of the form $g^{\alpha \cdot (p-1) \cdot x}$ or $g^{\alpha \cdot (p-1) \cdot y}$, where $\alpha \geq 1$, because a computer might not realize that these values have not been computed modulo p even though that they are large. (They equal 1 modulo p). So it is safe practice to always verify that g^x and g^y are positive integers smaller than $p-1$ and greater than 1.

The following attacks delve a bit deeper into computational number theory.

⁵ if g is a generator of \mathbb{Z}_p^* , $g^z = 1 \bmod p$ iff $z = 0 \bmod p-1$.

3.2 Generators of Arbitrary Order and the Pohlig-Hellman Algorithm

The Pohlig-Hellman algorithm [45] allows one to efficiently compute the discrete log of g^x if the prime factorization of g 's order consists of small primes. Precisely, given that the order of a group has the following prime factorization, $p_1^{e_1} p_2^{e_2} \cdot \dots \cdot p_r^{e_r}$, the Pohlig-Hellman algorithm's computational complexity is $O(\sum_{i=1}^r e_i (\lg(n) + \sqrt{p_i}))$. A secure DH implementation must make this algorithm impractical. A simple solution is to choose a prime p such that $p - 1$ contains large factors. Safe primes, primes of the form $p = Rq + 1$ (where R is some small positive value and q is a large prime), and Lim-Lee primes [36] which have the form $p = 2q_1 \cdot \dots \cdot q_n + 1$ (where the q_i s are all large primes) satisfy this property. (Remember that the order of any subgroup will divide $p - 1$, i.e. the order of \mathbb{Z}_p^* .)

3.3 Attacks Based on Composite Order Subgroups

The attacker can exploit subgroups that do not have large prime order [51]. This is best illustrated by an example. Suppose Alice and Bob choose a prime $p = 2q + 1$, where q is prime, and a generator g of order $p - 1 = 2q$. Oscar can intercept the messages g^x and g^y and exponentiate them by q . (He will replace g^x by g^{xq} and g^y by g^{yq} .) The secret key will be g^{xyq} which allows Oscar to find this value by exhaustive search. This is done by noting that the order of $g^q = g^{\frac{p-1}{2}}$ is⁶ 2 which implies that the secret key can only take one of two values! Hence, Oscar can use a brute force search (only two elements to try) in order to determine what the shared secret key is; for example, when Alice and Bob use it for symmetric encryption.

More generally, this attack can be easily mounted on primes of the form $p = Rq + 1$ (R small), the only difference being that there are R possible values to try in the exhaustive search.

The lesson to be learned from this attack is that we should choose a g that generates a large prime order *subgroup* or at the very least make sure that composite order subgroups are not vulnerable (e.g. the order's prime number factorization contains only large primes). Note that an attack of this type is part of the motivation for using DSA instead of Elgamal signatures. In essence DSA is an immunized version of Elgamal [4].

Notice that an insider attack can be mounted using this trick. Alice simply chooses x to equal q . In this case, even authentication mechanisms cannot protect Bob.

3.4 Pollard Lambda Algorithm

The Pollard Lambda method [46] enables one to compute z given g^z , when z is known to be in a certain interval $[b, b + w]$ in time $O(w^{1/2})$. This is an extremely relevant attack to consider when we want to limit the exponent range to improve efficiency. For example, when $x, y < 2^N \ll p$ the attacker can compute x and y (given g^x and g^y) in $O(\sqrt{2^N}) = O(2^{N/2})$. Hence, if we want the attacker to execute at least (2^N) operations⁷, x and y need to have been chosen uniformly at random in an interval of

⁶ the subgroup generated by $g^{\frac{p-1}{2}}$ is $\{g^{\frac{p-1}{2}} = p - 1, (g^{\frac{p-1}{2}})^2 = 1\}$.

⁷ Assuming that the Pollard Lambda technique is the best method in this situation.

size 2^{2N} . (Choosing x and b to be uniformly random $2N$ bit integers satisfies this last requirement.) We note that this attack has not been improved in a long time and many cryptographers feel that it is improbable that the state of the art for this kind of attack will change (this is a useful observation when choosing key sizes). Also remark that this attack can be mounted on subgroups of small order.

3.5 The Number Field Sieve Algorithm

It is obviously important to choose a group (i.e. p) large enough so that the best known algorithms for computing discrete logs are intractable. The state of the art index calculus based methods for computing discrete logs (number field sieves) have been steadily improving⁸ over the years and so it is harder to gauge how large p should be for long term security. In [41], Odlyzko proposes using a p of at least 1024 bits for moderate security and at least 2048 bits for anything that should remain secure for a decade. Note however that these values are controversial (see section 7.1).

3.6 Attacks on Prime Order Subgroups

In [36], an attack on prime order subgroups is presented (a slight extension of the ideas of [51]). The attack can be mounted if the protocol does not satisfy the sixth robustness principle of [3] which states:

Do not assume that a message you receive has a particular form unless you can check this.

The idea is that if we can get a participant with a secret key x to use an arbitrary group element instead of g^y when computing the DH shared secret key then we may be able to obtain some knowledge about x . If the attacker can obtain γ^x , for some generator γ whose order's prime factorization contains only small primes, then he can use the Pohlig-Hellman algorithm of subsection 3.2 to obtain x modulo the order of γ .

To obtain the actual value of the secret key (modulo p), a slight variation on the Pollard lambda method [51] might be feasible against a participant using a static private key.

The difficulty in this attack resides in the need to obtain γ^x . Lim and Lee [36] give a weaker version of the previous attack that enables the attacker to obtain the value of x modulo the order of γ in time linear in the order of γ .

This attack can be easily foiled if the participants check that the value they receive (i.e. usually g^z) has order q . This can be done by verifying that exponentiating the value by q yields 1. If p is of the form $p = 2q + 1$, with q prime, the best one can hope for is to determine the parity of the secret key.

If we have Certificate Authorities (CAs) certify the DH public keys, they must be wary of this attack. It is usually sufficient for the CA to verify that the user knows the secret associated with the public key when the participants use static public keys. This is usually done by having the user sign some message with the secret key. Unfortunately,

⁸ As opposed to the Pollard Lambda type algorithms for which there has not been substantial progress for about twenty five years [41].

a variation on the previous attack allows for an insider attack where a user can fool the CA when specific signature schemes are used (e.g. Schnorr signatures [48], see [36] for the details). Hence, when this type of attack can be mounted, we should check the order of the public keys.

4 Authentication

In the previous section we presented attacks related to the mathematical structure of the DH protocol primitives. In this section we address issues related to authentication. As the DH protocol can be broken by a simple man in the middle attack (if no authentication mechanism is used), it does not make sense to talk about DH protocol security without also discussing authentication.

Authentication consists of establishing authenticity, which is defined as: *factually accurate and reliable*. This is a somewhat slippery concept and there is no solid and formal definition, because the different settings and requirements change for every application. For example, validating the authenticity of a digital signature or a MAC is simple (just apply a verification function) whereas proving the authenticity of a message is more complicated. For example, if Alice sends a message to Bob, he might want to establish that:

1. the message has not been modified.
2. Alice sent the message.
3. the message was meant for him (i.e. addressed to him).
4. the message has not been “replayed”.
5. the message was sent within a certain time period.
6. etc.

Although we have some very powerful primitives that can help us in creating authentication mechanisms (digital signatures, MACs, symmetric encryption, etc.), using them in an effective manner is surprisingly difficult.

4.1 Message Replay Attacks

One of the deadliest attacks against authentication mechanisms is the message replay attack [39] in which the adversary simply takes a previously sent message and sends it again. This attack is deceptively powerful as can be seen by the next example: suppose a user sent a message to his wife saying, “I love you”. A few years later, after the user has been divorced, an attacker could re-send this same message which might lead to an awkward situation. If a correct authentication mechanism is used, the now ex-wife will not consider the message as being authentic. This example nicely illustrates the fact that digital signatures are not sufficient to establish message authenticity.

4.2 Message Redirection

If the destination is not specified in a message, an attacker can intercept it and send it to someone other than the intended recipient. Taking the previous subsection example’s

premise, the adversary sends the “I love you” message and delivers it to somebody other than the intended recipient; which again, might lead to an uncomfortable situation.

These schemes form the basis of many other, more involved, attacks.

4.3 Message Authentication Protocols

We now present one of the authentication mechanisms described in [6]. Note that the protocol is proved to be secure (see subsection 1.1). It has the property that if we have a scheme that is provably secure when the channels are authenticated, and replace the communication mechanism with the following protocol, then the resulting scheme will be provably secure in a setting in which the channels are not authenticated.

In the following protocol, we assume that the user’s public keys are certified by a certificate authority (CA), and by authentic we mean:

- The sender’s identity is established
- The intended recipient’s identity is established
- “Context” is established; message replay attacks such as the one in the example in subsection 4.1 are prevented.

Alice sends a message m to Bob, who establishes its authenticity.

1. \mathcal{A} sends m to \mathcal{B} .
2. \mathcal{B} replies with a challenge $N_{\mathcal{B}}$ and m , where $N_{\mathcal{B}}$ is a random number (nonce). Note that each *nonce* is only used *once*.
3. \mathcal{A} sends m and $SIG_{s_{\mathcal{A}}}(m, N_{\mathcal{B}}, \mathcal{B})$, where $SIG_{s_{\mathcal{A}}}(x)$ is a digital signature on x that uses \mathcal{A} ’s secret key $s_{\mathcal{A}}$.
4. If the signature verification procedure is successful then m is deemed authentic.

Let us look at this protocol a bit more closely and explain some of its features and propose some efficiency improvements.

First note that sender authenticity is established by the public key that is used to verify the signature. The public key and Alice’s identity are certified by a CA that Bob trusts.

The message m must, of course, be sent at some point. The protocol is still provably secure if the message is sent in one of the first or third rounds. For example, the first message could simply be a synchronization signal.

The need for a nonce $N_{\mathcal{B}}$ is quite interesting. Notice that without it, the protocol would be vulnerable to replay attacks. By slightly modifying the protocol we can, however, do without the nonce. If Bob takes note of all the messages he has received, and so can detect message replay attacks, we can omit the nonce. Unfortunately the amount of data Bob would need to keep could be huge. Also, deterministic signature schemes (e.g. RSA) are not well suited to this setting as two signatures on the same message are identical. (As opposed to probabilistic encryption schemes such as Elgamal.) Note that random numbers can be appended to the message in order to make signatures on the same message different. Another option is to have publicly available nonces. For example a counter can be used, in which case Bob needs to manage counters (synchronization is often difficult to implement). As long as nonces (counter values) are only

used once the protocol is correct. In both of these modifications, the second round of communication can be omitted and the first and third rounds combined.

The last interesting issue to point out is that Alice must sign \mathcal{B} , i.e. Bob's ID. If this is not done, the protocol is vulnerable to message redirection attacks.

Note that the CA *must* verify that the client knows the secret key associated with his public key, otherwise the protocol is vulnerable to message hijacking (i.e. claiming ownership of someone else's message). Note that self signed certificates⁹ can also be used to solve this problem as is done in PGP [44].

5 Attacks on Implementation Details

5.1 Attacks on Parameter Authentication

As a general principle, all parameters used in a cryptographic protocol should be authenticated. For example, suppose that the DH protocol could be used with different system parameters (e.g. g , p); if the participants do not authenticate their choice of parameters, an attacker might be able to fool them into using weak parameters. These types of attacks can be very subtle and can even be missed by top cryptographers and security experts. One need just look at the attack of [52] on the SSL protocol version 2 to be convinced of this¹⁰.

5.2 Context

In many situations it is necessary to make sure an adversary has not blocked (deleted) previous messages. This can be done by simply hashing all previous messages and appending the result with the current message. This establishes *context*. Note that if *all* messages are authenticated we can use a sequence number, which is more efficient.

5.3 Parallel Executions

In most if not all networking protocols, it is very important to preserve protocol run independence. That is, we do not want messages used in one protocol run to be used by another protocol execution. Session numbers, for example, can be used to prevent this kind of problem.

We are aware of a DH implementation that did not respect this design principle. If two parties initiated the DH protocol at the same time, each party obtained two shared DH secret keys and it was possible to have a situation in which none of the (four) supposedly "shared" DH secret keys were equal.

A key agreement confirmation (see 6.5) is a way of making sure problems such as the one described above do not occur.

⁹ certificates that are signed by the subject of the certificate.

¹⁰ SSL version 2 was vulnerable to what is called a version roll-back attack which is an attack on parameter authentication, see [52] for details.

5.4 Deleting the Ephemeral Secrets

It is important to delete the ephemeral secret keys (the secret exponents), to guard against memory being written to disk (swapping) and prevent unwanted access to these values (via a subpoena attack or a system break-in). Deleting private values is usually done by overwriting these with some constant (0s for example). We recommend that the values be deleted as soon as possible to guard against RAM reading techniques such as the ones described in [26].

5.5 Bleichenbacher Type of Attacks

D. Bleichenbacher described in [9] an attack against PKCS #1 v1.5. The attack exploited the fact that some servers implementations of the PKCS #1 v1.5 RSA encryption padding used an inadequate authentication mechanism: if a plaintext started with 0002, as described in the standard, they would blindly accept it as valid and continue, otherwise, they would return an error message to the client. Using a theorem due to Chor [18], Bleichenbacher devised a practical attack against some implementations of SSL v3.0.

Although we do not describe any padding methods, nor do we use the RSA encryption scheme, some of the proposed countermeasures (see [10]) to immunize protocols against this attack are relevant:

- Change keys frequently (as discussed in section 6.2) and make sure that different servers use independent keys.
- Use adequate authentication (as discussed in section 4.3). Servers written in SSL version 3 that used good authentication were not vulnerable to this attack.

5.6 Timing Attacks

An interesting attack was proposed in [34]; the attack relies on the fact that for most modular exponentiation algorithms the time taken is dependent on the inputs. In the Half Certified DH protocol, an attacker, by initiating many protocol runs with Alice and carefully choosing his “public keys”, could determine Alice’s secret key (x) by analysing timing information. Remember that Alice computes m^x in each protocol run (where m can be the attacker’s “public key” (simply a random value) and x is Alice’s secret key). Fortunately, the attack is only effective if the attacker can somewhat precisely determine Alice’s computing time. The attack can be countered by modifying the computations so that the exponentiation time does not depend as heavily on the input parameters.

Koehler [34] gives a method that uses the *blinding* techniques of [17] to randomize the modular exponentiation computing time¹¹ :

One Time Set Up: We calculate the *private* seeds.

- An integer, v_1^0 , is chosen at random from \mathbb{Z}_p^* .
- $v_f^0 = ((v_1^0)^{-1})^x$ is calculated. (Find inverse of v_1^0 and exponentiate by x .)

¹¹ thus blinding the attacker from knowledge about x .

j'th Exponentiation: Let m be the message to be exponentiated by x .

- Calculate $u = (m \cdot v_1^j)$. (the blinding part.)
- Calculate $t = u^x$. (u is not known to the attacker!)
- Calculate $t \cdot v_f^j$ which is equal to m^x . (the unblinding part.)

Computing the j'th seeds ($j > 0$):

- $v_1^j = (v_1^{j-1})^2$.
- $v_f^j = (v_f^{j-1})^2$.

The technique uses the fact that if v_1^0 is chosen randomly, then the series $(v_1^0, v_1^1, \dots, v_1^j, \dots)$ will have the property that v_1^j looks sufficiently random if nothing is known about the previous elements of the series¹². The algorithm has to keep v_1 and v_f in memory. Of course, these values must remain secret.

5.7 Denial of Service Attacks (Overloading)

One of the most damaging attacks in practice consists of overloading servers with requests¹³. This is a type of denial of service attack because the server is so busy processing bogus requests that he does not have time to reply to legitimate queries. The adversary usually exploits the fact that the servers are limited in terms of memory [16, 19] and/or computational power. The DH protocol is vulnerable to the following kinds of attack:

- The attacker can carry out a connection (memory) depletion attack (e.g. [16, 19]). Note that it is very important that the low level protocols for sending and receiving messages be immunized against this attack.
- The attacker can send huge amounts of public keys (which can simply be random numbers) so that the victim is compelled to carry out many modular exponentiations in order to compute the shared DH secret keys (computational).

The most robust solutions [5, 21, 29] to the problem involve having the connection initiators compute solutions to cryptographic puzzles (also known as hashcash or pricing functions). The amount of computations needed to solve these puzzles is small enough so that legitimate users can quickly compute the solution but large enough so that it is infeasible (or at least very hard) to solve a large number of them for use in overloading attacks.

If a server can validate the IP addresses of its clients, one can use a less robust protection scheme called SYN Cookies ([37], [13], [31]). SYN Cookies help prevent IP spoofing to a certain extent.

If a server is to accept unknown clients (or better yet anonymous clients), we suggest using the techniques of [29] which we now present and discuss. Note that $x_{\langle a, b \rangle}$ refers to the substring consisting of the a 'th through b 'th bits of x . \mathcal{C} is the client and \mathcal{S} is the server.

¹² this is much more efficient since modular squaring is a lot cheaper than choosing a new random value.

¹³ see for example <http://www.cisco.com/warp/public/707/newsflash.html>.

1. \mathcal{C} requests a puzzle from the server (stateless connection).
2. \mathcal{S} sends m (the number of sub-puzzles), k (a computation parameter), t (a timestamp) and $x = H(s, t, k, m, C)$. Note that $H()$ is a cryptographic hash function, s is \mathcal{S} 's secret and C is \mathcal{C} 's address.
3. For i equals 1 to m , \mathcal{C} finds z_i such that the first k bits of $x||i||z_i$ equal the first k bits of $H(x||i||z_i)$, where $||$ denotes concatenation. The z_i s (i.e. the solutions to the sub-problems), t and C are sent to \mathcal{S} .
4. \mathcal{S} receives these values and can efficiently check that the solutions are valid and that they have been computed in a timely manner. Note that the server can do these verifications in a stateless matter.

The value of C is usually implicitly determined, for example it might be included in the message headers. The server sends his replies to C and so the adversary must be able to intercept messages addressed to C which is difficult if the adversary is not located at C ¹⁴.

If no time parameters are used, an attacker could obtain a large number of puzzles, solve them (which can take a lot of time) and then overload the server. By encoding t in x using the secret s , the puzzles can be made to have a limited validity period which makes the previous attack infeasible (all t s should be different). (Note also that if connections have an unlimited lifetime, the server is vulnerable to denial of service attacks and so maximum connection lifetime must be taken into consideration when choosing our parameters.)

Attacks are usually of rare occurrences and so it makes sense to be flexible in our use of puzzles. Precisely, we should vary k depending on the situation: The busier the server is, the larger k should be (we can omit puzzles altogether in most situations). s should be large enough so that it cannot be obtained by a brute force attack (see [29]).

6 The DH Shared Secret Key

The shared secret obtained is usually used to derive session keys that will be used in other applications. Now, the operations these keys will be used for have their own requirements and security vulnerabilities. If we are not careful about the way we use the shared DH key, we might be vulnerable to other subtle attacks.

6.1 Key Derivation Function (KDF)

In most, if not all, instances we need to modify the shared secret key obtained in the DH protocol in order to use it with other cryptographic primitives. Here are the main motivations for “modifying” the shared DH secret key:

- The key sizes might not correspond. For example suppose we want to use our a -bit shared secret DH key with a crypto-system requiring a key size of b and $b \neq a$.

¹⁴ this prevents straightforward IP spoofing; this property is also achieved by SYN Cookies ([37], citeRFC1644, [31]).

- Although some bits of the shared secret are provably secure [12] the security of the vast majority of bits in the shared DH secret key is not known (i.e. it is not known whether an attacker can compute knowledge about them¹⁵). The shared DH secret is indistinguishable from an element chosen at random from the group if and only if the Decisional Diffie-Hellman problem in that group is hard (in several groups it is an easy problem), see [11].

Also notice that \mathbb{Z}_p^* does not span all the bit-strings of length p . Hence if we take a random number, chances are greater that the most significant bit equals 0.

Hence, it makes sense to spread the risk and have the bits in the new session key depend on *all* the bits of the shared DH secret key.

- Some attacks exploit algebraic relationships between keys (see section 6.3). Hence, it is important to destroy mathematical structure which can be done using a KDF.
- If we want to create more than one session key with a given shared secret DH key then, if the KDF is a carefully chosen one-way pseudo random number generator, the system can be resistant to known session key attacks (i.e. given a session key, it is hard to find other session keys derived using the same shared secret DH key).

For a KDF, we suggest the use of MGF1 which is described in PKCS #1 (see [30], an implementation of which can be found in [42]). Ignoring some representation details, MGF1 is essentially defined as follows:

$$\text{MGF1}(\text{seed}, \text{out_length}) := \text{HS}(\text{seed} \parallel 0) \parallel \text{HS}(\text{seed} \parallel 1) \parallel \dots \parallel \text{HS}(\text{seed} \parallel c).$$

where HS is a secure hash function and c is a value that depends on the number of bits required (*out_length*). *seed* will be taken to be the value of the DH shared secret from which we want to derive a key.

6.2 Key Freshness and Perfect Forward Secrecy

In many situations the shared DH secret key should be changed frequently. Here are the main reasons why we might want to obtain new shared secret keys often.

1. **Reduce Exposure** The probability that a given key is compromised is lower if it is not used often.
2. **Damage limitation** If the amount of traffic encrypted/authenticated with a given key is reduced then the amount of damage done if the key is compromised is reduced.
3. **Forward Secrecy** If old encryption keys are deleted, encrypted messages can no longer be decrypted. Hence, a third party cannot mount a subpoena attack (i.e. demand that old messages be decrypted).

As expected, tricks used to improve the efficiency of schemes in which keys are changed often have subtle problems (see subsection 6.4).

¹⁵ in fact, given only p, g, g^x and g^y in \mathbb{Z}_p^* , we can easily compute the Jacobi symbol of g^{xy} , and if g generates the whole group then we can also efficiently compute the last bit of xy .

6.3 Key Independence

As a general principle, we always want keys to be independent. Precisely, obtaining one secret key should *not* help an attacker uncover other keys. This property is called *known key security*. In the next subsection, we give an example of a protocol vulnerable to known key attacks.

6.4 An Example

We now present a condensed version of the KEA protocol [40] which is a part of the NSA's FORTEZZA suite of cryptographic algorithms and motivate the use of key derivation functions. Note that the explanations roughly follow those of [8].

1. \mathcal{A} gets \mathcal{B} 's static public key g^y and \mathcal{B} gets \mathcal{A} 's static public key g^x (x is secret to \mathcal{A} and y is secret to \mathcal{B}). (respectively) certified by a CA. (x is Alice's private key and y is Bob's private key.)
2. \mathcal{A} sends g^a , g^x and $\text{Cert}(\mathcal{A}, g^x)$ to \mathcal{B} and \mathcal{B} sends g^b , g^y and $\text{Cert}(\mathcal{B}, g^y)$ to \mathcal{A} . (a and b chosen randomly from the set $\{2, \dots, p-1\}$) Note that $\text{Cert}(x)$ is just a certificate certifying x .
3. If all verifications succeed, the shared DH secret key is taken to be $K = g^{ay} + g^{bx}$.
4. A key derivation function (derived from SKIPJACK) is then applied to K to obtain the key (the session key) that will be used in the other applications (e.g. encryption, MAC, etc.).

The protocol solves many of the problems mentioned in the previous subsections:

- **Key Freshness:** We can obtain as many fresh keys as we need without having the CA re-certify new public keys every time.
- **Forward Secrecy:** If Alice and Bob delete K and *both* the static and ephemeral secret keys (x and a respectively for Alice) we have forward secrecy.
- **Key Independence:** The protocol *seems* resistant to known key attacks.
- **Key Derivation Function:** The session key depends on all of the bits of the shared DH secret key. As will be seen shortly, the key derivation function is also important because it destroys the algebraic relationships between keys.

If the protocol did not use a key derivation function, it would be vulnerable to the Burmester triangle attack [14] which renders the protocol vulnerable to known key attacks. In the previous protocol, if a key derivation function is not used, it is vulnerable to the following attack:

1. \mathcal{O} first observes a protocol run between \mathcal{A} and \mathcal{B} . He obtains the ephemeral keys g^a and g^b . The key shared by \mathcal{A} and \mathcal{B} at the end of the protocol is $K_{\mathcal{AB}} = g^{ay} + g^{bx}$.
2. \mathcal{O} then engages \mathcal{A} in a protocol run. \mathcal{O} will use g^b as his ephemeral key and g^z as his static key. Assuming \mathcal{A} 's ephemeral key is $g^{\bar{a}}$, the shared key will equal $K_{\mathcal{AO}} = g^{bz} + g^{z\bar{a}}$.
3. \mathcal{O} carries out the same trick with \mathcal{B} but now uses g^a as his ephemeral key. Assuming that \mathcal{B} 's ephemeral key is $g^{\bar{b}}$, the shared key will be $K_{\mathcal{BO}} = g^{az} + g^{z\bar{b}}$.
4. If \mathcal{O} can obtain $K_{\mathcal{AO}}$ and $K_{\mathcal{BO}}$ he can determine $K_{\mathcal{AB}}$. This can be seen by noting that $K_{\mathcal{AB}} = K_{\mathcal{AO}} + K_{\mathcal{BO}} - g^{bz} - g^{z\bar{a}}$.

6.5 Key Agreement Confirmation

In some settings, the participants will not settle with just knowing that nobody except the intended party can compute the session key (i.e. a key derived from a shared secret DH key) but insist on having some kind of confirmation that a secret key has been (or can be) successfully created. A scheme provides *implicit* key confirmation if the participants can be convinced that they all *can* compute a common shared secret key, and provides *explicit* key confirmation if participants can be assured that a common shared secret key *has* been computed by all participants. The simple minded solution to providing explicit key agreement is to have the parties compute the MAC (using the new session key) of a known message. Unfortunately this means that the key will be distinguishable from a random key (we know the MAC of a known message). If key indistinguishability is required we need to use (as a MAC key) some other value, r , known only to the participants that cannot be easily linked to the session key. Precisely, given the session key it should be computationally infeasible to find r . See [8] for techniques that can be used to do this effectively.

Although explicit key confirmation appears to provide stronger assurances, implicit key confirmation is sufficient in practice and provides key indistinguishability. Also, it would seem that although it is possible to provide explicit confirmation of the derived shared secret *key* without using any previous shared secret, it is not obvious how to provide explicit confirmation of the DH *shared secret* in an efficient way without using a previously shared secret, so it's usefulness is questionable.

7 The Bottom Line

In this section, we give recommendations that are, for the most part, based on the lessons learned in the previous sections. These can be seen as general robustness principles¹⁶ that should be taken into account when implementing DH key agreement type protocols.

Note that DH protocol implementations should be especially wary of attacks that allow the attacker to obtain *static* secret keys (e.g. x, y). Compromising static secret keys allows the attacker to break all subsequent protocols using these values. Compromising g^{xy} does not help in compromising other shared DH secret keys $g^{xy'}$.

7.1 Diffie-Hellman Math

1. Spot Unconventional Messages

- Make sure that g^x , g^y and g^{xy} do not equal 1.
- Make sure that g^x and g^y are less than $p - 1$ and greater than 1.
- Choose x, y , from the set $\{2, \dots, p - 2\}$.

2. Be Careful About g 's Order

- The prime factor decomposition of the order of g should not be composed entirely of small primes.
- The subgroup generated by g should not have a small order subgroup. If at all possible, construct and use a generator that has a large prime order.

¹⁶ a DH version of [3].

3. **Make Sure the DH Public Keys Received Have the Correct Order**
 - The DH public key's (g^x) order should be checked. This can be easily done by verifying that $(g^x)^{\bar{o}} = 1$ where \bar{o} is g 's order. If $p = 2q + 1$ this is not necessary as explained in section 3.6.
4. **Make Sure the System Parameters Are Not Chosen Maliciously**
 - The system parameter's properties should be known (subgroup generated by g 's order, prime factorization of this number, etc).
 - Proofs that the parameters have not been chosen maliciously should be available. This can be done by kosherizing¹⁷ One can transform a typical strong prime number generator into one that generates kosherized primes. For a strong prime number generator that starts by choosing a random q and then verifies (using some efficient ways) that q and $2q + 1$ are prime, one can replace the choice of q by first choosing $q = \text{HS}(r)$, where r is a random value and HS is a secure hash function whose output is the same size as q . The kosherization is "proven" by giving r such that $q = \text{HS}(r)$ and $p = 2q + 1$.
5. **Choose Secure Parameters**
 - Cryptographic algorithms are only as secure as their weakest link and so it makes sense to try and balance the security. That is, attacks that exploit different parameters of the system should take roughly the same amount of time. For the DH protocol, the parameters to balance are : the value of p , the exponent's range and the size of the keys derived from the shared DH secret. We suggest looking at [35] for a table of balanced values. Note that these values are very controversial [49], the size of p is especially debatable since it assumes Moore's law type improvements in algorithmic number theory.
 - The parameters should be chosen in order to provide good long term security when required. Note that parameters that constitute "good" long term security is very controversial [35, 49]. Extremely conservative estimates are (from [35]):
 - For very good security until 2002 take: p 1024 bits, exponent range 127 bits and derived key length 72.
 - For very good security until 2025 take: p 2174 bits, exponent range 158 bits and derived key length 89.
 - For very good security until 2050 take: p 4047 bits, exponent range 193 bits and derived key length 109.
 - We suggest using strong primes or Lim Lee primes so as to guard against the attacks presented in section 3.5.
 - The number of symmetric keys derived from the shared DH secret key should also be taken into consideration when determining the size of p and of the exponent range since breaking the DH protocol breaks *all* derived keys. Precisely, if we derive n session keys of lengths n_1, n_2, \dots, n_k , our other parameters should, *in theory*, provide the same security as if we derived one session key of length $\lg(2^{n_1} + 2^{n_2} + \dots + 2^{n_k})$.

These measures protect against some insider attacks and man in the middle attacks.

¹⁷ kosherizing a public value refers to constructing the value in a way that there exists a proof of the fact that the value has not been chosen maliciously, this proof is to be verifiable by any other participant

Efficiency Considerations

1. Ideally, the generator g should be as small as possible in order to reduce the cost of modular exponentiation. Wiener and van Oorschot [51] claim that using $g = 2$ reduces the computation time for modular exponentiation by 20% (compared to randomly selected generators). For applications in which efficiency is crucial and the prime numbers can be generated beforehand, it makes sense to find a prime, p , such that a small value (e.g. 2,3,16) generates the desired subgroup. Note that if one is searching for a strong prime $p = 2q + 1$ such that $g = 2$ generates a subgroup of size q , one can simply test if $p = 7 \pmod 8$. In fact, $p = 7 \pmod 8 \iff^{18} 2$ is a quadratic residue¹⁹ $\pmod p \implies 2^{\frac{1}{2}}$ exists $\implies 2^{((p-1)/2)} = (2^{\frac{1}{2}})^{p-1} = 201 \pmod p \iff^{21} g = 2$ generates an order q subgroup.

2. Generating safe primes²² is more expensive than generating Lim-Lee primes ($p = 2q_1 q_2 \dots q_n + 1$). If lots of primes need to be generated *and* efficiency is an important requirement, we suggest using Lim-Lee primes [36] which are used in many cryptographic libraries (e.g. PGP [44], GNU PG [23] and Gutmann's cryptlib [25]). These primes have the form $p = 2q_1 q_2 \dots q_n + 1$ where the q_i s are large (for all $i \in \{1, \dots, n\}$). The generator can be taken to generate some prime order subgroup (e.g. of order q_i , for some i). A drawback to this method is that the range of values exponents can take is limited (i.e. exponents are taken modulo q_i instead of q) which restricts the range of possible DH secrets. Also note that the probability that a small generator generates an adequate subgroup is lower than for safe primes.

If the parameters are fixed, we suggest the use of Sophie Germain primes since they allow the use of larger exponents (thus resulting in larger shared secrets). Note that there exist particular primes that yield more efficient operations. [43] suggests the use of "special" safe primes which are used in the description of IKE [27] (a candidate DH protocol for IPsec). They have properties that enable efficient modular computations to a certain extent:

- The 64 high order bits are set to 1, so that the trial quotient digit in the classical remainder algorithm can always be set to 1.
- The 64 low order bits are also set to 1, which enables speed ups of Montgomery style remainder algorithms.
- The middle bits are taken from the binary expansion of π which provides a weak form of kosherization.
- $g = 2$ is a generator of a subgroup of order $(p - 1)/2$ (g has prime order).

These primes can be found in Appendix E.2 (1024 bits) and E.5 (1536 bits) of [43] and can be used in a DH.

¹⁸ see for example Fact 2.146 in [38]. We can eliminate the case where $p = 1 \pmod 8$ since p and q are prime. (left as an exercise.)

¹⁹ $x \in \mathbb{Z}_p^*$ is quadratic residue of \mathbb{Z}_p^* iff there exists a $y \in \mathbb{Z}_p^*$ such that $y^2 = x \pmod p$. If no such y exists, x is called a quadratic non-residue.

²⁰ the equality comes from the fact that $p - 1$ is the order of the group.

²¹ 2 's order is either 2, q or $p - 1$ (Lagrange's theorem, section 2.1). If $2^q = 1 \pmod p$, then 2 's order can only be 2 or q , but the only elements of order 2 are 1 and $p - 1$.

²² a prime p of the form $p = 2q + 1$ where q is prime (this q is often referred to as a Sophie Germain prime).

3. Exponentiations are usually much faster when the exponents are small and so we suggest using the smallest secure exponent range (see subsection 7.1).

7.2 Implementation Details

Correctly establishing authenticity is difficult and whenever possible, provably secure authentication protocols should be used (at the very least, the attacks mentioned previously must be taken into account). Particular care must be taken when improving a protocol's efficiency (e.g. removing "superfluous" messages).

Note the following tricky implementation level issues:

1. **Exact Destination:** The message recipient should be precisely specified. Identification fields could include IP address, port number, user ID, etc.
2. **Multiple Session Management:** It is of crucial importance for participants to separate concurrent protocol executions. Concurrent protocol executions should be *independent*. This problem can usually be dealt with by adding a session ID field to the messages. Note that this is a tricky problem to solve when sessions are related in some way, for example when counters are used instead of nonces.
3. **Certifying Public Static Keys:** The certifying authority should, of course, be trusted by both participants. As pointed out earlier the certificate authority should make sure that the certificate recipient knows the secret key associated with the public key being certified. Another option is to have the sender sign his identifier along with the rest of the message (i.e. self signed certificates).
4. **Authenticate Parameters:** *All* parameters should be authenticated.
5. **Previous Communications:** In many situations, all messages sent should "confirm" all previous messages. We want to avoid some messages being blocked. This can be done by appending a hash of all previous messages or by using sequence numbers.
6. **Delete:** When some sensitive information is no longer needed it should be securely deleted.
7. **Randomness:** The pseudo-random numbers must be chosen extremely carefully because systems can be broken if inadequate pseudo-random functions or badly chosen seeds are used (see [24]). A good pseudo-random number generator is Yarrow [32], since its design is based on many years of research and experience [33] and because it is easy to use (the programmer does not need to provide a seed for example).
8. **Timing Attacks:** When a system is vulnerable to timing attacks (see section 5.6), a special exponentiation routine should be used.
9. **Denial of Service Attacks:** When necessary (see section 5.7), parties should protect themselves against denial of service attacks.

7.3 Using the Shared DH Secret Key

Here are the general points related to the utilization of the shared DH secret key.

1. **Never Use the Key "As Is":** Always use a *suitable* key derivation function in order to get a session key.

2. **Key Independence:** It is important for the protocols to be resistant to known key attacks.
3. **Delete Old Keys:** If forward secrecy is desired old keys and all data that can be used to obtain them must be securely deleted.
4. **Be Careful with Confirmation:** If key indistinguishability is required, we cannot just send the MAC of a known message.

8 Conclusions

This work has attempted to present cryptographic protocol designers with the most important security issues related to the DH protocol. In doing so, we have addressed the shortcomings of the other approaches to secure cryptographic protocol design. It is hoped that documents with a form similar to this one for different cryptographic protocols will be produced. This would be a large step towards assuring cryptographic protocol security in real-world settings.

References

- [1] ABADI, M., AND GORDON, A. D. A calculus for cryptographic protocols: The spi calculus. *Information and Computation* 148, 1 (10 Jan. 1999), 1–70.
- [2] ANDERSON, R., AND NEEDHAM, R. Programming satan’s computer. In *Computer Science Today: Recent Trends and Developments* (1995), J. van Leeuwen, Ed., vol. 1000 of *Lecture Notes in Computer Science*, Springer.
- [3] ANDERSON, R., AND NEEDHAM, R. Robustness principles for public key protocols. In *Advances in Cryptology – CRYPTO ’95* (1995), D. Coppersmith, Ed., Lecture Notes in Computer Science, Springer-Verlag, Berlin Germany.
- [4] ANDERSON, R., AND VAUDENAY, S. Minding your p ’s and q ’s. In *Advances in Cryptology—ASIACRYPT ’96* (Kyongju, Korea, 3–7 Nov. 1996), K. Kim and T. Matsumoto, Eds., vol. 1163 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 26–35.
- [5] BACK, A. Hashcash. <http://www.cypherspace.org/adam/hashcash/>, mar 1997.
- [6] BELLARE, M., CANETTI, R., AND KRAWCZYK, H. Modular approach to the design and analysis of key exchange protocols. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)* (New York, May 23–26 1998), ACM Press, pp. 419–428.
- [7] BELLARE, M., AND ROGAWAY, P. Entity authentication and key distribution. In *Advances in Cryptology – CRYPTO ’93* (1994), D. R. Stinson, Ed., vol. 773 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany.
- [8] BLAKE-WILSON, S., AND MENEZES, A. Authenticated Diffie-Hellman key agreement protocols. In *Fifth Annual Workshop on Selected Areas in Cryptography (SAC ’98)* (1999), Lecture Notes in Computer Science, Springer Verlag, pp. 339–361.
- [9] BLEICHENBACHER, D. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. *Lecture Notes in Computer Science* 1462 (1998), 1–??
- [10] BLEICHENBACHER, D., KALISKI, B., AND STADDON, J. Recent results on PKCS #1 RSA encryption standard. *RSA Laboratories’ Bulletin*, 7 (1998).
- [11] BONEH, D. The decision Diffie-Hellman problem. *Lecture Notes in Computer Science* 1423 (1998), 48–63.

- [12] BONEH, D., AND VENKATESAN, R. Hardness of computing the most significant bits of secret keys in Diffie-Hellman and related schemes (extended abstract). In *Advances in Cryptology—CRYPTO '96* (18–22 Aug. 1996), N. Koblitz, Ed., vol. 1109 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 129–142.
- [13] BRADEN, R. RFC 1644: T/TCP — TCP extensions for transactions functional specification, July 1994. Status: EXPERIMENTAL.
- [14] BURMESTER, M. On the risk of opening distributed keys. In *Advances in Cryptology — CRYPTO '94* (1994), Y. G. Desmedt, Ed., vol. 839 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 308–317.
- [15] BURROWS, M., ABADI, M., AND NEEDHAM, R. A logic of authentication. *ACM Transactions on Computer Systems* 8, 1 (Feb. 1990), 18–36.
- [16] CERT. Advisory ca-96.21: Tcp syn flooding and ip spoofing attacks, 24 September 1996.
- [17] CHAUM, D. Blind signatures for untraceable payments. In *Advances in Cryptology: Proceedings of Crypto 82* (23–25 Aug. 1982), D. Chaum, R. L. Rivest, and A. T. Sherman, Eds., Plenum Press, New York and London, 1983, pp. 199–203.
- [18] CHOR, B.-Z. *Two issues in public key cryptography: RSA bit security and a new knapsack type system*. ACM distinguished dissertations. MIT Press, Cambridge, MA, USA, 1986. Originally presented as the author's thesis (doctoral — MIT, 1985).
- [19] DAEMON9. Project neptune. Phrack Magazine, 48(7): File 13 of 18, 8 November 1996. Available at www.fc.net/phrack/files/p48/p48-13.html.
- [20] DIFFIE, W., AND HELLMAN, M. New directions in cryptography. *IEEE Transactions on Information Theory* 22 (1976), 644–654.
- [21] DWORK, C., AND NAOR, M. Pricing via processing or combatting junk mail. *Lecture Notes in Computer Science* 740 (1993), 139–147.
- [22] ELGAMAL, T. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology: Proceedings of CRYPTO 84* (19–22 Aug. 1984), G. R. Blakley and D. Chaum, Eds., vol. 196 of *Lecture Notes in Computer Science*, Springer-Verlag, 1985, pp. 10–18.
- [23] FSF. Gnu privacy guard. <http://www.gnupg.org/>.
- [24] GOLDBERG, I., AND WAGNER, D. Randomness and the Netscape browser. *Dr. Dobbs's Journal of Software Tools* 21, 1 (Jan. 1996), 66, 68–70.
- [25] GUTMANN, P. <http://www.cs.auckland.ac.nz/pgut001/cryptlib/>.
- [26] GUTMANN, P. Secure deletion of data from magnetic and solid-state memory. In *6th USENIX Security Symposium* (San Jose, California, July 1996), USENIX.
- [27] HARKINS, D., AND CARREL, D. RFC 2409: The Internet Key Exchange (IKE), Nov. 1998. Status: PROPOSED STANDARD.
- [28] HUNGERFORD, T. W. *Algebra*. Holt, Rinehart and Winston, New York, 1974.
- [29] JUELS, A., AND BRAINARD, J. Client puzzles: A cryptographic defense against connection depletion attacks. In *NDSS '99 (Networks and Distributed Security Systems)* (2000), S. Kent, Ed., pp. 151–165.
- [30] KALISKI, . B., AND STADDON, J. RFC 2437: PKCS #1: RSA cryptography specifications version 2, Oct. 1998. Obsoletes RFC2313. Status: INFORMATIONAL.
- [31] KARN, P., AND SIMPSON, W. A. The photuris session key management protocol. Internet Draft, Dec. 1995. Version 0.8, expires June 96.
- [32] KELSEY, J., SCHNEIER, B., AND FERGUSON, N. Notes on the design and analysis of the yarrow cryptographic pseudorandom number generator. In *Sixth Annual Workshop on Selected Areas in Cryptography* (Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1999), Springer-Verlag, p. ????
- [33] KELSEY, J., SCHNEIER, B., WAGNER, D., AND HALL, C. Cryptanalytic attacks on pseudorandom number generators. *Lecture Notes in Computer Science* 1372 (1998), 168–188.

- [34] KOCHER, P. Cryptanalysis of Diffie-Hellman, RSA, DSS, and other cryptosystems using timing attacks. In *Advances in cryptology, CRYPTO '95: 15th Annual International Cryptology Conference, Santa Barbara, California, USA, August 27–31, 1995: proceedings* (Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1995), D. Coppersmith, Ed., vol. 963 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 171–183.
- [35] LENSTRA, A. K., AND VERHEUL, D. E. R. Selecting cryptographic key sizes. <http://www.cryptosavvy.com/>, Nov. 1999. Shorter version of the report appeared in the proceedings of the Public Key Cryptography Conference (PKC2000) and in the Autumn '99 PricewaterhouseCoopers CCE newsletter.
- [36] LIM, C. H., AND LEE, P. J. A key recovery attack on discrete log-based schemes using a prime order subgroup. In *Advances in Cryptology – CRYPTO '97* (Aug. 1997), B. S. K. Jr., Ed., *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 249–263.
- [37] MAUGHAN, D., SCHERTLER, M., SCHNEIDER, M., AND TURNER, J. RFC 2408: Internet Security Association and Key Management Protocol (ISAKMP), Nov. 1998. Status: PROPOSED STANDARD.
- [38] MENEZES, A. J., VAN OORSCHOT, P. C., AND VANSTONE, S. A. *Handbook of applied cryptography*. The CRC Press series on discrete mathematics and its applications. CRC Press, 2000 N.W. Corporate Blvd., Boca Raton, FL 33431-9868, USA, 1997.
- [39] MITCHELL, C. Limitations of challenge-response entity authentication. *Electronics letters*, 25 (August 1989), 1195–1196.
- [40] NSA. Skipjack and kea algorithm specification (version 2.0), May 1998. Also available at <http://csrc.nist.gov/encryption>.
- [41] ODLYZKO, A. M. Discrete logarithms: The past and the future. *Designs, Codes, and Cryptography* 17 (1999).
- [42] OPEN-SSL. <http://www.openssl.org/>.
- [43] ORMAN, H. RFC 2412: The OAKLEY Key Determination Protocol, Nov. 1998. Status: INFORMATIONAL.
- [44] PGP, I. <http://www.pgpi.org/>.
- [45] POHLIG, S., AND HELLMAN, M. An improved algorithm for computing discrete logarithms over $GF(p)$ and its cryptographic significance. *IEEE Transactions on Information Theory* 24 (1978), 106–110.
- [46] POLLARD, J. M. Monte Carlo methods for index computation (mod p). *Mathematics of Computation* 32, 143 (July 1978), 918–924.
- [47] RIVEST, R. L., AND SHAMIR, A. How to expose an eavesdropper. *Communications of the Association for Computing Machinery* 27, 4 (Apr. 1984), 393–395.
- [48] SCHNORR, C. P. Efficient identification and signatures for smart cards. In *Advances in Cryptology—EUROCRYPT 89* (10–13 Apr. 1989), J.-J. Quisquater and J. Vandewalle, Eds., vol. 434 of *Lecture Notes in Computer Science*, Springer-Verlag, 1990, pp. 688–689.
- [49] SILVERMAN, R. A cost-based security analysis of symmetric and asymmetric key lengths. RSA Laboratories Bulletin, april 2000.
- [50] STINSON, D. R. *Cryptography: Theory and Practice*. CRC Press, 1995.
- [51] VAN OORSCHOT, P. C., AND WIENER, M. J. On Diffie-Hellman key agreement with short exponents. In *Advances in Cryptology – EUROCRYPT '96* (1996), U. Maurer, Ed., *Lecture Notes in Computer Science*, Springer-Verlag, Berlin Germany, pp. 332–343.
- [52] WAGNER, D., AND SCHNEIER, B. Analysis of the SSL 3.0 protocol. Technical report, Counterpane Systems, 101 East Minnehaha Parkway, Minneapolis, MN 55419, 1996. Also published in *The Second USENIX Workshop on Electronic Commerce Proceedings*, USENIX Press, November 1996, pp. 29–40.

A Standards

A.1 PKCS #3

RSA Security ²³ has published a suite of cryptography standards which are called Public Key Cryptography Standard (PKCS). PKCS #3 deals with the DH protocol. Unfortunately, it does not help the protocol designer construct a secure version because it only specifies data formats.

A.2 ANSI X9.42 – Agreement of Symmetric Algorithm Keys Using Diffie-Hellman

working draft may 1998

A.3 IETF RFC 2522 – Photuris: Session-Key Management Protocol

march 1999

A.4 ANSI X9.63 – Elliptic Curve Key Agreement and Key Transport Protocols

working draft July 1998

A.5 IEEE P1363 – Standard Specifications for Public-Key Cryptography

working draft July 1998

A.6 ISO/IEC 11770-3 – Information Technology - Security Techniques - Key Management - Part 3: Mechanisms Using Asymmetric Techniques

draft (DIS), 1996

A.7 SKIPJACK and KEA algorithm specification

from FORTEZZA may 1998.

A.8 The Internet Key Agreement (IKE)

RFC 2409 November 1998.

A.9 The OAKLEY key Determination Protocol

RFC 2412 November 1998.

A.10 The TLS Protocol: Version 1.0

RFC 2246 January 1999

²³ see <http://www.rsasecurity.com>