SAS-Based Authenticated Key Agreement

Sylvain Pasini and Serge Vaudenay

EPFL, CH-1015 Lausanne, Switzerland

<http://lasecwww.epfl.ch>

Abstract. Key agreement protocols are frequently based on the Diffie-Hellman protocol but require authenticating the protocol messages in two ways. This can be made by a cross-authentication protocol. Such protocols, based on the assumption that a channel which can authenticate short strings is available (SAS-based), have been proposed by Vaudenay. In this paper, we survey existing protocols and we propose a new one. Our proposed protocol requires three moves and a single SAS to be authenticated in two ways. It is provably secure in the random oracle model. We can further achieve security with a generic construction (e.g. in the standard model) at the price of an extra move. We discuss applications such as secure peer-to-peer VoIP.

1 The SAS-Based Authenticated Key Agreement Problem

Secure communication channels are usually set up by authenticated key agreement protocols. This can be performed by relying on a public-key infrastructure, e.g. based on RSA [RSA78] or the Diffie-Hellman protocol [DH76]. Clearly, this is not well suited to the advent of mobile ad-hoc communications where ephemeral or bootstrap connections are needed “at once”: we certainly would not like to register a certificate to connect a PDA to a cell phone or to print to the neighbor available printer device. Secure communications can also be manually set up. For instance, peer-to-peer links using PGP can be set up by checking the digest of a public key over the telephone. Wireless devices can be securely connected by having the user to manually check a hashed value as well. To save the human user load, the string to be manually checked must be as short as possible. Recently, protocols based on Short Authenticated Strings (SAS) have been studied by Vaudenay [Vau05]. It was shown how to design and analyze a protocol to authenticate an arbitrary string assuming that we can authenticate a short one over a dedicated secure channel. Those protocols are based on commitment schemes. It was also briefly proposed how to design message cross-authentication protocols, namely protocols to authenticate arbitrary strings in two ways.

A SAS-based Authenticated Key Agreement (AKA) protocol can be easily designed by running the Diffie-Hellman protocol over an insecure channel, then by authenticating the digest of the protocol transcript using a SAS-based message cross-authentication protocol. This typically results in a 5-move protocol in addition to the bidirectional SAS transmission. In the present work, we show how to decrease the interaction cost. Namely, we design a generic construction which can use a 4-move protocol in addition to the bidirectional SAS exchange. This construction can rely on the standard model

(without random oracles). We also design an optimal 3-move protocol which is provably secure (with tight reduction) in the random oracle model.¹

2 Preliminaries

We adopt the security model from [Vau05, Vau06, PV06] based on the one from Bellare-Rogaway [BR93]. We consider a network of participants which are located at some nodes. A participant at node n is associated to a given identity ID_n . He locally maintains a database of (K_j, ID_j) pairs meaning that he can use the symmetric key K_j to securely communicate with ID_j in a private and authenticated way. Participants can run concurrent protocols. A protocol specifies a sequence of steps which consist of receiving a message and sending a response. An internal short-term state keeps track on previously completed steps. Once the protocol is completed, the short-term state is removed. A protocol starts with some specified inputs and an initial state (in terms of database content). It ends with some specified outputs (or an error message) and a final state. The difference between an input (resp. output) and an initial (resp. final) state is that the adversary has control on the first one but not on the second one, except if the node was corrupted or some information leaked. Protocol instances on a node n are denoted by a unique tag π_n^i . (Note that the state of a protocol related to a given tag changes with time as new steps are made.)

Nodes can communicate through an insecure broadband channel. In addition, they have access to peer-to-peer narrowband channels which can be used to authenticate short messages. A node receiving a message from one of these channels is ensured that this message was sent at some time in the past by a node whose identity is specified by the channel itself. In this paper, we concentrate on key agreement and cross-authentication protocols, so we assume that nodes share no prior exchanged keys.

2.1 Adversarial Model

By default, the adversary is assumed to have a full control on which node makes a new step of a given protocol instance, on the insecure channel, can influence the delivery of messages (without modifying them) over the authenticated channels, can choose the inputs of the protocols, and has access to the outputs. Occasionally, the adversary can violate the privacy of the internal state of a given node or even corrupt the node so that his behavior with respect to future runs of any protocol is no longer guaranteed. More formally, the adversary has access to the following oracles.

Launch. $\text{launch}(n, \text{role}, x)$ launches a new protocol instance on node n playing role (e.g. either Alice or Bob) with input x . It returns a new instance tag π_n^i . Note that the instance inherits of the current node state as its input state.

Send. $\text{send}(\pi, y)$ sends an incoming message y to the instance π . It returns an outgoing message z , or the final output of the protocol if it completed.

¹ After the present paper was submitted, a preprint was posted by Laur, Asokan, and Nyberg [LAN05]. This paper includes another 3-move protocol which is provably secure based on a generic commitment (e.g. in the standard model) but not optimal.

Test. $\text{test}(n, k, \text{ID})$ tells whether (k, ID) is an entry of the database of node n . In practice, this oracle may be implemented by an active adversary trying to impersonate node n to communicate with ID . If the attempt succeeds, it means that k was the right key to use.

Remove. $\text{remove}(n, \text{ID})$ removes any (k, ID) entry in the database of node n . In practice, this oracle may be implemented by an adversary making denial-of-services attacks in the communication link between n and ID so that n decides not to trust this connection anymore and to remove it.

Reveal. $\text{reveal}(n)$ reveals the full current state of node n . This models side channels or careless uses.

Corrupt. $\text{corrupt}(n)$ injects a malicious code in node n so that its behavior is no longer guaranteed.

The *attack cost* is measured by

- the number Q of launched instances of Alice or Bob, i.e. the *online complexity*.
- the additional complexity C , i.e. the *offline complexity*.
- the probability of success p .

We call *one-shot attacks* the attacks which launch only two instances in total, i.e. $Q = 2$.

By convention, we describe protocols by putting a *hat* on the notation for messages received by a node (i.e. inputs of the send oracle) which are not authenticated since they can differ from messages which were sent (i.e. outputs of the receive oracle) in the case of an active attack. A message m from a node of identity ID over an authenticated channel is denoted $\text{authenticate}_{\text{ID}}(m)$.

2.2 Key Agreement, Cross-Authentication, and Mutual Authentication

Authenticated key agreement. An Authenticated Key Agreement (AKA) protocol between Alice and Bob starts with no input, is independent from the current state, and ends with no output but a final state specifying an entry (k, ID) to be inserted in the database: Alice of identity ID_A ends with (k, ID_B) and Bob of identity ID_B ends with (k, ID_A) . An attack is successful if a $\text{test}(n, k, \text{ID})$ query positively answered where n and ID correspond to nodes on which no reveal nor corrupt query was made. For simplicity, we do not consider attacks making Alice and Bob end on some inconsistent states. Namely, *mutual authentication* is assumed to be (implicitly or explicitly) made by further communications.

To construct AKA protocols, we use the following building blocks.

Message cross-authentication. A Message Cross-Authentication (MCA) protocol between Alice and Bob of identity ID_A and ID_B starts with inputs m_A and m_B and ends with outputs (m_B, ID_B) and (m_A, ID_A) , respectively. An adversary is successful if some instance ended on an incorrupted node with a pair (m, ID) but no instance on the node of identity ID with input m was launched. Note that test, remove, and reveal oracles are not relevant in this case.

Message mutual-authentication. A Message Mutual-Authentication (MMA) protocol between Alice and Bob of identity ID_A and ID_B starts with inputs m_A and m_B and ends with outputs ID_B and ID_A , respectively. A honest run of an MMA protocol must have

$m_A = m_B$. An adversary is successful if some instance on an incorrupted node started with any m and ended with any ID such that no instance on the node of identity ID with input m was launched. As for MCA protocols, test, remove, and reveal oracles are not relevant. Obviously, we can transform an MCA protocol into an MMA protocol by just checking that the output message is equal to the input one on both sides.

MCA from MMA. We can also transform an MMA protocol with at least one move over the insecure channel into an MCA protocol at the price of an extra move: Bob of identity ID_B first sends his input message m_B and Alice of identity ID_A initiates an MMA protocol with input $m_A || \hat{m}_B$ by sending m_A together with the first MMA protocol message. Bob then follows the MMA protocol with input $\hat{m}_A || m_B$. The final outputs of Alice and Bob are $(\hat{m}_B, ID_B)$ and $(\hat{m}_A, ID_A)$ respectively.

To compare protocols we focus on the number of message moves over the insecure channel and on the length of authenticated messages. Furthermore, a protocol with two equal SAS to be sent in both directions (called symmetric SAS) will be considered as better than a protocol with two SAS of similar length (but not necessarily equal) to be exchanged. Indeed, some authentication channels may provide symmetric authentication at no extra cost.

2.3 Equivocable Commitment and Random Oracle Commitment

In this paper, we consider (tag-based) equivocable commitment schemes as defined by two algorithms commit and open and three oracles setup, simcommit, and equivocate.

Setup. $K_P \leftarrow \text{setup}$ generates a public key K_P to be used as a common reference string and a secret key K_S to set up the simcommit and equivocate oracles. The public key K_P is implicitly used by all other algorithms and oracles but omitted in the notations for simplicity.

Commit. $(c, d) \leftarrow \text{commit}(m, r)$ generates a commit value c and a decommit value d for a key r with a tag m . We assume that the distribution of the generated c is independent from r : the commitment is *perfectly hiding*.

Open. $r \leftarrow \text{open}(m, c, d)$ yields r if (c, d) is a possible output for $\text{commit}(m, r)$.

Simcommit. $(c, i) \leftarrow \text{simcommit}(m)$ simulates a commit value c for a tag m and produces extra information ξ to be used later. The distribution of c should be the same as for the distribution of c generated by any $\text{commit}(m, r)$. It also creates a unique identifier i (a nonce) and inserts (i, m, c, ξ) in a database. This oracle uses the secret key K_S and should be secured. Access to the database must be restricted to this oracle and equivocate.

Equivocate. $d \leftarrow \text{equivocate}(i, r)$ yields d such that $r = \text{open}(m, c, d)$ where (i, m, c, ξ) is in the database of simcommit. This entry is further removed. (Namely, a simulated c can be equivocated only once.)

Access to simcommit and equivocate oracles is restricted depending on the application. The normal usage of the commitment scheme should be limited to commit and open but we stress that our security model assumes that the adversary may cheat on *some* commitments by having access to simcommit and equivocate oracles. Indeed,

our notion of equivocal commitment relates to the notion of *simulation-sound* commitment [MY04].

The *hiding game* between a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$ runs as follows.

1. $\mathcal{C}$ runs setup and sends K_P to $\mathcal{A}$
2. $\mathcal{A}$ sends a tag m to $\mathcal{C}$
3. $\mathcal{C}$ commits to a random key with tag m and sends a commit value c to $\mathcal{A}$
4. $\mathcal{A}$ computes some r and sends it to $\mathcal{C}$
5. $\mathcal{C}$ releases a decommit value d and $\mathcal{A}$ wins if $r \leftarrow \text{open}(m, c, d)$

In that case, the adversary has access to the simcommit and equivocate oracles but cannot query simcommit with the selected tag m . Since the commitment is perfectly hiding, no adversary can win this game with a probability larger than 2^{-k} where k is the length of the key r .

The *binding game* between a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$ runs as follows.

1. $\mathcal{C}$ runs setup and sends K_P to $\mathcal{A}$
2. $\mathcal{A}$ sends a tag m and a commit value c to $\mathcal{C}$
3. $\mathcal{C}$ picks a random r and sends it to $\mathcal{A}$
4. $\mathcal{A}$ produces a decommit value d and wins if $r \leftarrow \text{open}(m, c, d)$

We say that the commitment with k -bit keys r is (T, ϵ) -secure if any adversary with complexity limited to T has a winning probability of at most $2^{-k} + \epsilon$. In that case, the adversary has access to the simcommit and equivocate oracles but cannot query simcommit with the selected tag m .

Secure equivocal commitment schemes can be easily constructed based on simulation-sound trapdoor commitments by MacKenzie-Yang [MY04] as detailed in [Vau05]. Constructions can be in the standard model with a common reference string, e.g. based on the security of DSA signatures [DSS00] or Cramer-Shoup signatures [CS02]. We can also build an efficient equivocal commitment scheme based on the random oracle model.

Random oracle commitment scheme. Let ℓ_c , ℓ_e , and k be three integers. The setup algorithm is unused, but we assume that we can use three oracles:

- H.** $c \leftarrow H(e, r, m)$ queried with an ℓ_e -bit string e and a k -bit string r , looks whether an entry (e, r, m, c) in a list exist. If not, the oracle creates one with a random ℓ_c -bit string c . In any case, the oracle answers c .
- Simcommit.** $(c, i) \leftarrow \text{simcommit}(m)$ simply picks a random ℓ_c -bit string c and a nonce i and stores (i, c, m) in a list.
- Equivocate.** $d \leftarrow \text{equivocate}(i, r)$ gets (i, c, m) and removes it from the list. The oracle then picks a random ℓ_e -bit string e . If $(e, r, m, \cdot)$ exists in the H list, the oracle fails. Otherwise, (e, r, m, c) is inserted. Clearly, if the number of oracle accesses to H and simcommit is limited by q , the probability that the oracle fails at least once is less than $q^2 \times 2^{-\ell_e-1}$.

The algorithm $\text{commit}(m, r)$ simply picks e at random, queries $H(e, r, m)$ and outputs $d = (e, r)$. The algorithm $\text{open}(m, c, d)$ simply checks that $H(d, m) = c$ and parses

$d = (e, r)$ to yield r . Unless equivocate fails, this scheme is clearly an equivocable commitment scheme as previously defined. Since all commit values c are generated in an independent way, there are no collisions with probability at least $1 - q^2 \times 2^{-\ell_c - 1}$. Clearly, being able to decommit any c to two values would lead H to a collision. Hence, the scheme is $(q, 2^{-k} + q^2 \times 2^{-\ell_e - 1} + q^2 \times 2^{-\ell_c - 1})$ -secure. In practice, simcommit and equivocate are unused. So, we can just instantiate H by a standard hash function, provided that instantiation of that kind of random oracle makes sense [CGH98].

3 Previous SAS-Based Key Agreement Protocols

A classical authenticated Diffie-Hellman [DH76] protocol over a multiplicative group spanned by a generator g consists, for Alice (resp. Bob) of picking a random integer x_A (resp. x_B), sending the Diffie-Hellman public keys, $y_A = g^{x_A}$ (resp. $y_B = g^{x_B}$) over the authenticated channel, computing $z_A = y_B^{x_A}$ (resp. $z_B = y_A^{x_B}$) and ending with state (z_A, ID_B) (resp. (z_B, ID_A)). In this case, authenticated messages are pretty long, but authentication is necessary to thwart man-in-the-middle attacks.

We first informally present an AKA protocol from Hoepman [Hoe04]. It is based on the Diffie-Hellman protocol and it uses an authenticated channel for the authentication of each Diffie-Hellman value. This protocol runs in three steps: commitment, authentication, and opening. (The original protocol has a fourth step: the key validation.) Instead of revealing its Diffie-Hellman public key, each party first commits on it, keeping it hidden. In the next step, each participant authenticates a piece of its Diffie-Hellman public key. Finally, they open their commitments and check their respective commitment and authenticated string before completing the regular Diffie-Hellman protocol.

Fig. 1. PGPfone 1995 Key Agreement Protocol

Another AKA protocol, depicted on Fig. 1, was used by Zimmermann for the PGPfone in 1995². Its advantage is to reduce the number of moves in the insecure channel and to make both authenticated strings equal. In this protocol, only the first participant Alice commits to its public key. The commitment is immediately opened when the other

² Personal communication.

public key is received. Finally, the authenticated string is a piece of the digest (denoted truncH on Fig. 1) of the Diffie-Hellman protocol transcript.

For both the Hoepman and the PGPfone protocols, the security is not formally proven ([Hoe04] only provides a sketch of argument for the security). Another approach consists of authenticating the transcript of a classical key agreement protocol by using an MMA protocol. The MANA protocols by Gehrman-Mitchell-Nyberg [GMN04, GN01, GN04] illustrates this. Finally, we study in what follows a generic construction reducing the amount of authenticated bits in AKA protocols. Using it with the Diffie-Hellman key agreement protocol and the MCA protocol of [Vau05], we obtain the DH-SC protocol of Čagalj-Čapkun-Hubaux [ČČH06]. Using an optimized MCA protocol we can save one protocol move. In what follows we describe the generic construction, analyze it, and study 3-move MMA and MCA protocols with symmetric SAS.

4 Reducing Key Agreement to Message Authentication

We can build an AKA protocol by exchanging Diffie-Hellman keys through a message cross-authentication protocol.

We propose a generic SAS-based construction for an AKA protocol that we call the constructed AKA protocol or simply *the* AKA protocol. For this, we use an initial AKA protocol (with longer strings to be authenticated), that we call the AKA_0 protocol, and an MCA protocol with short SAS. Consider that the AKA_0 protocol requires $n_k \geq 2$ moves, the $n_k - 1$ -th being from Alice to Bob, and the MCA protocol requires $n_a \geq 2$ moves over the insecure channel, the first one being from Alice to Bob. In the AKA protocol, the $n_k - 2$ first moves of the AKA_0 are performed over the insecure channel. Then, both participants assemble his view on the protocol transcript τ by concatenating all protocol messages (sent and received ones). Then, an MCA protocol starts. Alice wishes to authenticate τ concatenated with her $n_k - 1$ -th message α in the AKA_0 protocol. Bob wishes to authenticate the same $\tau || \alpha$ concatenated with his last message β in the AKA_0 protocol. (Note that Bob selects the message to be authenticated after receiving Alice's first message in the MCA protocol.) At the end, both participants use the authenticated messages to complete the AKA_0 protocol and end with final states as specified in the AKA_0 protocol. We have $n_k + n_a - 2$ moves in total.

Note that MCA can have $n_a < 2$. (For instance the trivial MMA protocol exchanging authenticated digests has no move and thus we can build an MCA with only one move.) In that case, we augment the MCA protocol by virtual moves and we obtain n_k moves in total. However, MCA protocols with $n_a < 2$ must have pretty large SAS to exchange the messages.

We can make a similar construction based on an n'_a -move MMA protocol instead of an MCA protocol. In that case, we can only encapsulate the last move β of the AKA_0 protocol in the MMA protocol, leading us to $\max(n_k, n_k + n'_a - 1)$ moves in total.

Theorem 1. *Let us consider an n_k -move AKA protocol (the AKA_0 protocol) and an n_a -move MCA protocol. The generic construction is essentially an AKA protocol with $\max(n_k, n_k + n_a - 2)$ moves in which the structure of authenticated messages is similar as in the MCA protocol. There exists a constant μ such that for any T , if ϵ_1 resp. ϵ_2 denotes the best success probability of an adversary bounded by T against the AKA_0*

protocol resp. the MCA protocol, then any adversary bounded by $T \times \mu$ against the AKA protocol has probability of success at most $\epsilon_1 + \epsilon_2$.

Using the Diffie-Hellman protocol and an n_a -move MCA protocol leads us to a $\max(2, n_a)$ -move AKA protocol in which the structure of authenticated messages is similar as in the MCA protocol. With the construction based on an MMA protocol, we obtain $\max(2, n'_a + 1)$ moves. In the case where we want to achieve small SAS, we must have $n_a \geq 2$, leading us to n_a moves using MCA protocols and $n'_a + 1$ moves using MMA protocols. Since $(n'_a + 1)$ -move MCA protocols can be made from n'_a -move protocols, we may decrease the total number of moves in AKA protocols by starting from MCA protocols directly.

Proof. For each instance of Alice, we let τ_A be the constructed transcript of the $n_k - 2$ first messages in the AKA protocol and we let α_A be her last message, i.e. the $n_k - 1$ -th message in the protocol. We further let $\hat{\tau}_B || \hat{\alpha}_B || \hat{\beta}$ be the accepted message from Bob at the end of the MCA protocol. Similarly, for each instance of Bob, we let τ_B be the constructed transcript of the $n_k - 2$ first messages in the AKA protocol, $\hat{\tau}_A || \hat{\alpha}_A$ be the accepted message at the end of the MCA protocol, and β be his last message in the AKA₀ protocol assuming that Alice's last one is $\hat{\alpha}_A$. We let $\alpha_B = \hat{\alpha}_A$. Bob's message to be authenticated is $\tau_B || \alpha_B || \beta$.

Given an adversary $\mathcal{A}$ against the AKA protocol, we construct a simulator $\mathcal{B}$ interacting with $\mathcal{A}$ and attacking the MCA protocol. We simply simulate instances running the AKA₀ protocol and launch the MCA protocol instances when appropriate. test, remove, reveal and corrupt queries can easily be simulated. Clearly, the attack against the MCA protocol does not succeed with probability at least $1 - \epsilon$. In those cases, we have $\tau_B = \hat{\tau}_B$, $\tau_A = \hat{\tau}_A$, $\alpha_A = \hat{\alpha}_A = \alpha_B = \hat{\alpha}_B$, and $\beta = \hat{\beta}$, just as if the instance of Alice and Bob had the AKA₀ protocol run over an authenticated channel.

We construct a simulator $\mathcal{C}$ interacting with $\mathcal{A}$ and attacking the AKA₀ protocol over an authenticated channel. The simulator simply replaces inputs to the send oracle by authenticated ones when possible, or fails, and simulates the MCA protocol. Clearly, running $\mathcal{A}$ in parallel with $\mathcal{B}$ and $\mathcal{C}$ with the same random source, we derive that whenever $\mathcal{A}$ succeeds, either $\mathcal{B}$ or $\mathcal{C}$ succeed. $\square$

A trivial MMA protocol consists of authenticating the digest of the input message from a collision-resistant hash function. This protocol can be transformed into an MCA protocol by using 2 moves (to exchange m_A and m_B) plus the authentication of a SAS in two ways as for the construction in Section 2.2. We obtain a 2-move AKA protocol with symmetric SAS, but the length of the SAS is quite long (typically, 160 bits).

A SAS-based cross-authentication protocol was proposed in [Vau05] by interleaving two SAS-based message authentication protocols. It is a 4-move MCA protocol with symmetric SAS and can thus be transformed into a 4-move AKA protocol with symmetric SAS based on Diffie-Hellman.

5 A New SAS-Based Message Mutual-Authentication Protocol

We propose a new protocol improving the number of exchanged messages. As depicted on Fig. 2, and without any attack, Alice and Bob start with the same message, i.e.

Fig. 2. A New SAS-Based Message Mutual-Authentication Protocol

$m_A = m_B$. Each participant chooses a k -bit random value R_A and R_B , respectively. Alice starts by committing on her random value R_A by sending c , keeping it hidden. Bob sends the random value R_B . Then, Alice opens her value by sending the decommit value d . Finally, both authenticate the SAS which has been computed using a simple XOR function. Using the generic construction with Diffie-Hellman we obtain a 4-move AKA protocol with symmetric SAS.

Theorem 2. *We consider adversaries against the MMA protocol of Fig. 2 who are bounded by complexity T , Q_A instances of Alice, and Q_B instances of Bob. We assume that we have an (T_C, ϵ) -secure equivocable commitment scheme. There exists a (small) constant μ such that any adversary wins either with probability limited to $Q_A \cdot (Q_A + Q_B)(2^{-k} + \epsilon)$ or with complexity $T \geq T_C - \mu$.*

Proof. Any adversary which would attack an instance of either Alice or Bob needs one SAS to send her/him so that she/he can complete. This required SAS can easily be obtained from any instance of Alice since she does not need any prior authenticated message. It can also be obtained from any instance of Bob in which case he must be sent another SAS before. The output SAS by Bob is equal to the sent one. Indeed, a successful adversary interaction defines the first attacked instance and a prior sequence initiated by one instance of Alice followed by a chain (possibly empty) of instances of Bob and ended by the attacked instance. Every (unattacked) instance of Bob in this sequence is sending a SAS identical to the received one to the next instance. Every intermediate instance of Bob terminates with an output message which must be equal to the input message of the previous instance in the sequence (otherwise, they would be successfully attacked). However, the final instance in the sequence outputs a message which is different than the input of the previous instance. Hence, every instance in the sequence but the final one has the same input message and all instances yield the same SAS. Clearly, sending the output SAS from the leading Alice to the tailing instance produces a successful attack with no intermediate instance of Bob.

Let $\mathcal{A}_0$ be an adversary who launches at most Q_A instances of Alice and Q_B instances of Bob. We transform it into an adversary $\mathcal{A}$ who launches an instance of Alice and a single target instance (of either Alice or Bob) as follows:

1. $\mathcal{A}$ first picks two random numbers I, J such that $1 \leq I \leq Q_A$ and $1 \leq J < Q_A + Q_B$.
2. We initialize counters i and j to 0 and run $\mathcal{A}_0$ step by step.
 - Every time $\mathcal{A}_0$ would like to make a launch query to launch an instance of Alice, we increment i . If $i = I$, we really launch it and call the instance Alice π . Otherwise, we increment j and if $j = J$, we really launch it and call the target instance π' . Otherwise, we simulate the oracle call.
 - Every time $\mathcal{A}_0$ would like to make a launch query to launch an instance of Bob, we increment j . If $j = J$, we really launch it and call the target instance π' . Otherwise, we simulate the oracle call.
 - If we have to send a SAS to π , we just simulate the oracle call.
 - If we have to send a SAS to π' and we already got a SAS from π which is equal to the expected one, we just send it. Otherwise, the attack fails.

Due to the previous discussion, if $\mathcal{A}_0$ succeeds, if π' is the first attacked instance for $\mathcal{A}_0$ and if π is the leading instance of Alice in the sequence, then $\mathcal{A}$ succeeds. Hence, the probability of success of $\mathcal{A}$ is at least $\frac{1}{Q_A(Q_A+Q_B-1)}$ times the probability of success p of $\mathcal{A}_0$.

We now have an adversary $\mathcal{A}$ with Alice and a target instance. We assume that the adversary complexity is bounded by $T_C - \mu$ for some constant overhead μ to be determined by the following reductions. We consider two cases: attacks targeting an instance of Bob and attacks targeting an instance of Alice. Let p_A resp. p_B be the probability of a target Alice resp. Bob and q_A resp. q_B be the success probability conditioned to both cases, respectively. The success probability of $\mathcal{A}$ is $p = q_A p_A + q_B p_B$ and we have $p_A + p_B = 1$.

In both cases, we define a simulator $\mathcal{B}$ who simulates the two instances as follows. We first pick a random k -bit SAS. When an instance of Alice is launched for the first time by the adversary $\mathcal{A}$, we simulate a commitment c by using `simcommit`. Then the corresponding $\hat{R}_B$ is sent to this instance of Alice, the commit value is equivocated so that it opens to the key $\text{SAS} \oplus \hat{R}_B$. This simulation of Alice is perfect and has the

Fig. 3. Simulator Playing the Hiding/Binding Game

property to determine the final SAS at the beginning. If the attack succeeds, the other instance will have to deal with a commit value with a different tag. Depending on whether the other instance is an Alice or a Bob, we simulate it so that we can win the hiding game or the binding game against a challenger C as depicted on Fig. 3. In the case of a target Alice, the adversary succeeds if $\hat{R}_B$ leads the target instance to derive SAS. In that case we can correctly derive R_A and win the hiding game. Since the equivocable commitment is always perfectly hiding, we deduce $q_A = 2^{-k}$. We could have played the binding game in a trivial way and won with the same probability 2^{-k} . In the case of a target Bob, the adversary succeeds if $\hat{d}$ decommits to a key which leads Bob to the right SAS, thus to the key $\hat{R}_A$. In that case, we win the binding game with probability q_B . To summarize, we made an adversary playing the binding game with probability of success p . Therefore, $p \leq 2^{-k} + \epsilon$. $\square$

6 A New SAS-Based Message Cross-Authentication Protocol

We propose a new protocol based on the previous one, but improving the number of exchanged messages through the broadband insecure channel. Our protocol uses an almost strongly universal hash function family h [Sti91, Sti94]. In practice, one can use $h_K(x) = \text{trunc}(\text{hash}(K||x))$ where hash is a collision-resistant hash function and trunc truncates to the leading p bits. Our protocol also uses a commitment scheme to commit on a κ -bit key K . Contrarily to our previous protocol, the committed key K can now be pretty large. Using the generic construction with Diffie-Hellman we obtain a 3-move AKA protocol with symmetric SAS. Note that we added an identity test on Alice's side to avoid trivial reflection attacks.

Theorem 3. *Let ℓ_e, ℓ_c be the parameters of the random oracle commitment scheme. Let q be the upper bound on the number of H queries. Let $\epsilon = q^2 2^{-\ell_e} + q^2 2^{-\ell_c}$. Let h be an ϵ_h -almost strongly universal hash function family with p -bit digests, i.e. $\Pr[h_K(a) =$*

Fig. 4. A New SAS-Based Message Cross-Authentication Protocol

$\alpha, h_K(b) = \beta] \leq 2^{-2p} + 2^{-p}\epsilon_h$ for any a, b, α, β such that $a \neq b$ with a random K .³ We consider adversaries against the message cross authentication protocol of Fig. 4 who are bounded by Q instances of Alice or Bob and by q queries to H . The success probability is limited by $\frac{Q(Q-1)}{2}(2^{-p} + \epsilon + \epsilon_h)$.

By launching Q instances of either Alice or Bob with pairwise different input messages and by picking independent uniformly distributed $\hat{R}$, all SAS are independent and uniformly distributed so we have one matching with probability $1 - 2^{-Qp} \cdot 2^p! / (2^p - Q)!$ which is roughly $\frac{Q(Q-1)}{2}2^{-p}$ when $Q \ll 2^{\frac{p}{2}}$. Hence, this bound is essentially tight. Note that the above attack can apply to any MCA protocol of similar structure (see [PV06]), so our protocol is optimal.

Proof. We let $\epsilon_1 = \epsilon_2 = \epsilon_h$ (h is almost uniform). We have $\Pr[h_K(a) = \alpha] \leq 2^{-p} + \epsilon_1$ and $\Pr[h_K(a) \oplus h_K(b) = \alpha] \leq 2^{-p} + \epsilon_2$ for any a, b, α such that $a \neq b$ and with K uniformly distributed (h is ϵ -almost XOR universal [Kra94]). In what follows, only those properties will be used. Namely, we could replace the condition on h by those two properties.

We define a new character: the flipped Bob who proceeds as Bob but first issues a SAS equal to $R \oplus h_K(m_B)$ then receives a SAS for verification. In a new protocol, Alice and the flipped Bob can interact with two crossing SAS exchange.

We consider an adversary successfully running his attack with many instances for the original MCA protocol. We say that a given instance is attacked if it completed the protocol during which a SAS was received, with an output which is not consistent with the input of the instance who issued the received SAS. (Note that a successful adversary must have an attacked instance.) An attacked (target) instance (of either Alice or Bob) must receive one SAS from a (sending) instance. Note that those two instances must be different. (Indeed, no instance of Bob can send a SAS to himself otherwise it would have to be received before being sent. Similarly, no instance of Alice can accept a SAS coming from herself.) Clearly, both instances must agree on the SAS to complete. Hence, if the SAS sent by the target instance is forwarded to the sending instance then both instances fully interact. We can guess the pair of instances with probability $2/(Q(Q-1))$. Hence, we can simulate all instances except the two guessed ones. Since the SAS verification phase is the last step on both instances, there is no trouble to make the two instances exchange their SAS. We thus transform the initial adversary against the MCA protocol with success probability p into a one-shot adversary against our new protocol with success probability at least $2p(Q(Q-1))^{-1}$.

The interaction of the transformed adversary with an instance of Alice consists of two steps

- A_1 sending her her message m_A (for the launch query) and getting her commit value c (for the first send query)
- A_2 giving her Bob's alleged message $\hat{m}_B$ and random value $\hat{R}$ and getting her decommit value d .

³ Note that this definition of almost strongly universal hashing is slightly different from [Sti91, Sti94] in the sense that perfect uniformity is not required.

Alice's SAS equals $\hat{R} \oplus h_K(\hat{m}_B)$ where K is the result of $\text{open}(m_A, c, d)$. The second step must be performed after the first one.

The interaction of the adversary with an instance of Bob consists of two steps

- B_1 sending him his message m_B (for the launch query) and Alice's alleged message $\hat{m}_A$ and commit value $\hat{c}$ and getting his random value R (for the first send query)
- B_2 giving him Alice's alleged decommit value $\hat{d}$.

The adversary wins if the two instances complete and compute the same SAS and if the input message of one instance is different from the output message of the other instance.

In what follows we show that all cases can be simulated so that we can win a hard game, proving that the probability of success is at most $2^{-p} + \varepsilon + \max(\varepsilon_1, \varepsilon_2)$.

Cases Alice-Alice. We number 2 the instance of Alice whose A_2 step is the last. Since the commitment is perfectly hiding, this Alice leaks no information about K^2 (variable K for Alice number 2) until this very last step. Hence, K^2 is independent from the rest and $\hat{R}^1 \oplus \hat{R}^2 \oplus h_{K^1}(\hat{m}_B^1) = h_{K^2}(\hat{m}_B^2)$ with probability at most $2^{-p} + \varepsilon_1$.

On Bob's Incoming $\hat{c}$ (Step B_1). In the random oracle commitment model, we only consider the event where no collision occurred. Hence, a commit value $\hat{c}$ issued by the adversary for an instance of Bob is either a real output by H and can only be opened in a single way, or no output from H . In the latter case, we can consider $(\perp, \perp, \perp, \hat{c})$ as a new entry in the H list and count it as an extra oracle call. This way, $\hat{c}$ can never be opened. Hence, with probability at least $1 - (q+1)(q+2)2^{-\ell_e-1} - (q+1)(q+2)2^{-\ell_c-1}$, which is larger than $1 - \varepsilon$, the commit value(s) $\hat{c}$ by the adversary are either openable in a single fixed way or not openable. If they are not openable, the adversary fails. If openable $\hat{c}$ are issued by an oracle call to H by the adversary, we can thus virtually replace the adversary release of $\hat{c}$ by an adversary release of $\hat{K}$ and step B_2 can be ignored. If openable $\hat{c}$ are issued by other oracle calls to H , it can only be by a simulation of Alice, leading us to $c = \hat{c}$, thus $\hat{K} = K$ and $m_A = \hat{m}_A$.

Cases Bob-Bob. We number 2 the instance of Bob whose B_1 step is the last one. Those cases produce no oracle calls to H by Alice, so $\hat{K}^1$ and $\hat{K}^2$ are selected by the adversary before the B_1^2 step. Note that R^1 is already released. The attack succeeds if $R^2 = R^1 \oplus h_{\hat{K}^1}(m_B^1) \oplus h_{\hat{K}^2}(m_B^2)$ where R^2 is independent of the righthand term and selected at random by the second Bob. Clearly, this succeeds with probability 2^{-p} .

Cases Alice-Bob. Without loss of generality, we can assume that B_2 is the last step.

In cases $A_1A_2B_1B_2$, R is selected in step B_1 so the adversary succeeds with probability 2^{-p} .

In cases $A_1B_1A_2B_2$ with $c \neq \hat{c}$ or in cases $B_1A_1A_2B_2$ (necessarily with $c \neq \hat{c}$), the adversary has no information about K until step A_2 and succeeds when $\hat{R} \oplus R \oplus h_{\hat{K}}(m_B) = h_K(\hat{m}_B)$. Hence succeeds with probability at most $2^{-p} + \varepsilon_1$.

In cases $A_1B_1A_2B_2$ with $c = \hat{c}$, we must have $m_A = \hat{m}_A$. This can only be an attack for $m_B \neq \hat{m}_B$. The adversary has no information about K until step A_2 and succeeds when $\hat{R} \oplus R = h_K(m_B) \oplus h_K(\hat{m}_B)$, hence with probability at most $2^{-p} + \varepsilon_2$. $\square$

With the same analysis as in [Vau05], in a network of N participants, each limited to R runs of the protocol, and a maximal attack probability *at large* p , we should use $\rho \approx \log_2 \frac{N^2 R^2}{2p}$. When p is the probability to attack a target node, we should use $\rho \approx \log_2 \frac{NR^2}{2p}$. With $N \approx 2^{20}$, $R \approx 2^{10}$, and $p \approx 2^{-10}$, we obtain $\rho \approx 49$. In an ATM-like environment, we can take $N = 2$, $R = 3$, and $p = 3 \cdot 10^{-4}$, leading us to $\rho \approx 15$. In between, we believe that $\rho = 20$ bits provides enough security in a small community of human users.

7 Conclusion

We have shown how to construct efficient SAS-based AKA protocols based on existing ones and SAS-based MMA or MCA protocols. We have proposed a new 3-move MMA protocol using a generic commitment scheme. It can make a secure and efficient SAS-based AKA protocol with 4 moves over the insecure channel. We have also proposed a new 3-move MCA protocol using random oracle commitments. It can make a secure and efficient SAS-based AKA protocol with 3 moves in the random oracle model. For both constructions, we can have e.g. a SAS of 20 bits. Note that our two constructions use the same authenticated strings in both directions.

Applications of such protocols can be traditional key agreement, but run in an ad-hoc way. For instance, it can be used to exchange PGP public keys to be authenticated by a human-to-human telephone conversation. It can also be used to secure peer-to-peer VoIP communications. Other straightforward applications can be the Bluetooth-like establishment of symmetric key between associated wireless devices, e.g. for wireless USB.

References

- [BR93] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO '93: 13th Annual International Cryptology Conference*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249, Santa Barbara, California, U.S.A., 1993. Springer-Verlag.
- [ČČH06] Mario Čagalj, Srdjan Čapkun, and Jean-Pierre Hubaux. Key agreement in peer-to-peer wireless networks. *Proceedings of the IEEE, Special Issue in Security and Cryptography*, 94(2):467–478, 2006.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology revisited (preliminary version). In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 209–218, New York, NY, USA, May 1998. ACM Press.
- [CS02] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT '02: International Conference on the Theory and Applications of Cryptographic Techniques*, volume 2332 of *Lecture Notes in Computer Science*, Amsterdam, The Netherlands, April 2002. Springer-Verlag.
- [DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, November 1976.
- [DSS00] Digital signature standard (DSS). Federal Information Processing Standard, Publication 186-2, U.S. Department of Commerce, National Institute of Standards and Technology, 2000.

- [GMN04] Christian Gehrmann, Chris J. Mitchell, and Kaisa Nyberg. Manual authentication for wireless devices. *RSA Cryptobytes*, 7(1):29–37, January 2004.
- [GN01] Christian Gehrmann and Kaisa Nyberg. Enhancements to Bluetooth baseband security. In *Nordsec '01*, Copenhagen, Denmark, November 2001.
- [GN04] Christian Gehrmann and Kaisa Nyberg. Security in personal area networks. *Security for Mobility*, pages 191–230, 2004.
- [Hoe04] Jaap-Henk Hoepman. The ephemeral pairing problem. In Ari Juels, editor, *Financial Cryptography: the 8th International Conference (FC '04)*, volume 3110 of *Lecture Notes in Computer Science*, pages 212–226, Key West, FL, USA, February 2004. Springer-Verlag.
- [Kra94] Hugo Krawczyk. LFSR-based hashing and authentication. In Yvo Desmedt, editor, *Advances in Cryptology – CRYPTO '94: 11th Annual International Cryptology Conference*, volume 839 of *Lecture Notes in Computer Science*, pages 129–139, Santa Barbara, California, U.S.A., August 1994. Springer-Verlag.
- [LAN05] Sven Laur, N. Asokan, and Kaisa Nyberg. Efficient mutual data authentication using manually authenticated strings. Cryptology ePrint Archive, Report 2005/424, 2005. <http://eprint.iacr.org/>.
- [MY04] Philip MacKenzie and Ke Yang. On simulation-sound trapdoor commitments. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EURO-CRYPT '04 : International Conference on the Theory and Applications of Cryptographic Techniques*, volume 3027 of *Lecture Notes in Computer Science*, pages 382–400, Interlaken, Switzerland, May 2004. Springer-Verlag.
- [PV06] Sylvain Pasini and Serge Vaudenay. An optimal non-interactive message authentication protocol. In David Pointcheval, editor, *Topics in Cryptology – CT-RSA '06: The Cryptographers' Track at the RSA Conference 2006*, volume 3860 of *Lecture Notes in Computer Science*, pages 280–294, San Jose, California, U.S.A., February 2006. Springer-Verlag.
- [RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, Februar 1978.
- [Sti91] Douglas Stinson. Universal hashing and authentication codes. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO '91: 11th Annual International Cryptology Conference*, volume 576 of *Lecture Notes in Computer Science*, pages 74–85, Santa Barbara, California, U.S.A., August 1991. Springer-Verlag.
- [Sti94] Douglas Stinson. Universal hashing and authentication codes. *Designs, Codes and Cryptography*, 4:369–380, 1994.
- [Vau05] Serge Vaudenay. Secure communications over insecure channels based on short authenticated strings. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO '05: The 25th Annual International Cryptology Conference*, volume 3621 of *Lecture Notes in Computer Science*, pages 309–326, Santa Barbara, California, U.S.A., August 2005. Springer-Verlag.
- [Vau06] Serge Vaudenay. On Bluetooth repairing: Key agreement based on symmetric-key cryptography. In Dengguo Feng, Dongdai Lin, and Moti Yung, editors, *Information Security and Cryptology: First SKLOIS Conference, CISC'05*, volume 3822 of *Lecture Notes in Computer Science*, pages 1–9, Beijing, China, December 2006. Springer-Verlag.