Separating two roles of hashing in one-way message authentication

L.H. Nguyen and A.W. Roscoe

Oxford University Computing Laboratory
{Long.Nguyen/Bill.Roscoe@comlab.ox.ac.uk}

Abstract. We analyse two new and related families of one-way authentication protocols, where a party wants to authenticate its public information to another. In the first, the objective is to do without shared passwords or a PKI, making use of low-bandwidth empirical/authentic channels where messages cannot be faked or modified. The analysis of these leads to a new security principle, termed *separation of security concerns*, under which protocols should be designed to tackle one-shot attacks and combinatorial search separately. This also leads us develop a new class of protocols for the case such as PKI where a relatively expensive signature mechanism exists. We demonstrate as part of this work that a popular protocol in the area, termed MANA I, neither optimises human effort nor offers as much security as had previously been believed. We offer a number of improved versions for MANA I that provides more security for half the empirical work, using a more general empirical channel.

1 Introduction

We examine some protocols which attempt to transmit a (possibly very long) message from one party to another in such a way that the origin and integrity of the message are authenticated. Initially we set out to do this with just one-way communication and authentication strings without the presence of any initial security infrastructure. This illustrates the power of authenticated empirical channels that are authentic, unspoofable or unfakable, but on the other hands, can be overheard by anyone.

To set this work in context, recall the classic one-way authentication protocol which works where there is a PKI. Here a message M and the name of the sender A are accompanied by the digital signature, or *message authentication code*[1] (MAC) $\{hash(M)\}_{sk(A)}$ and possibly the public-key

[1] In this paper we investigate a variety of techniques for providing authentication and integrity evidence for a message. We are inclined, therefore, to use the name "MAC" for a rather wide class of such techniques including ones based on asymmetric cryptography and further concepts that we will discover, rather than just referring, for example, to cases where the participants share a symmetric key.

certificate of the sender A. The receiver knows M really is from A because he can form the cryptographic hash of M and discover if it really was A who signed this value with her secret key.

Although the whole of such a message may be assumed to be sent over a standard Dolev-Yao channel, there is in fact a closer tie-in with the subject matter of this paper than there might appear to be. For public key encryption and decryption are computationally expensive: this means that there is a strong incentive to keep the bandwidth of information transmitted under this form of cipher to a minimum. We might therefore regard the single message described above as the combination of a (perhaps large) message M over an insecure channel with the smaller one $hash(M)$ over an authenticated one.

Since in many cases the empirical channels we will be using are human mediated, the chief difference from the above analysis is be that our empirical channels are much lower bandwidth yet: the amount of security delivered for a given amount of empirical communication becomes the most important measure of a protocol's effectiveness.

We start by describing a number of non-interactive one-way authentication schemes that use empirical channels in different ways, for example: protocols of Pasini and Vaudenay [11], and Balfanz [1]. These require the transmission of more bits over empirical channels than desirable.

Subsequently, we observe the development of MANA I by Gehrmann, Mitchell and Nyberg [3] that can reduce the number of bits required to be transmitted over the empirical channel significantly. This was the first example in the literature where useful combinatorial search was largely prevented. However, we will see that the scheme neither optimises human effort nor offers as much security as has previously been believed. We offer a number of improved versions for MANA I that provides more security for half the empirical work, using a more general empirical channel.

Having done this work in the context of "authentic" channels, we are able to formulate a general principle from it: the principle of *separation of security concerns*, under which protocols should be designed to resist a one-shot attack and combinatorial search separately. In turn, this principle allows us to devise a new family of protocols that work in the context of a PKI or similar. These schemes will, we believe, often have efficiency advantages over the conventional signature described above thanks to the use of short output *digest* functions that can be computed more efficiently than cryptographic hashes.

1.1 Notation

Capital letters such as A and B are used to identify parties. In common with much of the literature we are citing, the combination of two pieces of data will frequently be written $x \parallel y$ or the ordered pair (x, y). We will assume node A has some public information M that it wants to have authenticated to node B, this might include name/addressing information, its uncertificated public key or Diffie-Hellman token. The notations for several types of channel are given below. These are taken from [8].

- $\longrightarrow_N$, all messages transmitted over the Dolev-Yao network can be overheard, deleted or modified by the intruder.
- $\longrightarrow_{WE}$, this *weak empirical* channel cannot be forged, but it can be blocked, overheard, delayed or replayed [14, 11].
- $\longrightarrow_E$, this is like a weak empirical channel except that it cannot be replayed. It can be delayed, but not sufficiently long so that a message from one session can be used in another [6, 7].[2]
- $\longrightarrow_{SE}$, this is similar to a normal empirical channel, but it also provides stall-free transmission, and cannot be delayed, removed or blocked by the intruder. This is termed a *strong empirical channel* [14, 11].
- $\longrightarrow^t_{BE}$ is the same as $\longrightarrow_{SE}$ except that messages cannot take more than time t to arrive. In other words, no empirical message can be accepted more than t time units after it was sent. We will call this a *bounded delay empirical channel.*

Some of the proposed mechanisms such as Authentication without hashing presented in Section 5 and the Improved MANA I make use of a new cryptographic primitive termed a *digest* function. While the desired properties of a digest function are similar to those required of cryptographic hash functions, universal hash functions, and the MAC or check function of [3], it has a less challenging specification than a hash, and is frequently intended to be short output (perhaps 16 to 32 bits). We have previously described digest functions in [6, 7]. The uses we make of them in Sections 5.1 and 5.2 of the present paper are qualitively different, however, since (i) it is no longer necessary that the output of the function is short, and (ii) the only motivation for using them is that they are faster to compute than a hash.

[2] In practice, it is not hard to avoid messages transmitted over the empirical channel being delayed from one to a latter session because it is normally the humans who run the channel. For example, if the humans involved are not away at any time during a protocol run, then empirical messages cannot be delayed from one to another session.

The following is the specification of a short-output, b-bit digest. The specification for a general digest is the same except that 2^{-b} is replaced by some small ϵ and collision w.r.t different keys is also taken into consideration.

As the key k varies uniformly over its range:

1. $digest(k, M)$ is uniformly distributed for any fixed M.
2. For any fixed θ and $M \neq M'$: $\mathbf{Pr}(digest(k, M) = digest(k, M')) \leq 2^{-b}$

The computational complexity of digest functions is crucial for Section 5, as indicated above. The following complexity model is taken from [6–8]. It is clear that the cost of computing the b-bit output $hash_b(M)$ or $digest(k, M)$ increases linearly with the length of M. It also seems clear that it will increase significantly with b, and a simple model in which each word of a running temporary value of length b is combined with each input word suggests our overall model might be $b \times length(M)$. Since well-known hash algorithms tend to be fixed width, and vary significantly in their individual costs, it is hard to be too definite about this rule, although the nature of the individual algorithms tends to support our hypothesis. We will discuss this issue further in Section 6.

We now summarise an idealised framework for the digest function, proposed in [6, 7]. This has been formally proved to satisfy the above specification exactly. In practice (as opposed to *idealised*) the random numbers required by this scheme would be simulated by a pseudo-random number generator.

Suppose we want to construct a b-bit digest of a $(K-1)$-bit message M. The first thing we do is to pad M with an extra 1-bit at the end, so its length becomes K with $M_K = 1$. For $i = 1, \ldots, b$ and $j = 1, \ldots, K$, suppose $R_{i,j}$ are independent uniform boolean-valued random variables whose values are derived from k.

Using matrix product, we define $digest(k, M) = M \odot R$ where the symbol $\odot$ represents the binary product of the vector M and the matrix R. Instead of deriving a completely random matrix R from the key k, a Toeplitz matrix – where $R_{i,j} = R_{(i+1),(j+1)}$ for all values of i and j, in other words it is constant on any diagonal – can be used to reduce the required number of random bits from $K \times b$ to only $K + b - 1$ without loss of security as suggested in [7, 5].

2 Long authentication string protocols

The analysis of the use of digital-signature MACs in the introduction shows they are closely analogous to the following protocol, devised by

Balfanz [1]. In this scheme, A wants to authenticate its information M to B. Here $hash_{160}()$ denotes a 160-bit output cryptographic hash function.

Balfanz non-interactive protocol, [1]
1. $A \longrightarrow_N \quad B : A, M$
2. $A \longrightarrow_{WE} B : hash_{160}(A, M)$

The 160-bit hash sent over the weak empirical channel can be delayed and the information M is under the control of the intruder, hence s/he might carry out an off-line attack to find a different M' with the same hash value.[3] That is something which the standard specification of a hash function deems infeasible as it takes about $2^{160/2} = 2^{80}$ hash calculations on average to find such a collision using the birthday paradox.

In order to improve the number of authenticated bits, Pasini and Vaudenay [11] make use of a probabilistic commitment scheme[4] that is at least as secure as a standard cryptographic hash function to commit to the authenticated information. The 80-bit hash of the commitment is then sent over the weak empirical channel. Here the hash function is required to be weakly collision resistant (i.e. the second preimage resistance property: an intruder cannot find a second value v' such that $hash(v) = hash(v')$ for fixed v).

Pasini-Vaudenay non-interactive protocol, [11]
1. $A \longrightarrow_N \quad B : c \parallel d = commit(A, M)$ $\qquad\qquad B$ computes $A \parallel M = open(c, d)$
2. $A \longrightarrow_{WE} B : hash_{80}(c)$

In [11], Pasini and Vaudenay argue that this provides the same degree of authentication as the Balfanz protocol (namely 2^{80} hash computations) while halving the number of empirical bits thanks to the probabilistic commitment scheme that avoids the possibility of a birthday attack. However, it seems fair to remark that, even 80 bits will seem tedious for most humans to compare carefully in practice.

[3] In the original protocol [1], there is no restriction on the order of sending and receiving Messages 1 and 2.

[4] The commitment scheme used in Pasini-Vaudenay [11] consists of two functions. $c \parallel d = commit(A, M)$ and $A \parallel M = open(c \parallel d)$. A intends to bind a fresh long random nonce R_A and M together without revealing R_A by publishing the commitment c. Eventually sending d (the *decommitment*) reveals R_A, and binds this value firmly to M in the eyes of the receiver. As R_A is a long random nonce the security of the scheme in term of both *binding* and *hiding* is equivalent to a standard cryptographic hash function.

2.1 Objectives in designing authentication protocols

When designing an authentication protocol, particularly one based on hash functions such as the two above, we typically need (inter alia) to meet the following pair of objectives:

A Combinatorial attacks that involve searching for hash collisions etc are made too difficult to carry out with any reasonable hope of success.

B Whatever guess-work or strategy the attacker can carry out (perhaps involving **A**), his chances of success are sufficiently low.

In traditional uses of hashes, these two are inextricably linked, and indeed we would normally characterise the required strength of a hash function as being what is required to overcome both of these simultaneously (and this is the case in both Balfanz and Pasini-Vaudenay).

For example, it is often reasonable (and this is the case with Balfanz) to assume that an attacker can carry out a birthday-style attack, in which the expected number of collisions he can find with N hash calculations in a hash space of size H is $\frac{N \times N}{C \times H}$ for some positive constant C that takes into account both the nature of the search and assumed imperfections in the hash function. It follows that in order to keep the probability of success less than $1/T$, it is necessary (to a close approximation) to make the hash space greater than N^2T/C in size. Notice here that the parameter N comes from **A** above and that T comes from **B**, and the way in which H varies with the two of them is different.

Note that this demonstrates that if we have two uses of a hash function, one of which is vulnerable to birthday attacks and the other only to a plain search (as is achieved in Pasini-Vaudenay), it is not actually true that the two protocols give the same degree of security when a hash function of half length is used for the second, as is perhaps implied by the respective lengths quoted above by Balfanz and Pasini-Vaudenay. The lengths of the hash actually required are $(2\log N + \log T - logC)$ and $(\log N + \log T - logC')$, here C and C' can be the same or different to each other.

3 Short authentication string protocols

Gehrmann, Mitchell, and Nyberg [3] took a different approach to preventing combinatorial search. They use empirical channels to transmit the b-bit output of a check function $m_k()$ together with a b-bit key that has been instrumental in its computation.

MANA I (Gehrmann, Mitchell and Nyberg), [3]
1*a*. $A \longrightarrow_N B : A, M$
1*b*. $B \longrightarrow_E A :$ 1-bit committed
A picks a b-bit random number k
2. $A \longrightarrow_E B : k, m_k(A \parallel M)$

To eliminate the 1-bit empirical signal in MANA I,[5] Vaudenay proposes to use a strong empirical channel ($\longrightarrow_{SE}$), which achieve stall-free or instant delivery, to send the key and the check-value.[6] Thus $2b$ bits are transmitted in all. In the following description, we will modify the scheme slightly by using a digest function to compute the check-value. The rest of this analysis applies to both versions.

V-MANA I, [14, 11]
1. $A \longrightarrow_N B : A, M$
A picks a b-bit random number k
2. $A \longrightarrow_{SE} B : k, digest(k, A \parallel M)$

The use of the strong empirical channel that provides stall-free transmission leads to a significant fewer number of authenticated bits transmitted from A to B: these are the first example we have seen of protocols that, given the properties we have assumed of the digest function, come close to preventing the intruder performing any useful combinatorial search. This is because the distribution properties of the digest mean that it is impossible for the intruder to look for a M' that will digest to the same value as M in ignorance of k.

However, the protocol is far from optimal in the human work since any one can modify M blindly in the 1st message transmitted over the insecure normal network, and hope that the digests come out the same in the 2nd message. This will occur with a probability of 2^{-b} irrespective of the value of the key provided that the b-bit digest meets the specification defined in Section 1.1. What this means is that $2b$ empirical bits only guarantee at best a 2^{-b} security level.

Whilst the security proofs presented in [3, 11] are largely correct, what these authors have not discovered is that the bit-length b they choose for the key is too short compared to the digest output and the authenticated information: it is impossible to construct a digest function such that the

[5] In the original description of MANA I, the pair of parties additionally need to agree on the success of the protocol with the help of human. Since this is irrelevant to our security analysis, we ignore the step in our description of the protocol.

[6] We can replace $\longrightarrow_{SE}$ with a bounded delay empirical one ($\longrightarrow^t_{BE}$) provided B checks that he has received Message 1 before Message 2 could have been sent.

probability of any one-short attack is no better than 2^{-b}.[7] In fact there is a known theoretical bound on the bit-length of the key [13] that can guarantee the digest meets its specification: $bitlength(k) \geq bitlength(M) - b$.[8]

This result suggests we should aim always to have k noticeably longer than the digest in this style of protocol. Of course to do this without ruining efficiency in human effort, we need to find ways of communicating k over $\longrightarrow_N$ rather than empirically.

4 Improvements to (V-)MANA I

Given two weaknesses discussed in the previous section, we will present improved versions of V-MANA I that optimise the use of the expensive strong empirical channel. These improvements can also apply to MANA I. In other words, human comparison/handling of a b-bit short authentication string (SAS) always corresponds to a probability of 2^{-b} of a successful one-shot attack. Whilst this can only be done at the expense of introducing another (third) message sent over the Dolev-Yao channel we argue that this is not at all a bad trade-off since our highest priority is to minimise the empirical cost.

In contrast to V-MANA I, the key k generated by A in the following protocol can be as long as we want to ensure that the digest function meets the specification in Section 1.1. In addition, we can weaken the assumption that empirical messages' transmission is instantaneous to being of bounded delay as follows.

Improved version of V-MANA I (direct binding) *New*
1. $A \longrightarrow_N \ B : M, hash(k)$
2. $A \longrightarrow_{BE}^{t} B : digest(k, M)$
3. $A \longrightarrow_N \ B : k$

Note that the message order here and in other improved schemes of V-MANA I is more important than in all preceding protocols: we specify that B will not accept Message 2 within t of receiving Message 1 and that A will not send Message 3 within t of sending Message 2. This is to ensure that B was committed to Message 1 when Message 2 was sent, and that Message 3 cannot be received by anyone before B has accepted (if he does) the only Message 2 that A will ever send that relates to it.

[7] We will give a detailed analysis of the (off-line) computation complexity and its related probability of a successful one-shot attack on this protocol in Appendix B.

[8] We should remark that the bound can be met except for an infinitesimal tolerance for very much smaller lengths than this [9]. However, we suspect that it will be good practice for it to be significantly longer than b.

Furthermore, we can replace the bounded delay empirical channel and the need to wait by a simple acknowledgement from B to A. The resulting protocol is actually the pairwise version of HCBK protocol [12].

Improved version of MANA I (direct binding) [12, 6, 7]
$1a.\ A \longrightarrow_N B : M, hash(k)$
$1b.\ B \longrightarrow_E A :$ 1-bit commitment
$2.\quad A \longrightarrow_E B : digest(k, M)$
$3.\quad A \longrightarrow_N B : k$

We note that this scheme is flexible since the digest and key (Messages 2 and 3) can be released in any order as long as A has received the commitment signal from B in the 1^{st} message. It will often be the case that a bounded delay empirical channel and a one-bit acknowledgement signal are alternatives in this style of protocol design/structure.

Since the SAS in these schemes are functionally dependent on the authentic information M, we term these as the *direct binding* version of Improved (V-)MANA I.

Readers who are interested in the formal security proof as well as variants using indirect binding and Diffie-Hellman can find them in Appendix A.

5 Separation of security concerns

Protocols such as our improved versions of MANA I as well as HCBK [12, 6, 7] only work because it has been possible to separate the two concerns or objectives **A** and **B** as mentioned in Section 2.1. Specifically, these protocols avoid combinatorial attack by pre-committing participants to non-deterministic values such as the keys k, and keep the probability of a one-shot attack working low by choosing a good digest method and a short string of sufficient length.

In these protocols it was *necessary* that we separated these concerns, because it was unreasonable to expect humans to transmit or compare a value as complex as a normal cryptographic hash accurately (or in good temper!). It is interesting to note, however, that it brings a quite unexpected benefit: of the various calculations performed by the participants in the direct binding version of Improved MANA I or HCBK, only the calculation of the short string or digest actually depends on the message M being transmitted. It is reasonable to expect that, because the objective of this calculation is only to meet goal **B** rather than both this and what will almost always be the harder one **A**, it can be done more cheaply as a function of the length of M. A substantial gain is reflected in the complexity model described in Section 1.1.

Since the cost of this calculation is the only one that rises (almost certainly linearly) with the length of M, and all other aspects are constant, we come to the following striking conclusion: *when M is large, protocols based on the computation of a short digest can be more efficient than a traditional message signature scheme or MAC based on a standard cryptographic hash of the whole of M.*

This leads us to propose the following principle:

- **Separation of Security Concerns**: where a single cryptographic primitive is being used to satisfy several different security goals, one should consider if efficiency gains can be made by meeting these goals separately. This particularly applies if the primitive is being applied to a large block of data.

A good illustration of this is the way the objectives of message transmission and authentication are met separately in the useful and popular structure: $A \longrightarrow B : \langle M, \text{MAC}(A, M)\rangle$, or $\langle M, sign_A(hash(M))\rangle$. However, for largish M (approximately 10K bytes in our experiments based on SHA-1 and 1024 bit RSA) the time for hashing overtakes the time taken for the signature, and, for much larger messages than this, will dominate.

A particular consequence of the above principle derived from protocols such as our Improved (V-)MANA 1 is the following:

- **Factorisation of cryptographic hashing**: where a cryptographic hash function is being applied to a substantial item of data, analyse whether its security goals can be achieved more cheaply via a combination of a digest function to limit the chances of a one-shot attack, and some constant-time supplementary operations that limit the chances of an attacker to a single try.

5.1 Authentication without hashing

Consider the following protocols as an alternative to the conventional method of authenticating messages with a MAC of the section above.

In the first, A can compute $digest(k, M)$ simultaneously with sending Message 1, but only sends this value to B once the latter has signalled that it is completely committed to the value M by sending a nonce.

Authentication without hashing I (interactive) *New*
1. $A \longrightarrow_N B : M$
2. $B \longrightarrow_N A : N_B$
3. $A \longrightarrow_N B : sign_A(k, digest(k, M), N_B)$

Provided that B has not sent Message 2 until it knows (and is therefore committed to) M, it knows that A has not revealed the hash key k to anyone before that point, as Message 3 depends on N_B. Note that N_B communicated over Dolev-Yao channel is playing the same role of the 1-bit acknowledgement over empirical channel in our Improved MANA I.

In our second protocol, the role of the nonce N_B is replaced by a time stamp ts whose role is to prove that k was not revealed until B was committed to M. A must therefore wait a suitable period between completing Message 1 and sending Message 2.

Authentication without hashing II (non-interactive) *New*
1. $A \longrightarrow_N B : M$
2. $A \longrightarrow_N B : sign_A(k, digest(k, M), ts)$

In this scheme, B cannot accept the protocol run unless receipt of M was complete by time ts, which resembles to the requirement of the bounded delay empirical channel. Notice that this version is suitable for broadcasting a message to many B's simultaneously, but cannot (unlike a traditional digital signature) be used over and over again at different times. This is because the use of the same digest key at different times will allow an intruder to do a combinatorial search for a second M' such that $digest(k, M) = digest(k', M')$, and then deploy this against later recipients of the signature. We therefore will sometimes refer to this family of protocols as *one-time message authentication.*

A further disadvantage of the above protocols is that they do not permit the recipient to begin digesting until after the key k has been received. We believe, however, that they give both parties a significant reduction in processing time over an ordinary cryptographic hash function.

Below we offer one mechanism that overcomes the second difficulty and another one that overcomes both of them, both at extra processing cost.

A can allow B, or B can allow A, the chance to begin digesting immediately by using a confidential mechanism for the agreement of key as can be shown in the two following similar protocols.

Authentication without hashing III (non-interactive) *New*
$1\alpha.\ A \longrightarrow_N B : \{k\}_{pk(B)}$
$2.\quad A \longrightarrow_N B : M$
$3.\quad A \longrightarrow_N B : sign_A(B, hash(k), digest(k, M))$

Authentication without hashing III (interactive) *New*
$1\beta.\ B \longrightarrow_N A : \{k\}_{pk(A)}$
$2.\quad A \longrightarrow_N B : M$
$3.\quad A \longrightarrow_N B : sign_A(B, hash(k), digest(k, M))$

Notice that B's name and $hash(k)$ appearing in Message 3 prove to B that Message $1\alpha/\beta$ had k encrypted for B, not any other node. Furthermore, it represents proof to B that the key k is unknown to any one except A and B, we can remove the time stamp as well as timing constraints here. As a result, there is no restriction on the order of sending or receiving these messages: the order above is advantageous because it allows both parties to compute $digest(k, M)$ without delay. Clearly any other way of transmitting secret information from A to B could be used in place of the initial public key encryption.

If there is no need for both parties to digest simultaneously, the three messages can be combined into a single one, and indeed the secret transmission of k can be moved inside the final signed package

Authentication without hashing IV (non-interactive) *New*
$A \longrightarrow_N B : M, sign_A(B, \{k\}_{pk(B)}, digest(k, M))$

5.2 Flexi-MACs

None of the protocols above are relevant to the important practical problem of allowing one user to publish a piece of data together with some form of MAC that any recipient (the expectation being that there will be many of them) can check at any time. In methods I and II the key k is only valid for a short period, whereas in III and IV it is designed only for a single recipient.

We offer the following concept as a partial solution to this problem. It actually requires *more* effort on the part of the sender than the conventional approach, but of course we hope that this will be more than counterbalanced by the large number of recipients who can check it quickly.

All our protocols work by making A choose a key and not allowing the intruder to know the key until B is committed to M and the digest. It does not seem to be possible to achieve this in the circumstances we are now considering, so we turn it around and allow B to choose the key. But of course we are expecting B to be analysing data recorded by A (e.g. on a DVD), not with A herself, so this also sounds impossible. We can, however, simulate it by making A compute a large number of digests of M under different, randomly chosen, keys (set K), which she includes in a single signed block as her "Flexi-MAC" of the message M. B can then select any number of these values that it wishes to at random and check them.

Provided that the verification of each of these individual digests is much easier to compute than a single cryptographic hash, this should still be achievable more quickly than verifying a standard signature. It will also have the advantage that a single signature can be checked to different degrees depending on the perceived security threat and the time/computing power available. The mechanism can be summarised as follows.

- Flexi-MAC(M) $= \{(k, digest(k, M)) | k \in K\}$ concatenated with $sign_A(hash(\{(k, digest(k, M)) | k \in K\}))$
- To verify: select a random set C of the keys represented in K, of a chosen size, and check that the received M digests to the right value for each $k \in C$.

For example, suppose that our "Flexible MAC" consists of 1,000 signed key-and digest combinations, and that we believe that it is inconceivable that an attacker can have found a collision over more than three of these keys simultaneously. Then if the recipient chooses 1,2,3 of the keys at random, it follows that the chances of an attack succeeding are respectively bounded by 0.3%, 0.0006% and $1/(1.6 \times 10^8)$.

The effectiveness of this scheme will depend on how efficient and effective particular digest functions prove to be, and on how much assurance is required by every recipient B separately: we might be quite happy for any given recipient to have a 0.3% of being duped, either because of the application or because different B's share information: if a faked DVD is produced it is likely to be checked many times. We hope that there may be interesting applications in the area of DRM.

Note that by choosing the different keys k used in the Flexi-MAC randomly and after being committed to M, A has gained the same advantage as that of Pasini/Vaudenay relative to Balfanz: birthday attacks are eliminated and so a chosen plain-text will not be an advantage to an attacker in this sense.

In all the protocols we have suggested based on the uses of digest functions not transmitted by humans, there is not the same imperative for them to be *short*. Any reasonable fixed length will suffice. What we still require, of course, is that they are efficient to calculate.

6 Conclusion

In this paper we have analysed the security of two new and related classes of one-way authentication protocols.

We have derived the principle of *separation of security concerns*, that it can be inefficient to use a complex primitive for two difference, and factor-able, purposes. These concerns have an impact on the required length of the SASs manually handled by the human as illustrated in the first family of protocols based on human interaction: Long authentication string: Balfanz and Pasini-Vaudenay; Short authentication string: (V-)MANA I and its family of improved protocols.

The principle also has an impact on computational efficiency gained in our second family of protocol (authentication without hashing) as has been illustrated in Section 5.

The advantage provided by the schemes presented in Section 5.1 is only real if (a) we can substantiate our claim that digest functions can be computed significantly faster than hashes and (b) this advantage is not made irrelevant by issues such as the ratio between communication bandwidth (whether from memory or a peripheral) and processing speed. The second of these issues will vary greatly from application to application.

Both the Toeplitz model presented in Section 1.1 and the results of [9] suggest that the linear model of digest complexity quoted earlier is a little optimistic about how fast one might expect to compute a short one. The former does not satisfy this model since consumption of pseudo-random numbers is independent of the digest width. The latter shows that in order to create a near-perfect digest of any length, an accumulator of length a little larger than the output length must be maintained; this is a more important problem for short as opposed to long outputs. However, both pieces of work suggest that it should be possible to define very good digest functions that do not deviate from our model by very much in the range of output lengths that are likely to concern us. We will report on experiments on computing digest functions, and the comparison with hash functions, in a subsequent paper.

References

1. D. Balfanz, D. Smetters, P. Stewart, H. Wong. *Talking to strangers: Authentication in Ad Hoc Wireless Networks.* In Symposium on Network and Distributed Systems

Security, 2002.

2. M. Bellare, P. Rogaway. *Entity Authentication and Key Distribution.* CRYPTO 93, LNCS vol. 773, pp. 232-249.
3. C. Gehrmann, C. Mitchell, K. Nyberg. *Manual Authentication for Wireless Devices.* RSA Cryptobytes, vol. 7, no. 1, pp. 29-37, 2004.
4. J.-H. Hoepman. *Ephemeral Pairing Problem.* In 8th Int. Conf. Fin. Crypt., LNCS 3110, Springer, pp. 212-226.
5. H. Krawczyk. *LFSR-based Hashing and Authentication.* CRYPTO 1994, LNCS vol. 839, pp. 129-139.
6. L.H. Nguyen, A.W. Roscoe. *Efficient group authentication protocol based on human interaction.* Proc of FCS-ARSPA 2006, pp. 9-31.
7. L.H. Nguyen, A.W. Roscoe. *Authenticating ad hoc networks by comparison of short digests.* Journal of Information and Computation. Vol. 206, Issues 2-4, Feb-Apr 2008, pp. 250-271.
8. L.H. Nguyen, A.W. Roscoe. Authentication protocols based on low-bandwidth unspoofable channels: a comparative survey. Submitted to Journal of Computer Security.
9. L.H. Nguyen, A.W. Roscoe. *New theoretical bounds for Universal hash functions.* To appear. See `www.comlab.ox.ac.uk/people/publications/personal/Bill.Roscoe.html`
10. L.H. Nguyen, A.W. Roscoe. *Separating two roles of hashing in one-way message authentication,* See `www.comlab.ox.ac.uk/people/publications/personal/Bill.Roscoe.html`
11. S. Pasini, S. Vaudenay. *An Optimal Non-interactive Message Authentication Protocol.* CT-RSA'06, LNCS vol. 3860, pp. 280-294.
12. A.W. Roscoe. *Human-centred computer security.* Unpublished manuscript, 2005. `www.comlab.ox.ac.uk/people/bill.roscoe/publications/113.pdf`
13. D.R. Stinson. *Universal Hashing and Authentication Codes.* CRYPTO 1991, LNCS vol. 576, Springer-Verlag, pp. 74-85, 1992.
14. S. Vaudenay. *Secure Communications over Insecure Channels Based on Short Authenticated Strings.* CRYPTO 2005, LNCS vol. 3621.

A Security proofs of Improved MANA I

Adversarial model. Except for the authentic or empirical channels, we assume that the adversary has full control on the normal communication channel ($\longrightarrow_N$), for example: (s)he can block, delete, delay, redirect, fake and modify any information transmitted over this normal channel. In our security proofs, we will adopt the strong security model of Bellare and Rogaway [2] and its simplified version due to Vaudenay [14] which allow the adversary to have full control on which node a new instance of the protocol is launched.

In addition to this, following are the formal definitions of the terms: *one-shot* and *powerful* adversaries that will be used extensively in our security proofs.

1. **A powerful adversary** can launch multiple instances of participants (Alice and Bob in our case). As commonly known in the literature, the number of times that he can launch an instance of any participant is limited by a finite number, for example $\mathcal{Q}_A$ for Alice and $\mathcal{Q}_B$ for Bob. The time complexity of this adversary is bounded by a finite number say T. This is the kind of adversary we want to prove our protocols are resistant to in all security proofs presented here.
2. **A one-shot adversary** is a special case of the powerful adversary where the number of each participant's instances he can launch is at most once, in other words, $\mathcal{Q}_A = \mathcal{Q}_B = 1$.

The strategy of the proofs is to prove that our schemes are secure against a one-shot intruder in the first step, and then use Theorem 1 stated below to lift the one-shot intruder security proof to the powerful intruder's model.

The following theorem (and its proof) is a combined and slightly modified result of Lemma 6, introduced by Vaudenay [14], and Theorem 5 of Pasini and Vaudenay [11].

Theorem 1. [14, 11] *We consider three improved versions of Improved (V-)MANA I presented in Section 4 and Appendixes A.2 and A.3 with claimant Alice (A) and verifier Bob (B). We consider powerful adversaries such that the number of instances of Alice (resp. Bob) is at most $\mathcal{Q}_A$ (resp. $\mathcal{Q}_B$). If there exists a one-shot adversary against any of these protocols, bounded by time complexity t, has a probability of success smaller than p then the powerful adversary with time complexity $T = t \cdot \mathcal{Q}_A$ has a probability of success of $P \leq p \cdot \mathcal{Q}_A$.*

Proof. We say that an instance of A is compatible with an instance of B if B's instance succeeded and received all messages in the right order, where one of which is transmitted over the (bounded delay) empirical channel from the corresponding A's instance. The number of possible compatible pairs of instances is therefore upper bounded by $\mathcal{Q}_A \mathcal{Q}_B$, which could be reduced to $\mathcal{Q}_A$. This is because in both improved versions of MANA I and V-MANA I the single SAS (digest or random nonce) transmitted over empirical channels by definition given in Section 1.1 cannot be delayed from one to another session, and thus the number of compatible pairs of instance is upper bounded by the number of times that A's instance is launched by the powerful intruder.[9]

[9] In improved versions of V-MANA I, B can always be offline. As a result, the intruder can simulate all instances of B and picks one who will make the attack succeed.

When an attack is successful, there exists at least one compatible pair of instances of A and B. We note that these pairs of instances are independent of one another,[10] and the probability of success of each pair is limited to p in a time complexity t. Therefore, we can conclude that the powerful intruder bounded by a time complexity $T = t \cdot \mathcal{Q}_A$ is successful with a probability of $P \leq p \cdot \mathcal{Q}_A$.

A.1 Security proof of Improved (V-)MANA I (direct binding)

Improved version of V-MANA I (direct binding) *New*
1. Alice $\longrightarrow_N$ Bob : $M, hash(k)$
2. Alice $\longrightarrow^t_{BE}$ Bob : $digest(k, M)$
3. Alice $\longrightarrow_N$ Bob : k

Followings are extra security properties of cryptographic primitives that we need to have in order to prove the direct binding version secure.

General specification of digest functions. Instead of using 2^{-b}, we will refer to ϵ_d as the digest collision probability to make our security proofs presented in this appendix more general. We will assume $\epsilon_d = 2^{-b} + \mu$, where μ is some negligible value.

Specification of $digest()$
As the key k varies uniformly over its range:
1. $digest(k, M)$ is uniformly distributed for any fixed M.
2. For any fixed θ and $M \neq M'$:
$\mathbf{Pr}[digest(k, M) = digest(k, M')] \leq \epsilon_d$

Inversion-resistant hash functions. Inversion-resistance means that the Inversion-resistance game is hard. Assume that $hash()$ is a (T_h, ϵ_h) inversion-resistant hash function then any adversary $\mathcal{A}$ bounded by a time complexity T_h wins the Inversion-resistance game with probability at most ϵ_h. Since the bitlength of hash output is typically 160, we can assume that $\epsilon_h \ll 2^{-b}$. Here b is the length of digest, for example 16 or 20 bits.

[10] In any of these protocols, the random nonce or key instrumental in the computation of the SAS is generated by the trustworthy party A. Owing to the specification of empirical channels in Section 1.1 which does not allow messages to be delayed from one to another session, we therefore conclude that all compatible pairs of instances are probabilistically independent of one another.

Inversion-resistance game
The challenger $\mathcal{C}$ picks a random input x
1. $\mathcal{C} \longrightarrow \mathcal{A} : hash(x)$ 2. $\mathcal{A} \longrightarrow \mathcal{C} : x'$
Winning condition: $x = x'$

Weakly collision-resistant digest function. Since digest functions have a short output of b-bit, any adversary bounded by a time complexity T_d of order $\Theta(2^b)$ wins the weak collision-resistance game with probability of 1. In other word, digest is $(T_d, 1)$-weakly collision-resistant.

Weak collision-resistance game against $digest()$
The challenger $\mathcal{C}$ picks k and M at random.
1. $\mathcal{C} \longrightarrow \mathcal{A} : (k, M)$ 2. $\mathcal{A} \longrightarrow \mathcal{C} : (k', M')$
Winning condition: $digest(k, M) = digest(k', M')$ and $M \neq M'$

Theorem 2. *Consider the direct binding version of Improved V-MANA I authentication protocol. We assume that the function $hash()$ is a (T_h, ϵ_h) inversion-resistant hash function, and $digest(\_,\_)$ is $(T_d, 1)$ weakly collision-resistant. Then any powerful adversary against the protocol with time complexity bounded by $(T_h + T_d)\mathcal{Q}_A$ and with number of Alice's (resp. Bob's) instances bounded by $\mathcal{Q}_A$ (resp. $\mathcal{Q}_B$) has a probability of success at most $(\epsilon_h + \epsilon_d)\mathcal{Q}_A$.*

We note that the following proof can be slightly modified to cope with the direct binding version of Improved MANA I presented in Section 4.

Proof. We first find the time complexity and probability of success of a one-shot adversary.

A one-shot adversary has no advantage of sending $(M', hash(k'))$ after the digest is released in the 2^{nd} message thanks to the timing constraints of the bounded delay empirical channel. As a consequence, after $(M, hash(k))$ is sent in the 1^{st} message, there are two possibilities that can happen with the adversary, namely being *able* or *unable* to invert the hash to find out the key.

1. With a probability of ϵ_h, the adversary can invert the hash to discover the key within a time complexity of T_h. After that, with certainty (i.e. probability of 1) he will be able to find a different M' that produces a collision on the digest under key k in a time complexity of T_d.

Game against the direct binding version of Improved V-MANA I – Inverting hash function
1. Alice $\longrightarrow_N$ $\mathcal{A}$: $M, hash(k)$ $\mathcal{A}$ successfully inverts hash function to find out k $\mathcal{A}$ $\longrightarrow_N$ Bob : $M', hash(k)$
2. Alice $\longrightarrow^t_{BE}$ Bob : $digest(k, M)$
3. Alice $\longrightarrow_N$ Bob : k
Winning condition: $M \neq M'$ and $digest(k, M) = digest(k, M')$

2. Conversely, with a probability $(1 - \epsilon_h)$, the adversary fails to invert the hash value in time complexity T_h. Thus the adversary must select a pair (k', M') blindly, and hope that the fresh and unpredictable key k chosen by Alice will result in a digest collision.

Game against the direct binding version of Improved V-MANA I – Failing to invert hash function
1. Alice $\longrightarrow_N$ $\mathcal{A}$: $M, hash(k)$ $\mathcal{A}$ fails to invert hash. $\mathcal{A}$ picks a random pair (k', M') $\mathcal{A}$ $\longrightarrow_N$ Bob : $M', hash(k')$
2. Alice $\longrightarrow^t_{BE}$ Bob : $digest(k, M)$
3. Alice $\longrightarrow_N$ $\mathcal{A}$: k $\mathcal{A}$ $\longrightarrow_N$ Bob : k'
Winning condition: $M \neq M'$ and $digest(k, M) = digest(k', M')$

Clearly, the probability of success of this case is limited to $(1 - \epsilon_h)\epsilon_d$ thanks to the digest specification.

We conclude that any one-shot adversary bounded by a time complexity $t = T_h + T_d$ has a probability of success

$$p \leq \epsilon_h + (1 - \epsilon_h)\epsilon_d = \epsilon_h + \epsilon_d - \epsilon_h\epsilon_d < \epsilon_h + \epsilon_d$$

We now can apply Theorem 1 to deduce that any powerful adversary has a probability of success at most $(\epsilon_h + \epsilon_d)\mathcal{Q}_A$ within a time complexity of $(T_h + T_d)\mathcal{Q}_A$.

A.2 Improved (V-)MANA I (indirect binding) and security proof

In the following description, R is a b-bit random nonce of party A, which is inputted to a commitment scheme.[11]

Improved version of V-MANA I (indirect binding) *New*
1. Alice $\longrightarrow_N$ Bob : M, c
$\quad\quad\quad\quad (c, d) = \text{commit}(M, R)$
2. Alice $\longrightarrow^t_{BE}$ Bob : R
3. Alice $\longrightarrow_N$ Bob : d

While there is no relation between the SAS and M (i.e. they are completely independent in the sense of probability: indirect binding strategy), the security of the protocol comes from the use of a commitment scheme that firmly binds random nonce R to M in the eye of the receiver when the decommitment is released in Message 3. This, however, leads to an increase in computation cost because the many kilo-bytes information M must be processed by a long output commitment scheme that is more expensive than a short output digest function in direct binding solution, and this can be justified by the computational cost model of cryptographic primitives given in Section 1.1.

It is interesting to note that this protocol might also be regarded as the non-interactive version of the pairwise (indirect binding) authentication protocol of Vaudenay [14].

Similar to the direct binding version, we can replace the bounded delay empirical channel with a simple acknowledgement to have the following scheme.

Improved version of MANA I (indirect binding) *New*
1a. $A \longrightarrow_N B : M, c$
$\quad\quad\quad\quad (c, d) = \text{commit}(M, R)$
1b. $B \longrightarrow_E A$: 1-bit commitment
2. $A \longrightarrow_E B : R$
3. $A \longrightarrow_N B : d$

In order to support the security proof of the protocol, all we need is the following security property of a commitment scheme.

[11] Even though there is a difference between this and the commitment scheme used in Pasini-Vaudenay protocol in Section 2 regarding the extra input, a short random nonce R of b-bit, both of them can be implemented by the same algorithm. The style of this commitment scheme is also used by Vaudenay in [14].

Binding property of commitment scheme. A commitment scheme is (T_c, ϵ_c)-binding if any adversary $\mathcal{A}$ bounded by a time complexity T_c wins the following game with probability at most ϵ_c. When $T_c = +\infty$ and $\epsilon_c = 2^{-b}$, we say that the scheme is perfectly binding [14].

Binding game The adversary $\mathcal{A}$ picks a random pair (M, c). The challenger $\mathcal{C}$ picks b-bit random number R.
1. $\mathcal{A} \longrightarrow \mathcal{C} : (M, c)$ 2. $\mathcal{C} \longrightarrow \mathcal{A} : R$ 3. $\mathcal{A} \longrightarrow \mathcal{C} : d$
Winning condition: $\mathrm{open}(M, c, d) = R$

Theorem 3. *Consider the indirect binding version of Improved V-MANA I authentication protocol. We assume that the commitment scheme is (T_c, ϵ_c)-binding. Then any powerful adversary against the protocol with time complexity bounded by $T_c\mathcal{Q}_A$ and with number of Alice's (resp. Bob's) instances bounded by $\mathcal{Q}_A$ (resp. $\mathcal{Q}_B$) has a probability of success at most $\epsilon_c\mathcal{Q}_A$.*

The following proof applies to the indirect binding version of Improved V-MANA I, but it can be slightly modified to cope with the indirect binding version of Improved MANA I.

Proof. A one-shot adversary against this protocol follows the game depicted below in which it runs a man-in-the middle attack.

Game against the indirect binding version of Improved V-MANA I protocol
0. $\mathcal{A} \longrightarrow$ Alice : M
1. Alice $\longrightarrow_N \mathcal{A}$: M, c $(c, d) = \mathrm{commit}(M, R)$ $\mathcal{A} \longrightarrow_N$ Bob : M', c'
2. Alice $\longrightarrow^{t}_{BE}$ Bob : R
3. Alice $\longrightarrow_N \mathcal{A}$: d $\mathcal{A} \longrightarrow_N$ Bob : d'
Winning condition: $M \neq M'$ and $\mathrm{open}(M, c, d) = \mathrm{open}(M', c', d') = R$

Clearly this can be reduced to an adversary which plays the binding game against the commitment scheme. As a result, the success probability of a one-shot attack is equivalent to ϵ_c within a time complexity T_c.

Binding game
1. $\mathcal{A} \longrightarrow \mathcal{C} : M', c'$
2. $\mathcal{C} \longrightarrow \mathcal{A} : R$
3. $\mathcal{A} \longrightarrow \mathcal{C} : d'$
Winning condition: $\text{open}(M', c', d') = R$

We now can apply Theorem 1 to deduce that any powerful adversary has a probability of success at most $\epsilon_c \mathcal{Q}_A$ within a time complexity of $T_c \mathcal{Q}_A$.

A.3 Improved (V-)MANA I (Diffie-Hellman style) and security proof

We are going to describe another improved scheme whose main idea is taken root from the pairwise (direct binding) authentication protocol of Hoepman [4].

In the following description, k is a long secret key (160-bit) of A that corresponds to his Diffie-Hellman token g^k of which he wants to have authenticated. In order for the following protocol to be secure, the Diffie-Hellman token g^k must be fresh at each session, unpredictable and kept secret to A when its longhash and b-bit shorthash are revealed in the first two messages.

Improved version of V-MANA I (Diffie-Hellman style) *New*
1. Alice $\longrightarrow_N$ Bob : $hash(g^k)$
2. Alice $\longrightarrow^t_{BE}$ Bob : $shorthash(g^k)$
3. Alice $\longrightarrow_N$ Bob : g^k

The main difference between this and the previous two schemes is that there is no M sent in the 1^{st} message because the Diffie-Hellman token revealed in the 3^{rd} message plays the dual-role of both M and the long secret key. Also because of this, the digest function (used in the direct binding version) which requires 2 inputs can be replaced by a single input *shorthash*(); though the combination of this and the exponentiation of Diffie-Hellman needs to satisfy a specification similar to that of the digest.

In addition to an inversion-resistant *hash*(), we need to have the following security property for the *shorthash*() function.

Weakly collision-resistant shorthash function. Since *shorthash*()

has a short output of b-bit, any adversary bounded by a time complexity T_{sh} of order $\Theta(2^b)$ wins the weak collision-resistance game with probability of 1. In other word, $shorthash()$ is $(T_{sh}, 1)$-weakly collision-resistant.

Weak collision-resistance game against $shorthash()$ The challenger $\mathcal{C}$ picks M at random.
1. $\mathcal{C} \longrightarrow \mathcal{A} : M$ 2. $\mathcal{A} \longrightarrow \mathcal{C} : M'$
Winning condition: $shorthash(M) = shorthash(M')$ and $M \neq M'$

Theorem 4. *Consider the Improved V-MANA I protocol in Diffie-Hellman style. We assume that the function hash() is a* (T_h, ϵ_h) *inversion-resistant hash function, and shorthash() is* $(T_{sh}, 1)$ *weakly collision-resistant. Then any powerful adversary against the protocol with time complexity bounded by* $(T_h{+}T_{sh})\mathcal{Q}_A$ *and with number of Alice's (resp. Bob's) instances bounded by* $\mathcal{Q}_A$ *(resp.* $\mathcal{Q}_B$*) has a probability of success at most* $(\epsilon_h + \epsilon_d)\mathcal{Q}_A$.

Proof. The proof for this theorem is nearly identical to the proof of Theorem 2, and therefore it is left as an exercise for the readers.

B Combinatorial attack on (V-)MANA I

V-MANA I, [14, 11]
1. $A \longrightarrow_N B : A, M$ A picks a b-bit random number k 2. $A \longrightarrow_{SE} B : k, digest(k, A \parallel M)$

We term b and r the bit-lengths of the digest output and the key k (in this protocol, $b = r = 16$ bits). The intruder first chooses some number c different keys $\{k_1, \cdots, k_c\}$. Based on an off-line brute force search at the cost of $\Theta(2^{bc/2})$ computation steps, related to the birthday paradox, he can expect to find two different M and M' such that for all $k \in \{k_1, \cdots, k_c\}$,[12] we have:

$$digest(k, A \parallel M) = digest(k, A \parallel M')$$

Assuming that the adversary can influence A to send M in the 1st message of the protocol, there is then an attack it can attempt.

[12] It might be clearer if we define $H_{\{k_1,\cdots,k_c\}}(X) = digest(k_1, X) \parallel \cdots \parallel digest(k_c, X)$, and if $digest$ is an ideal digest function, then so is the function H w.r.t its $c \cdot b$ output-bits. As there is no limit on the bit-length of the input X, it normally takes $2^{cb/2}$ computation steps to search for a collision.

– The adversary blocks the message A, M that A sends, checking that it is the particular value that was desired.
– Immediately afterwards (to reduce the chance of A sending the empirical message too soon) it impersonates A to send A, M' to B.

Recall that in the above protocol, the key length r and digest length b are equal. The following calculations where these numbers are kept separate will allow us to draw more general conclusions.

After sending the 1^{st} message, A picks a random key k: with a probability of $\frac{c}{2^r}$, $k \in \{k_1, \cdots, k_c\}$, and the attack is successful. On the other hand, with a probability of $\frac{2^r - c}{2^r} \cdot 2^{-b}$, k is not in this set and so the attack is only successful with a probability of (presumably) 2^{-b}.

Overall, at the cost of $\Theta(2^{cb/2})$ due to the birthday paradox, the chance of a successful one-shot attack is:

$$\mathbf{Pr}_r(c) = c \cdot 2^{-r} + \frac{2^r - c}{2^r} \cdot 2^{-b}$$

When $r = b$ this is significantly larger than the desired probability of 2^{-b}.

The above vulnerability indicates we need to increase the bit-length r of the key to avoid this type of attack. When r increases, 2^r will quickly become significantly bigger than 2^b, this will allow the likelihood of a successful one-shot attack $\mathbf{Pr}_r(c)$ to converge to 2^{-b}. However, this is not feasible in this protocol because the key must be sent with the digest value over the strong empirical channel that is severely limited in bandwidth.