

Secure Sessions from Weak Secrets

Bruce Christianson¹, Michael Roe², and David Wheeler³

¹ Computer Science Department, University of Hertfordshire, Hatfield

² Microsoft Research Limited, Cambridge

³ Computer Laboratory, University of Cambridge,
England, Europe

Abstract. Sometimes two parties who already share a weak secret k such as a password wish to share also a strong secret s such as a session key without revealing information about k to an active attacker. We assume that both parties can generate strong random numbers and forget secrets, and present new protocols for secure strong secret sharing, based on RSA, Diffie-Hellman, and El-Gamal. As well as being simpler and quicker than their predecessors, our protocols also have stronger security properties. In particular, our protocols make no cryptographic use of s and so do not impose subtle restrictions upon the use which is subsequently made of s by other protocols. Neither do we rely upon the existence of hash functions with serendipitous properties. In the course of presenting these protocols, we also consider how to frustrate some new types of cryptographic and system attack.

1 Introduction

Sometimes there is a requirement to establish a secure session between two parties who initially share only a weak long-term secret. “Secure” includes the requirement that the parties can be sure that they are talking to each other, as well as properties of integrity and secrecy. By “weak secret” we mean a secret that is chosen from a moderately small set, so that an attacker could search through all possible values. Passwords are often weak secrets, as the total number of words in a dictionary is searchable.

A weak secret cannot be used directly as a cryptographic key to secure the session, as this is vulnerable to a known plaintext attack. If the attacker knows (or can guess with high probability of being right) the message plaintext m corresponding to a known encrypted text $E_k(m)$ then they can search through all possible values of the password until they find the value k which decrypts the cyphertext to m . This reveals the password, which can then be used to decipher the session.

Suppose that the parties who wish to communicate have good random number generators. This means that they can generate secrets which are strong (chosen from a set which is too large to search) but not shared. We would like to have a protocol which starts with a weak shared secret and a pair of strong non-shared secrets and which ends up with a secret which is both strong and

shared. We refer to such a protocol as a Strong Secret Sharing Password (S3P) Protocol.

Previous attempts at solving this problem include Bellare and Merritt's Encrypted Key Exchange [3], Jablon's SPEKE [9], and Lucks' Open Key Exchange [12]. A related, but slightly different approach is taken by Gong *et al* in [7].

In the next section, we give a careful specification of the properties which we desire S3P protocols to have, and of the assumptions under which we believe our protocols to have these properties. In Section 3, we present an S3P protocol based on RSA and show that it resists a number of attacks, some known and some novel. In section 4 we consider various DH-based S3P protocols which have stronger security properties than the conventional versions, and discuss some novel threats. In section 5 we consider the implications of using EG in place of DH when the requirement is to transfer an existing secret rather than to agree a new one. In section 6, we discuss what should be done when an active attack is detected, and introduce the notion of a robust protocol wrapper. In the final section we summarize our conclusions.

2 S3P Protocol Properties and Assumptions

In this section, we give a careful specification of the properties which we desire S3P protocols to have, and of the assumptions under which we believe our protocols to have these properties.

The S3P protocols described in this paper are peer-to-peer protocols, which operate directly between two principals rather than between a principal and a server. In accord with tradition, we assume that the two parties trying to operate the S3P protocol are unambiguously known to each other as A and B .

2.1 Properties

We turn now to describing the features which we desire the S3P protocol to have. At the start of the protocol A and B share a weak secret k . Following a run of the protocol which A believes to have ended correctly, it should be the case that B really did participate in that run of the protocol, and that the two of them really do now share a fresh strong secret s . The corresponding statement should also be true for a run which B believes to have ended correctly. The protocol should not reveal information about the weak secret k in any case.

The protocol should be secure against active attacks in which the attacker creates or modifies messages. Leakage or cryptographic compromise of a strong secret s shared using a protocol run should not reveal information about the password k . If several strong secrets s_i are shared by different runs using the same password k then obtaining one such s_i should not help an attacker to obtain s_j with $j \neq i$.

An attacker should not be able to obtain any information about whether a guessed value of the password is correct without making an active attack: effectively the attacker should be forced to masquerade as one of the participants to the other. An active attack should be detectable by at least one of the

genuine participants, unless the guessed value is correct, and each such failed attack should eliminate no more than one possible value of the password (ie the unsuccessful guess) from the attacker's list of possible password values. Approaches such as [1] do not satisfy this requirement. Finally, if the password is compromised (by whatever means) this should not assist the attacker to obtain strong secrets agreed using the protocol with that password prior to the point of compromise, or to obtain subsequently agreed strong secrets by passive attack.

2.2 Assumptions

We assume that neither party can reliably maintain the integrity of a strong secret s from one protocol run to another: in other words if A tries to use a strong secret from one run in another run, then there is a good chance that s either leaks, or is forgotten, or changes (or is changed) without A noticing that it has. This assumption may correspond to the fact that the parties move frequently from one piece of hardware to another, or may be because the hardware is initialized in some way between protocol runs to erase secret information. We discuss this issue further in section 6.

We assume that both parties can reliably maintain the integrity of public, slowly varying data such as software and public keys: other protocols are available to assist with this [11].

The protocols which we consider include the operations “generate a random bit pattern n ” and “forget the bit pattern m ”. We assume that both ends have good irreproducible random bit generators and can forget secrets. By the first assumption we mean that our threat model does not consider the possibility of an attacker determining n by examining other bit patterns produced (previously or subsequently) by the same or other generators. By the second assumption we mean that our threat model does not consider the possibility of an attacker subsequently determining m from an examination of the hardware which has been instructed to forget it. Note that the hardware which must forget includes the random generator. This assumption is probably the most difficult requirement to realize in practice.

We wish to make no assumptions about what the strong shared secret will be used for. The S3P protocol run used to agree the strong shared secret s ends as soon as both parties can be sure that s has been appropriately shared. In the light of known chosen protocol attacks [10] we wish to impose no restrictions upon the nature of the cryptographic protocols or algorithms to which s is handed off for subsequent use, or upon the length of s itself. It may be intended to reveal s (for example s may be used as a one-time pad) or it may be that s is not intended to be used as a key at all, but as a salt or initial value. In particular, the S3P protocol should not assume that s is strong: it may be feasible for an attacker to search for s in the time available between steps of the protocol. This may be because s is required to be weak (less than 40 bits, suppose) or because the protocol is running very slowly (the messages may be carried by a diskette sent through the post, for example.)

2.3 System Context

In the protocol descriptions that follow, we omit from each message the header information identifying which protocol is being used, which parties it purports to operate between, which run of the protocol the message relates to, and the sequence number of the message within that protocol run. These data will be conveyed by the outer wrapping protocol, which will be discussed in section 6. Also, we have not yet specified explicitly what conditions cause a participant to treat a particular run as having failed (eg receiving an incorrect bit pattern or timing out.) We shall also consider these points further in section 6, but first we describe the inner S3P protocols themselves.

3 RSA Based Protocol

In this section, we present new S3P protocols based on RSA. Correct protocols of this form have previously resisted construction, to the point where some have doubted their possibility. We show that this RSA-S3P protocol resists the known attacks as well as some novel ones.

3.1 Basic RSA-S3P Protocol Description

A generates an RSA modulus $N = pq$ with p, q prime and so that $(p-1)/2, (q-1)/2$ each contains a large prime factor. We assume that the bit-lengths of p, q and N are prescribed exactly.

We assume that there is a publicly known function e which converts a password k into a large prime number $e(k)$ suitable for use as an RSA exponent. By large we mean that the bit length of $e(k)$ is strictly larger than that of N by some suitable algorithm. For example, supposing that the password k is encoded in such a way that the bit length of k is small relative to that of \sqrt{N} , then one suitable algorithm would be to search through ascending values of i until a prime of the form $1 + k(a + ib)$ is found, where a, b are published co-prime constants guaranteed to exceed N .

The RSA-S3P protocol runs as follows:

$$A \rightarrow B : N \tag{1}$$

$$B \rightarrow A : z^{e(k)} + 2^{e(k)} \bmod N \tag{2}$$

$$A \rightarrow B : n_a \tag{3}$$

$$B \rightarrow A : n_b \tag{4}$$

Here $z = c|s|n_a|n_b$ where s is the session key, c is a strong random number called a *confounder*, and n_a, n_b are random numbers called *nonces*. The vertical bar $|$ denotes concatenation of bit strings, with the high order bits on the left.

Note that a fresh key N is required for each run of the protocol, but the function $e(k)$ can remain constant. Only A need verify the strength of the public key N , although B must check that N has the correct bit-length. A must forget

$d(k)$, the decryption key corresponding to $e(k)$, as well as p and q . Both A and B must forget c . B must forget the whole of z if the protocol fails at step (3).

In the RSA protocol, the key s and all nonces are chosen by B . Although s need not be strong, it must contain no redundancy and must not be predictable. Prior knowledge of the value of s which B will choose allows an attacker masquerading as A to determine k . Also note that s must appear random and so can't be a public key. Although the confounder c must be strong, the nonces n_a and n_b need not be strong, although they should be significantly harder to predict than k .

Jolyon Clulow points out that A must reject an offered value of 0 in message (2), or an attacker can force $z = N - 2$ without knowing or learning k .

3.2 Design Discussion

The presence of c prevents an attacker using a compromised session key and a copy of message (2) to search for k .

It is vital that there be no redundancy in the plaintext $z = c|s|n_a|n_b$ which is encrypted in message (2). If there were, then an attacker masquerading as A could use this to search for k in time to generate message (3) correctly and complete the protocol run. Note that consequently z must be a random number in the range $1 \dots N$. The statistical distribution of the high order bits of c is thus skewed, because N is not a power of 2. However, this effect dies away exponentially with bit order, so the low order bits of c plus all bits of s cannot contain enough skew to be useful to an enemy after any achievable number of protocol runs. It must be infeasible for an attacker to search over the low-order bits of c . Otherwise after s is revealed for a completed run, passive search would reveal k by a match on message (2).

The purpose of messages (3) and (4) is to convince A and B that they are not experiencing an active attack. Instead of using the nonces n_a and n_b in messages (3) and (4), we could use cryptographic hashes of them instead. But nothing is gained by doing this, and it requires us to exhibit a cryptographic algorithm with suitable subtle properties, a commitment which we prefer to avoid. (3) and (4) by an attacker masquerading as B , who finds s in time to send message (4). Effectively, if s is used in this way and is not sufficiently strong, then the attacker gets an undetected guess at k .

It is tempting to try and shorten the protocol run to three messages by combining the texts of messages (2) and (4) into a single message. This doesn't work, because there must be no redundancy in message (2).

3.3 Number-Theoretic Attacks

Factorization of N by an attacker gives the attacker k and, worse, allows the attacker to obtain old values of s . The public key N must therefore be many times longer than s , consequently a large number of bits is available for n_a, n_b and c . However the protocol must specify the exact bit-length of N , and B must check that N has the correct number of bits, in order to ensure that an attacker does not insert an extra factor into N in order to gain residue information about $e(k)$.

The term $2^{e(k)}$ in the second message is required to block the *Bleichenbacher attack*: if this term is omitted, then an attacker masquerading as B can send $z^e \bmod N$ in message (2), where $e = \prod_i e(k_i)$ and the k_i are candidate passwords. Comparison of n_a with $z^{e/e(k_i)}$ for each i now reveals the correct k_i , and this can even be done in time for the false B to generate message (4) correctly and finish the protocol run. Variations are possible in message (2), for example using $3^{e(k)}$ as the added term or using the exclusive-or \oplus in place of modulo- N addition.

We turn now to a consideration of the constraints on the function $e(k)$. Suppose that e is a fixed prime, and consider a modified form of the RSA-S3P protocol where message (2) contains $z^e + k \bmod N$ in place of $z^{e(k)}$. If an attacker masquerading as A can choose $N = pq$ where p is a prime of the form $re + 1$, then the Euler totient $\phi(N) = re(q - 1)$ so for almost all values of z we have $(z^e)^{r(q-1)} = 1 \pmod{N}$. Exhaustive search following a single foiled active attack now reveals k .

This attack on the modified protocol can be blocked by making the requirement that $e > N$, which ensures that e is relatively prime to $\phi(N)$.

Suppose p_i are small odd primes and the attacker would like to find i with $p_i = e(k)$. Define $p = 1 + 2 \prod_{i \text{ even}} p_i$, $q = 1 + 2 \prod_{i \text{ odd}} p_i$ and arrange the indexing of p_i so that p and q are prime. Setting $N = pq$ we have (almost certainly) that $p_i | e(k)$ iff $(z^{e(k)})^{P/p_i} = 1 \pmod{N}$ where $P = 4 \prod_{\text{all } i} p_i$. Other number-theoretic attacks are considered by Patel [13].

To block attacks of this form it suffices to ensure that $e(k) > N$ for all k , since the bit-length of N is prescribed by the protocol and checked by B . Alternatively, we could just insist that all $e(k)$ have a one somewhere in the high order bits. This still allows an attacker to eliminate two values of k per active attack, but no more.

We also need to ensure that the mapping from k to $e(k)$ is, as nearly as possible, one-to-one. The fact that a and b are relatively prime and $ab > N$ ensures this. The density of primes below N implies that a prime will be found on average for i of order $\ln N$ and almost always for i very much less than \sqrt{N} . There are other subtle constraints upon the algorithm for e . For example, if the function $e(k)$ were instead defined to be the first prime after $a + b.k$, then the density of prime numbers implies that $e(k) = a + b.k + i$ for small i . The fact that i is typically bounded by a small multiple of $\ln N$ gives the attacker information about $e(k) \bmod b$. An attacker pretending to be A can choose N to contain a factor of the form $br + 1$, and hence deduce information about $z \bmod br + 1$ which can be used to recover z from message (2) with a greater probability of success than guessing k .

4 Diffie-Hellman Based Protocol

In this section 4 we consider various DH-based S3P protocols which have stronger security properties than the conventional versions, and which make less use of supercryptography. We also discuss some novel threats.

4.1 Basic DH-S3P Protocol Description

Let q be a publicly known large prime of prescribed length, and let g be a publicly known residue modulo q . To prevent various known subtle attacks [3,9,13] we assume that $p = (q - 1)/2$ is a prime and g is a generator modulo q , so that g^n has period $2p$. Note that in case $p \bmod 4 = 1$ we can take $g = 2$ [8, Theorem 95]. We assume that k has a smaller bit-length than q and is encoded in such a way that any two candidate values for k differ in at least three bit positions. The last requirement is for technical reasons which will be discussed later, but can be guaranteed by adding a small amount of redundancy to k .

A and B select strong random numbers x, y respectively. By strong we mean that exhaustive search is infeasible. The basic version of the DH-S3P protocol runs as follows:

$$A \rightarrow B : g^x + k \bmod q \tag{1}$$

$$B \rightarrow A : g^y \bmod q \mid n_b \tag{2}$$

$$A \rightarrow B : n_a \tag{3}$$

where we write $g^{2xy} \bmod q = z = c|s|n_a|n_b$ with semantics as in the RSA protocol.

Both A and B should check to ensure that $z \neq 0$ and $z \neq 1$. If $z = 0$ or $z = 1$ then the run fails, since otherwise an active attacker could masquerade as B by using these values in message (2).

The nonces n_a, n_b and the confounder c need not be strong, although they should be significantly harder to predict than k . We require that x, y strong and large relative to $\log_2 q$.

A must forget x and g^x , while B must forget y . Both A and B must forget c . A must forget the whole of z if the protocol run fails at message (2). The DH protocol ensures that s appears random, but does not allow it to be chosen or predicted by A or B .

4.2 Design Discussion

Whereas the RSA-based protocol required four messages, the Diffie-Hellman variant can be done in three, effectively by combining both texts uttered by B in the same message. Consequently, in marked contrast to the RSA case, the second message in the DH protocol contains verifiable redundancy in the form of n_b . The reason an attacker cannot use this to break the protocol is that the redundancy is only detectable by an entity who knows x or y . These are not searchable by hypothesis, and the value of x is not deducible from message (1) even with a guessed value for k .

As with RSA, eventual cracking of the chosen public key will give the attacker k and, worse, allow the attacker to obtain old values of s . For this reason the public parameter q must be many times longer than s , and so a large number of bits is again available for n_a, n_b and c .

The purpose of the Hamming-distance restriction alluded to earlier upon the encoding of k is to prevent the attacker testing multiple values of k in a single

run. If $g = 2$, then the attacker knows the discrete logarithms of small powers of 2, and can use this fact to test simultaneously a set of candidates for k , each of which differs in only one bit from some value k_0 . The attacker sends k_0 as the first message and inspects $g^{2^{ix}}$ for all i less than the bit-length of k , to see if one contains n_b . If all candidate values for k are at least three bits apart then this attack is defeated. For example including a Hamming code would add only 10 bits to a 1,000 bit k .

4.3 Choosing the Modulus

In this sub-section, we consider possible alternative approaches to the choice of q and g . This in turn leads to some variations on the DH-based S3P protocols.

To avoid narrowing attacks, we require that q be a prime of the form $2p + 1$ for some prime p , and that g be a primitive root modulo q . Such parameters are relatively expensive to generate, and in the DH protocol both A and B must check that q, g are suitable values, since using poor values can reveal k . In the RSA case only A need check. However, while the RSA protocol needs a new key N for each run, the DH protocol can use same parameters g and q many times. Consequently the parameters g and q could be relatively long-term and chosen by A and B jointly prior to the first run of the protocol, or else chosen, certified and published by some party, protocol or algorithm trusted for this purpose by both A and B .

Alternatively, A could choose the public parameters q, g and send them to B in the first message. B must carry out a deterministic test to verify that the parameters have the required properties. A deterministic test should be used, since many non-deterministic tests assume random rather than malicious choice of candidate primes. To enable B to carry out such a test efficiently A can send a witness along with the parameters. However, we still need q to be many times longer than s . If A chooses q, g each run then it may be more efficient to find a prime q of the more general form $q = rp + 1$ where p is a large prime and r is relatively small, since these primes can be sieved for more quickly, although it is then more difficult to find a generator g . The previous case corresponds to $r = 2$, whereas for this case $r = 2^n$ for a small n might be better. To avoid narrowing attacks when q is of this more general form, take $g^{rxy} \bmod q = z = c|s|n_a|n_b$ to force z into the large subgroup, and check $z \neq 1$ and $z \neq 0$. Another option is to replace $g^y \bmod q$ by $g^y + k \bmod q$ in message (2).

A further possibility is where a fresh q of the prescribed length is chosen in some way for each particular run of the protocol. For example, the value of q may depend in a deterministic way upon both k and an unpredictable random value r , so that $q = q(k, r)$, where r is produced during or just prior to the protocol run. This unpredictable value of r need not be kept secret, and may be published by a beacon, or agreed by A and B using some other protocol. In this case the protocol requires no superencypherment by k at all. The first message contains just $g^x \bmod q$ with the provision that trial values for x must be picked and discarded until one is found for which the high order bit of $g^x \bmod q$ is zero, and similarly for y . This ensures that the protocol run gives no information

about which q was used, and hence leaks no information about k . By forcing all $q(k, r)$ to have high-order bits $100\dots 0$ for some fixed number of zeros, we can make the probability of a high order one in $g^x \bmod q$ as small as desired.

4.4 Modified DH-S3P Protocol

We conclude this section by considering in more detail the case where the random number r is produced during the S3P protocol run itself. The simplest method is for A to send r in the first message along with $g^x \bmod q$. However the protocol which follows is designed to illustrate a more paranoid scenario. We assume that A and B wish to use part of the value of s to settle a bet [15]. Even if they have no doubt of one another’s honesty, they must be able to prove to a sceptical third party that neither of them has the capability to influence the value of the shared secret s in a predictable way. The random values such as x, y actually generated during the course of the protocol run must be destroyed, and so cannot subsequently form part of an audit trail.

A picks strong random numbers m, x and y' . B picks strong random numbers m', x' and y . In the protocol description which follows, $q = q(0|m, k)$, $q' = q(1|m', k)$ and g, g' are the corresponding generators. We assume that $q = r.p + 1$, $q' = r'.p' + 1$ for large primes p, p' . The functions `unzip0` and `unzip1` denote the even and odd-numbered bits respectively.

The DHm-S3P protocol runs as follows:

$$A \rightarrow B : m \mid g^x \bmod q \tag{1}$$

$$B \rightarrow A : m' \mid \text{unzip0}((g')^{x'} \bmod q' \mid g^y \bmod q) \tag{2}$$

$$A \rightarrow B : (g')^{y'} \bmod q' \tag{3}$$

$$B \rightarrow A : \text{unzip1}((g')^{x'} \bmod q' \mid g^y \bmod q) \mid n_b \tag{4}$$

$$A \rightarrow B : n_a \tag{5}$$

Here $z = c|s|n_a|n_b = (g^{rxy} \bmod q) \oplus (g^{r'x'y'} \bmod q')$ where \oplus denotes bitwise XOR. Effectively the DHm-S3P protocol runs two instances of the basic protocol back to back, but reveals information only about the exclusive-or of the two results. This means that an attacker must crack discrete log for both q and q' simultaneously, rather than searching log tables one at a time. The `unzip` functions force A to commit y' before learning g^{rxy} , but after B commits to y .

5 El Gamal Based Protocol

In this section we consider the implications of using EG in place of DH, when the requirement is to transfer an existing secret rather than to agree a new and unpredictable secret. The EG variation of the S3P protocol allows B to pick the session key and nonces, as was the case in the RSA protocol. The EG-S3P protocol runs as follows:

$$A \rightarrow B : g^x + k \bmod q \quad (1)$$

$$B \rightarrow A : g^y \bmod q \mid z \cdot g^{2xy} \bmod q \mid n_b \quad (2)$$

$$A \rightarrow B : n_a \quad (3)$$

where $z = h|s|n_a|n_b$ as for the DH protocol, except that instead of the confounder c , z contains a known fixed bit pattern h chosen so that it is not invariant under shifts or subtraction from q . The constraints on h will be further discussed below. The password k is encoded as in the DH protocol. As in the Diffie-Hellman case, A and B should check that $g^{2xy} \notin \{0, 1\}$. A should also check that z contains the expected value for h .

A must forget x and g^x , while B must forget y . A must forget the calculated value of g^{2xy} and z if the protocol run fails at message (2). Apart from this, h need not be kept secret. The EG protocol allows s as well as n_a, n_b to be chosen by B so as to contain redundancy or known text.

The El Gamal variant shares some features with the RSA case and some with the DH case. As with DH, it is a three-message protocol and the middle message must contain redundancy. In the RSA protocol the value of s is chosen by B but s must contain no redundancy discernible to the attacker: otherwise k is in danger. The DH protocol ensures that s appears random, but does not allow it to be chosen or predicted by the participants. The EG protocol allows B to choose s , and for s to contain redundancy in any form desired. Indeed for the EG protocol even prior knowledge of s by the attacker does not assist in an active attack against k . Also, in the EG protocol n_a and n_b may contain redundancy or known text. As with the other protocols, the nonces n_a and n_b need not be strong, although they should be significantly harder to predict than k . However in the EG protocol h is not a confounder at all. Instead, it contains redundancy to prevent a person in the middle modifying messages (2) and (3) in such a way that the protocol appears to complete successfully, but with A and B disagreeing on s . For example an attacker can multiply z by two in the second message and shift n_b and n_a one bit to the left and right respectively, with a 50% chance of escaping detection. The result is that s is shifted left one bit in transmission. Division and complementation are also possible. Such attacks can be prevented by placing a fixed bit pattern in h , for example a single 1-bit with n 0-bits on either side will suffice, provided 2^{-n} is small relative to the chance of guessing k .

The EG protocol, like the DH protocol, can use the same parameters g and q many times. As in the DH protocol, both A and B must check that q, g are suitable values, since using poor values will reveal k . As with RSA and DH, eventual cracking of the chosen public key will give the attacker k and, worse, allow the attacker to obtain old values of s . For this reason the public key must be many times longer than s , and so a large number of bits is available for n_a, n_b, h .

An alternative approach (which we do not pursue here) is to obtain n_a, n_b from g^{2xy} as in the DH protocol, rather than from z .

6 System-Level Considerations

In this section, we discuss what should be done when an active attack is detected, and introduce the notion of a robust protocol wrapper. We also discuss the system context for the deployment of S3P protocols, and the hardware support required.

6.1 Action Following a Detected Attack

An important feature of all the S3P protocols we consider is that it is not acceptable to ignore an active attack. If active attacks are ignored, the attacker can make one active attack for each possible k , and will eventually succeed. If suitable emergency action is taken in the event of an active attack being detected (eg switching to a more expensive but physically secure channel, or to another, previously agreed, password, after a certain number of failed runs), then the attacker never gets enough information to improve his chances of guessing correctly by more than some previously agreed security parameter.

In an extreme case we can confine the attacker to two guesses, one with each of A and B . In a less extreme case, with (say) a million equally likely values for k , we could choose to allow the attacker 32 guesses with each of A and B . The attacker has less than a one in ten thousand chance of obtaining the true value of k . Of course, this strategy requires some assumptions about the physical locations of A and B , and their ability to remember the number of active attacks detected over an appropriate time scale such as the expected life of the long term password k . We also need to specify, in any particular system context, how these numbers are stored and whether they are secret.

In particular, if the protocol is used by many pairs of participants, then an attacker can make a small number of attacks against each of a very large number of passwords, and will almost certainly succeed against one. The effects of this form of penetration, and the countermeasures for containing it, depend upon the interactions between the system-level protocols for which the strong shared secrets are used.

An attractive alternative to using a deterministic counter and a threshold is to invoke emergency action with a certain constant probability after each detected attack. For example, if we set this probability at 2% then the alarm will almost certainly be raised after 70 detected attacks, regardless of who detects them. Since we assume that all parties who use the S3P protocols are able to generate good random numbers, this stochastic technique imposes no new system constraints.

6.2 Robust Protocol Wrappers

The primary system context which we consider for the S3P protocol is one of paranoia rather than hostility. In other words, we assume that the world is full of very clever and hardworking attackers, but at the same time we are confident that things will go right most of the time. In effect, we assume that the S3P

protocol is nested inside another protocol, which we call the wrapper, and that the outer wrapping protocol works nearly all the time unless there really is an active attack by an extraordinarily malicious and ingenious entity. The inner S3P protocol is intended both as a trip-wire to indicate whether the outer wrapper has been deliberately breached, and as a last-ditch defence.

A primary purpose of the outer protocol wrapper is to ensure that the inner S3P protocol is under no accidental misapprehension about whether an offered bit pattern represents an attempt to engage in the S3P protocol, and if so in which run, at what stage, and as whom. A similar two-layer scheme for distinguishing accident from malice was used by Lomas and Christianson [11] and a related notion of robustness is discussed by Anderson and Needham [2]. The S3P run must fail if any such presented bit pattern is incorrect, otherwise the enemy gets a free guess.

We also assume that “eventually” a run of the protocol which does not proceed will be regarded as having failed by at least one of the participants. However this timeout may be very long. One reason for this is that we do not wish to have too many false alarms, but there is another reason. We wish also to allow a system context in which a run of the S3P protocol is transported by a slow non-cryptographic outer protocol such as fax, snail-mail, or sneakernet. This gives rise to two further considerations: re-entrancy and interactive breaking.

If an S3P protocol run can take a long elapsed time, then the S3P protocol must be re-entrant. The total number of active runs (plus the number of previously detected failures) must be less than the threshold value for the number of active attacks which we are prepared to tolerate. This ensures that all runs which terminate successfully are safe. In particular, runs can be pipelined or used back-to-back between the same two parties on tamper-proof hardware such as smart cards which are kept locked up when not in use.

The possibility of a long elapsed time also provides one motivation for our consideration of the possibility that a value of s could be broken between steps of the S3P run, for example if it were used to encrypt a known plain text as part of the S3P protocol. In any protocol involving key agreement, it is possible to specify the agreement of a much longer shared key than required, and to forget all but the required number of bits. Many applications of other published protocols would benefit from doing this.

Just as we do not assume s to be strong, neither do we *require* k to be weak: although the S3P protocols were originally designed to work with passwords k drawn from a space of order 2^{20} possibilities, the protocols also have particularly nice properties when a 40-bit shared key k is being traded up to a series of 120-bit keys s_i .

6.3 Tamper-Proof Hardware Platform

One possible system context for S3P is where we wish to ensure that the right person is using a particular box. The box may be designed to be used by several different people (eg a workstation in a shared area) or by only one person (eg a mobile telephone or a hand-held authenticator). The box may be stateful (eg

able to retain session keys) or stateless (all mutable information is deliberately erased between uses). However a box may be stolen or tampered with. We wish to ensure that the box can only be used by a person who knows the correct password.

To deploy the S3P protocol we assume that the box is tamper-evident, and unviable to forge. We assume that the user checks the tamper-evident seal before entering the password at the start of each run. We assume that the box forgets the password once the S3P protocol run ends or fails, and that while the run is in progress the box is tamper-proof, in the weak sense that that the box will irrevocably destroy (forget) secrets rather than allow them to be read. This property might require that the box is used in a different environment from the one in which it is stored between runs.

Under these assumptions, the S3P protocol suffices to ensure that the box cannot be used by a person who does not know the password. In the case where the box may be used by more than one person, each user may have a different password. Note that tamper-proofing is not required except while the protocol is running. Tamper-evidence suffices the rest of the time even in the stateful case. State which persists between runs can therefore be used to support the outer wrapping protocol, so long as the state of the box between runs can reveal no information about the password. However the inner protocol must not rely upon the outer protocol preserving state correctly. This point will be illustrated at the end of the section by discussion of a reflection attack.

6.4 Blocking Reflection Attacks

We conclude this section with a brief consideration of how to block a reflection attack. The inner S3P protocol is assumed to be stateless, and so a new run cannot reliably determine which other protocol runs are still active when it begins. Suppose that *A* attempts to run the protocol with *B*. The attacker takes the first message from *A* and replays it to *A* as if it came from *B* initiating a different run of the protocol. *A*'s reply to this is in turn reflected as a reply to *A*'s initial message, and so on for the subsequent messages. If *A* is not careful, she will end up sharing a fresh strong secret with herself, rather than with *B*, in violation of our requirement that the other intended participant must actually be involved in any apparently successful run.

Of course, this attack cannot actually succeed against the protocols as we have described them here, since *A* always sends the odd-numbered messages and *B* the even. But suppose we wish to allow either party to initiate the protocol, possibly re-entrantly, so that *A* may legitimately use *k* to speak the lines attributed to *B* in the script.

We can block the reflection attack by associating the nonces firmly with principals. In respect of each password, one party is (by mutual agreement) the *a*-end and the other is the *b*-end. Suppose that *Carol* is the *a*-end and that *Ted* is the *b*-end. Then whenever *Carol* has to place a nonce in a message she always uses n_a , regardless of whether she is playing the part of *A* or of *B*.

7 Discussion

The protocols given in this paper are provocatively weak. For example, we use bit selection in place of a hash function, modulo addition to perform superencypherment, and a base of 2 for certain exponentiations.

This weakness is quite deliberate. From a practical point of view, our protocols could doubtless be strengthened by the judicious inclusion of “one-way” hash functions, or the use of more complex forms of superencypherment and convolution. We have instead put forward very concrete versions of the protocols, with primitives which rely upon specific number-theoretic relations between, for example, modular exponentiation, and addition or “unzip”. Following Ockham, we wish to understand these simple protocols before we propose anything more complicated.

We have not provided correctness proofs for these protocols here. This is an area in which we anticipate future progress. The major present difficulty lies in determining precisely how the threat model interacts the desired properties of the protocol with those of the underlying cryptoalgorithm. For example, one standard *reductio* approach might be to prove that, if the RSA-protocol reveals more information than it should about k , then it also gives an attacker the ability to decrypt unpredictable RSA cyphertexts. However attacks in the spirit of Bleichenbacher show that such an outcome need not constitute a break of the RSA cryptoalgorithm, and hence ought not simply to be presumed counterfactual. Conversely, a protocol continues to satisfy the assertion of a correctly proved predicate, even after the protocol has been broken by another means.

On the positive side, cryptographic innovations which we claim for this paper include the successful use of RSA as a vehicle for encrypted key exchange, section 4.4. As a minor point, we also draw attention to the lack of superencypherment by k in the second message of the basic version of our DH protocol. However, we regard as primary our original contribution to consideration of the system context given in section 6, including the introduction of robust protocol wrappers and their application to the case of “stateless” platforms.

The protocols in this paper owe an obvious debt to the original discussion by Bellovin and Merritt [3], which opened up a number of fertile research directions. We make no attempt to give a systematic account of all this related work here. (An excellent roadmap is provided by Jablon’s website at <http://www.IntegritySciences.com>.)

Some of the material in this paper appeared in preliminary form in University of Cambridge Computer Laboratory Technical Report 458 (1998). We would like to thank Daniel Bleichenbacher, David Jablon, David Wagner and everyone else who provided attacks and related comments on these early versions of the protocols.

8 Conclusions

Although the primary purpose of the S3P protocol is to share strong secrets, the design of the protocol does not assume that s is strong. The S3P protocol can also

be used simply to allow a remote authentication service to authenticate a user to a “stateless” host which is local to the user. In this case s may be an authenticator for the audit trail. In our design we make no restrictions upon what the shared secret s is used for once the S3P protocol run has ended: s may be revealed, used as a one-time pad, a cryptographic key, as a salt or an initial value. At a slightly more general level, we remark that it appears very difficult abstractly to model security protocols in a formal way that takes adequate account both of the cryptographic properties assumed, and of the security service provided. Protocols may legitimately be used in ways not explicitly considered by their designers, and the safety of the resulting applications can depend in an unknown way upon the safety of obscure number-theoretic hostages which were abstracted away in the construction of the threat model.

Our S3P protocols make no use of hash functions or symmetric cryptography. However our protocols rely completely for their properties upon the security of the public key systems used. Consequently it is necessary for the moduli to be uncrackable for at least the lifetime of all secrets (weak or strong) used or agreed with that modulus. This provides a sufficient number of bits to provide a strong secret and a strong confounder, together with two nonces. In contrast with the confounder, the nonces can be searchable, so long as the most likely nonce is less likely than some system threshold parameter. Again, we remark at a general level the need to balance the bit-budget carefully when tuning the performance of security protocols which use public key cryptography.

Our S3P protocols also rely upon the ability of those using them to generate irreproducible random bit patterns, and to delete information irrecoverably. These are both interesting technical challenges. In particular the task of finding a suitable source of randomization, upon which (in the context of a particular system) it would be impractical to eavesdrop is one which would repay further study. As a final remark of a general nature, we stress the importance of explicit consideration, when specifying the threat model, not only of the hardware platform supporting the security protocol, but also of the system context, and the security policy under which the hardware will be configured.

28 January 1998; revised July 1998, January & November 1999, June 2000, October 2004. Contact: b.christianson@herts.ac.uk

References

1. Anderson, R., Lomas, M., 1994, Fortifying Key negotiation Schemes with Poorly Chosen Passwords, *Electronics Letters*, 30(13) 1040–1041.
2. Anderson, R., Needham, R., 1998, *Programming Satan’s Computer*, Springer LNCS 1000.
3. Bellare, S.M., Merritt, M., 1992, Encrypted Key Exchange: Password-Based Protocols Secure Against Dictionary Attacks, *Proc IEEE Computer Society Symposium on Research in Security and Privacy*, Oakland 92, 72–84.
4. Diffie, W., Hellman, M., 1976, New Directions in Cryptography, *IEEE Transactions on Information Theory*, 22(6) 644–654.

5. ElGamal, T., 1985, A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms, *IEEE Transactions on Information Theory*, 31(4) 469–472.
6. Gong, L., 1995, Optimal Authentication Protocols Resistant to Password Guessing Attacks, *Proc 8th IEEE Computer Security Foundations Workshop*, 24–29.
7. Gong, L., Lomas, M., Needham, R., Salzer, J., 1993, Protecting Poorly Chosen Secrets from Guessing Attacks, *IEEE Journal on Selected Areas in Communications*, 11(5) 648–656.
8. Hardy, G.H., Wright, E.M., 1978, *An Introduction to the Theory of Numbers*, 5th edition, Oxford University Press.
9. Jablon, D.P., 1996, Strong Password-Only Authenticated Key Exchange, *Computer Communications Review*, 26(5) 5–26.
10. Kelsey, J., Schneier, B., Wagner, D., 1998, Protocol Interactions and the Chosen Protocol Attack, *Security Protocols 5*, Springer LNCS 1361, 91–104.
11. Lomas, M., Christianson, B., 1995, To Whom am I Speaking? Remote Booting in a Hostile World, *IEEE Computer*, 28(1) 50–54.
12. Lucks, S., 1998, Open Key Exchange: How to Defeat Dictionary Attacks Without Encrypting Public Keys, *Security Protocols 5*, Springer LNCS 1361, 79–90.
13. Patel, S., 1997, Number Theoretic Attacks on Secure Password Schemes, *Proc IEEE Computer Society Symposium on Research in Security and Privacy*, Oakland 97, 236–247.
14. Rivest, R., Shamir, A., Adleman, L., 1978, A Method for Obtaining Digital Signatures and Public Key Cryptosystems, *Communications of the ACM*, 21(2) 120–126. *Operating Systems Review*, 29(3) 22–30.
15. Wheeler, D., 1997, Transactions using Bets, *Security Protocols 4*, Springer LNCS 1189, 89–92.