A Simulation-Based Treatment of Authenticated Message Exchange*

Klaas Ole Kürtz, Henning Schnoor, and Thomas Wilke

Christian-Albrechts-Universität zu Kiel, 24098 Kiel, Germany
{kuertz,schnoor,wilke}@ti.informatik.uni-kiel.de

Abstract. Simulation-based security notions for cryptographic protocols are regarded as highly desirable, primarily because they admit strong composability and, consequently, a modular design. In this paper, we give a simulation-based security definition for two-round authenticated message exchange and show that a concrete protocol, 2AMEX-1, satisfies our security property, that is, we provide an ideal functionality for two-round authenticated message exchange and show that 2AMEX-1 realizes it securely. To model the involved public-key infrastructure adequately, we use a joint-state approach.

1 Introduction

Simulation-based security definitions for cryptographic protocols, see, for instance, [1,2,3,4], are attracting much attention, the reasons being that such security definitions "guarantee security even when a secure protocol [...] is used as a component of an arbitrary system" [1] and that they enable "modular proofs of security" [2]. As a consequence, a variety of cryptographic primitives such as asymmetric encryption and digital signatures have been treated following the simulation-based approach. There are, however, only few complex cryptographic protocols that have been tackled within the simulation-based framework. We are aware of [5,6,7,8,9], where, for instance, Kerberos and the Yahalom protocol are treated.

In this paper, we deal with two-round authenticated message exchange protocols following the simulation-based approach. We (i) provide an ideal functionality for two-round authenticated message exchange protocols, $\mathcal{F}_{\mathrm{2AM}}$, (ii) provide an implementation, $\mathcal{P}_{\mathrm{2AMEX-1}}$, corresponding to a particular such protocol, 2AMEX-1, and (iii) prove the implementation of 2AMEX-1 to be secure, that is, prove that $\mathcal{P}_{\mathrm{2AMEX-1}}$ securely realizes the ideal functionality, in symbols $\mathcal{P}_{\mathrm{2AMEX-1}} \leq^{\mathrm{BB}} \mathcal{F}_{\mathrm{2AM}}$. (The superscript stands for black-box simulatability.)

The protocol 2AMEX-1, see [10], which is a generic protocol for message authentication in a web service setting, is complex in several respects: it distinguishes between short-lived clients and long-lived servers; it uses digital signatures and therefore makes use of a public-key infrastructure; it requires only bounded memory; it uses nonces and timestamps to counter replay attacks; each client and each

* This work was partially supported by the DFG under grant KU 1434/4-2.

A. Datta (Ed.): ASIAN 2009, LNCS 5913, pp. 109–123, 2009.
© Springer-Verlag Berlin Heidelberg 2009

server has its own local clock. In [10], 2AMEX-1 was proved to be secure in the Bellare-Rogaway framework as presented in [11].

Several simulation-based approaches have been developed over the last decade (see above). We could have used any of these approaches, but we have adopted the one by Küsters, see [4], because it provides a very flexible addressing mechanism and easy-to-use joint-state theorems, see [12]. The latter is especially useful in the analysis of 2AMEX-1, because it allows us to show with only little effort that 2AMEX-1 works securely with a simple, but realistic public-key infrastructure. Although Küsters' setting comes in handy in many respects, it also has some shortcomings, which become evident from our analysis and are discussed in this paper.

Due to the space limit, we can only provide a high-level account of Küsters' model, provide a brief description of the functionalities we have developed, and give a short sketch of the proof of our main result. A full version of this paper can be found in [13]. We start with the sketch of Küsters' model in Section 2, go on with a description of the setting and the ideal functionalities in Section 3 and a description of the implementation for 2AMEX-1 in Section 4, and conclude with our main result and a discussion in Sections 5 and 6.

We are grateful to Max Tuengerthal for helpful comments.

2 Simulation-Based Security

In this section, we give a high-level description of the simulation-based framework from [4], which is referred to as the *IITM framework*, where IITM stands for *inexhaustible interactive Turing machine*.

In the IITM framework cryptographic protocols and the environment they are run in (including the adversary) are modeled as concurrent, polynomial-time, probabilistic, interactive, replicable Turing machines. Here, "concurrent" refers to an interleaving semantics, that is, only one IITM is active at a time and there is a mechanism that determines which IITM is activated next; "replicable" refers to a mechanism which allows certain machines, the so-called banged machines, to be instantiated several times (and run concurrently); "interactive" means that the machines can communicate by sharing tapes, more precisely: an output tape of one machine can be the input tape of another machine. From a security point of view, it is important that systems of IITM's can be simulated in polynomial time. To achieve this, it is, however, not enough to require that the individual IITM's are polynomial-time, because two IITM's "playing ping pong" could double their outputs on each activation, leading to an overall exponential running time. For that reason the IITM framework imposes certain restrictions on how machines are interconnected, based on a partition of tapes into consuming and enriching. Roughly speaking, the overall length of the output of one IITM up to a certain point may be polynomial in the overall length of the input on enriching tapes up to the same point, but there must not be any cycle of enriching tapes. This is less restrictive than requiring that each IITM runs in time polynomial in the security parameter; it allows to process inputs of arbitrary size.

To illustrate the IITM framework consider Figure 1 and first focus on the box labeled $\mathcal{F}_{2\mathrm{AM}}$. This box represents a model of two-round authenticated message exchange protocols (details follow in the next section); it contains four machines which represent an actual protocol: C, S, EI, and NG, of which the first three are banged (can be replicated), and the last one is not. Every instance of machine C is connected with machine NG, in both directions. The corresponding input tape of NG is enriching, while the input tape of C is not.

There are two types of connections crossing the borders of $\mathcal{F}_{2\mathrm{AM}}$: solid connections representing tapes classified as I/O tapes and dashed connections representing tapes classified as network tapes. I/O tapes should roughly be thought of as tapes communicating with "users" of the system, whereas network tapes are tapes where the adversary can interfere.

In Figure 1, the adversary, represented by an IITM denoted $\mathcal{A}$, is not connected directly with $\mathcal{F}_{2\mathrm{AM}}$. Rather, there is a mediator between $\mathcal{A}$ and $\mathcal{F}_{2\mathrm{AM}}$, namely an IITM $\mathcal{S}$ called simulator. The situation is typical for simulation-based security: a simulator "translates" network traffic to make a system (in this case $\mathcal{F}_{2\mathrm{AM}}$) seem equivalent to another one (usually a "real" system $\mathcal{P}$, see below) to an outside observer consisting of an environment machine $\mathcal{E}$ (taking over the role of all users) and an adversary $\mathcal{A}$.

Another feature of Figure 1 not discussed yet has to do with how different instances of the same machine are addressed. Underlining the name of a machine indicates the usage of a generic addressing mechanism provided by the IITM framework, which works by using prefixes of messages as identifiers for instances.

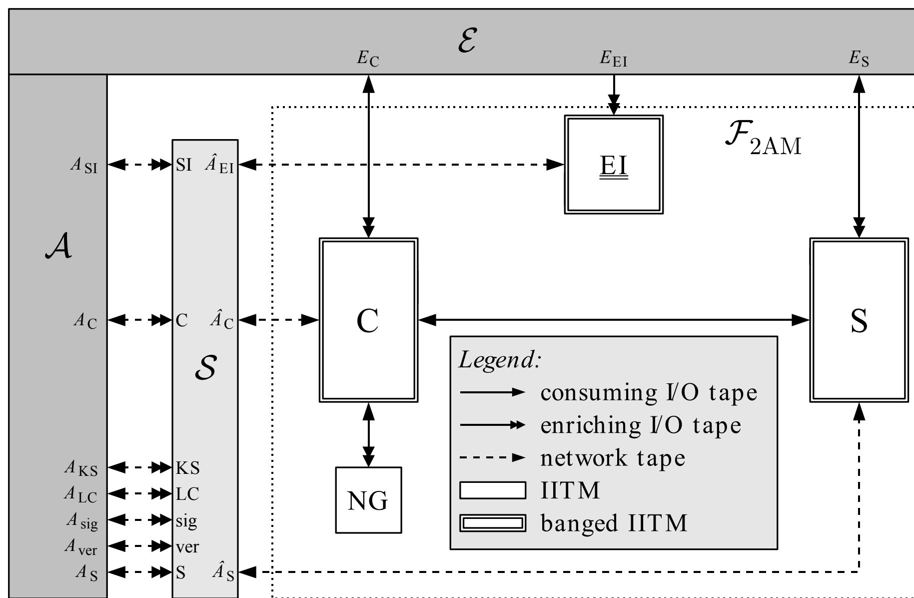

Fig. 1. Ideal functionality for two-round message authentication

In Figure 1 the machine EI is underlined twice, which adds two prefixes for addressing, that is, a hierarchical addressing mechanism is used. We use it to model multi-user multi-session instances.

The formal way to specify the system represented by the box $\mathcal{F}_{2\mathrm{AM}}$ in Figure 1 is by the expression

$$\mathcal{F}_{2\mathrm{AM}} = !\mathcal{F}_{\mathrm{C}} \mid !\mathcal{F}_{\mathrm{S}} \mid \mathcal{F}_{\mathrm{NG}} \mid !\underline{\underline{\mathcal{F}_{\mathrm{EI}}}} \ , \tag{1}$$

where $\mathcal{F}_{\mathrm{C}}$, $\mathcal{F}_{\mathrm{S}}$, $\mathcal{F}_{\mathrm{NG}}$, and $\mathcal{F}_{\mathrm{EI}}$ denote (descriptions of) the underlying IITM's, and $\mid$ denotes an operator for composing machines.

In the IITM framework, security of a protocol is defined as follows. First, one describes a system of IITM's, $\mathcal{F}$, which works in an ideal fashion in every setting where an environment and an adversary are connected to it, that is, how one would expect a perfect protocol to work. Such a system is called an ideal functionality. Then, given a real protocol, one describes a system of IITM's, $\mathcal{P}$, which works just the way the real protocol would work in every setting where an environment and an adversary are connected to it. Now, $\mathcal{P}$ is considered secure if there is a simulator IITM $\mathcal{S}$ with the following property. For every environment machine $\mathcal{E}$ and every adversary machine $\mathcal{A}$, the system composed of $\mathcal{P}$, $\mathcal{E}$, and $\mathcal{A}$ is computationally indistinguishable from the system composed of $\mathcal{F}$, $\mathcal{E}$, $\mathcal{A}$, and $\mathcal{S}$. As explained above communication between these machines is restricted as follows: all external network connections of $\mathcal{F}$ are handled by the simulator $\mathcal{S}$; the adversary may only communicate with $\mathcal{F}$ using the network interface provided by the simulator; and the environment may only communicate with $\mathcal{F}$ using I/O connections. Hence, the system composed of $\mathcal{F}$ and the simulator (translating network messages) is "equivalent" to $\mathcal{P}$. In other words, every attack on the real protocol can be transferred into the ideal system.

If the above condition is satisfied, then $\mathcal{P}$ securely realizes (or implements) $\mathcal{F}$, denoted by $\mathcal{P} \leq^{\mathrm{BB}} \mathcal{F}$ (for *black-box simulation*).

3 Two-Round Authenticated Message Exchange

We start with a description of the general scenario. In a session of a two-round authenticated message exchange protocol (2AM protocol) a client sends a request to a server and expects to receive an appropriate response. This is, for instance, the case for web service calls, see, e. g., [14,15] and remote procedure calls, see, e. g. [16,17]. Observe that for these protocols to make sense the request and response messages include payloads.

In a 2AM protocol the request and the response messages are required to be secured in such a way that (i) both client and server can verify that the messages they receive are authentic, (ii) the server accepts no message twice (payloads, on the contrary, may be received twice, but only in different messages), and (iii) if the client receives a response, it can be sure which of his requests the response refers to. Note that the same client may have multiple sessions with the same or different servers in parallel, but each session has only two rounds.

Tapes: C $\longleftrightarrow$ E_{C}, C $\dashrightarrow$ $\hat{A}_{\mathrm{C}}$, C $\longleftrightarrow$ S, C $\longleftrightarrow$ NG
Initialization: $c = s = r = \varepsilon$, $n = 0$, $state = \mathsf{Init}$, $cor = \mathsf{false}$
Steps: `loop`
Send a request to the server:
`if` $(c', (\mathsf{Client}, s'), \mathsf{Init})$ received from E_{C}
Let $state = \mathsf{OK}$, $c = c'$ and $s = s'$.
Send $(c, (\mathsf{Client}, s), \mathsf{GetNonce})$ to NG.
Recv $(c, (\mathsf{Client}, s), \mathsf{Nonce}, r')$ from NG, let $r = r'$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Nonce}, r)$ to E_{C}.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Request}, p_c, 1^{n'})$ from E_{C}, let $n = n'$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Request}, p_c, n)$ to $\hat{A}_{\mathrm{C}}$.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Request}, \mathsf{Send})$ from $\hat{A}_{\mathrm{C}}$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Request}, p_c)$ to S.
Receive and process a response from the server:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Response}, p_s)$ received from S
If $state \neq \mathsf{OK}$ or $|p_s| > n$, abort.
Let $state = \mathsf{Stopped}$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Response}, p_s)$ to E_{C}.
Corruption: $\mathtt{Corr}(cor, \mathsf{true}, state \neq \mathsf{Init}, \varepsilon, \hat{A}_{\mathrm{C}}, \{E_{\mathrm{C}}\}, E_{\mathrm{C}})$
CheckAddress: Accept the initialization message only once. Check for c, s, and r as soon as each one has been set.

Fig. 2. The client functionality $\mathcal{F}_{\mathrm{C}}$

3.1 Overview of the Ideal Functionality

Our model of the ideal functionality for 2AM protocols consists of four functionalities, see Figure 1: a client $\mathcal{F}_{\mathrm{C}}$ (defined in Figure 2), a server $\mathcal{F}_{\mathrm{S}}$ (defined in Figure 3), a nonce generator $\mathcal{F}_{\mathrm{NG}}$, and an enriching input functionality $\mathcal{F}_{\mathrm{EI}}$. The ideal functionality $\mathcal{F}_{\mathrm{2AM}}$ is the composition of these functionalities, as defined in (1).

One instance of the client functionality handles exactly one session between a client identity and a server, i. e., after initialization it basically (i) receives a request from the environment and encapsulates it in a message to a server, and (ii) receives a response from the server and forwards its contents to the environment. One instance of the server functionality also handles exactly one session; as with the client, it consists of receiving a request and sending a response. The nonce generator generates globally unique session identifiers (numbers used once, nonces) to distinguish multiple sessions between two parties. The enriching input functionality passes bits from an enriching input tape to the adversary. These bits are necessary to give the adversary additional capabilities as explained in Section 4.3.

3.2 Ideal Client Functionality

When the environment wants to start a new session, it provides the client with the identity of a server the client is supposed to communicate with. The client then responds with a nonce, which can be viewed as a handle, i. e., it allows the environment to distinguish different sessions this client is involved in.

The environment can now pass the payload of the request message to the client as well as enough resources to process a possible response from the server. The client then notifies the adversary that a message is ready to be sent. If the

Tapes: $\mathrm{S} \longleftrightarrow E_{\mathrm{S}}$, $\mathrm{S} \dashleftarrow\dashrightarrow \hat{A}_{\mathrm{S}}$, $\mathrm{S} \longleftrightarrow \mathrm{C}$
Initialization: $s = c = r = p_s = \varepsilon$, $n = 0$, $state = \mathsf{Init}_0$, $cor = \mathsf{false}$
Steps: `loop`
Initialization by the environment:
`if` $(s', (\mathsf{Server}), \mathsf{Init}, 1^{n'})$ received from E_{S}
If $state \neq \mathsf{Init}_0$, abort. Let $s = s'$ and $n = n'$.
Send $(s, (\mathsf{Server}), \mathsf{Init}, n)$ to $\hat{A}_{\mathrm{S}}$.
Recv $(s, (\mathsf{Server}), \mathsf{Init}, \mathsf{OK})$ from $\hat{A}_{\mathrm{S}}$.
Let $state = \mathsf{Init}_1$.
Receive and process a request from the client:
`if` $(c', (\mathsf{Client}, s, r'), \mathsf{Request}, p_c)$ received from C
If $state \neq \mathsf{Init}_1$ or $|p_c| > n$, abort. Let $state = \mathsf{OK}$, $c = c'$, and $r = r'$.
Send $(s, (\mathsf{Server}, c, r), \mathsf{Request}, p_c)$ to E_{S}.
Receive a response payload from the environment:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Response}, p)$ received from E_{S}
Let $p_s = p$. Send $(s, (\mathsf{Server}, c, r), \mathsf{Response}, p_s)$ to $\hat{A}_{\mathrm{S}}$.
Deliver a response to the client:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Response}, \mathsf{Send})$ received from $\hat{A}_{\mathrm{S}}$ and not cor
If $state \neq \mathsf{OK}$, abort. Let $state = \mathsf{Stopped}$.
Send $(s, (\mathsf{Server}, c, r), \mathsf{Response}, p_s)$ to C.
Send an error message to the environment:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Response}, \mathsf{Error})$ received from $\hat{A}_{\mathrm{S}}$
Send $(s, (\mathsf{Server}, c, r), \mathsf{Response}, \mathsf{Error})$ to E_{S}.
Corruption: $\mathtt{Corr}(cor, \mathsf{true}, state \neq \mathsf{Init}_0, \varepsilon, \hat{A}_{\mathrm{S}}, \{E_{\mathrm{S}}\}, E_{\mathrm{S}}, s)$
CheckAddress: Accept the initialization message only once. Check for s, c, and r as soon as each one has been set.

Fig. 3. The server functionality $\mathcal{F}_{\mathrm{S}}$

adversary (ever) allows the transfer, the message is written to the incoming tape of the server. This models the adversary's ability to delay or drop messages on the network.

When the server transfers a response (which is not too large), the client simply unwraps it and forwards the contents to the environment. The details are spelled out in Figure 2.

A special mode of computation of IITM's, CheckAddress, is used in the last line of IITM definitions like Figure 2 to determine whether an incoming message is addressed to the current instance of the client IITM. If a message is rejected by all running instances, a new instance of the client IITM is started since the client IITM is banged in $\mathcal{F}_{\mathrm{2AM}}$. In addition, we use the corruption macro `Corr` from [12] (with a slightly extended addressing mechanism) to allow a uniform treatment of corruption of clients and servers in both the ideal and the real world, see Appendix A.

3.3 Ideal Server Functionality

To start a session on the server side, the environment sends a message to the server with the identity it is supposed to receive messages for and the maximal length of an incoming request message.

Upon receiving a request from a client, the server unwraps it and forwards the request payload to the environment. Now the environment can respond by passing a response payload to the server functionality. The server asks the adversary, who has three options: It can either approve the sending of the payload,

in which case the server delivers the message directly to the client. Secondly, the adversary can ignore the response, in which case the server sends no message at all. Thirdly, the adversary can also explicitly deny processing the payload, which results in an error message being sent to the environment.

The first two options again model that the adversary may intercept and delay network traffic. The third type of reaction models that in our implementation the server may reject messages due to bounded memory and notify the environment of the rejection.

4 Implementation of the 2AMEX-1 Protocol

In this section, we describe a system of IITM's implementing the 2AMEX-1 protocol, which is a 2AM protocol in the above sense and described in detail in [10]. First, we give an informal introduction into the protocol.

4.1 The Protocol 2AMEX-1

In 2AMEX-1, an authenticated message exchange between a client with identity c and a server with identity s works roughly as follows.

1. a) c is asked by the environment to send the request p_c
   b) c sends $\{(\mathsf{From}\colon c, \mathsf{To}\colon s, \mathsf{MsgID}\colon r, \mathsf{Time}\colon t, \mathsf{Body}\colon p_c)\}_{sk_c}$ to s
   c) s checks whether the message is admissible and if not, stops
   d) s forwards the request (r, p_c) to the environment
2. a) s receives a response (r, p_s) from the environment
   b) s checks whether the response is admissible and if not, stops
   c) s sends $\{(\mathsf{From}\colon s, \mathsf{To}\colon c, \mathsf{Ref}\colon r, \mathsf{Body}\colon p_s)\}_{sk_s}$ to c
   d) c checks whether the message is admissible and if not, stops
   e) c forwards the response p_s to the environment

Here, r is the nonce as described in the previous section, which is also used as a handle by the server (see steps 1. d) and 2. a)), t is the value of a local clock of the client, p_c is the payload the client sends, p_s is the payload the server returns, and $\{\cdot\}_{sk_c}$ and $\{\cdot\}_{sk_s}$ stand for signing the message by the client and server, respectively. Repeating the message id of the request allows the client to verify that p_s is indeed a response to the request p_c.

The interesting parts are steps 1. c) and 2. b). We assume that there is a constant $\mathrm{cap}_s > 0$, the so-called capacity of the server, and a constant tol_s^+ that indicates its tolerance with respect to inaccurate clocks. At all times the server keeps a time $t_s^{\min}$ and a finite list L of triples (t, r, c) of pending and handled requests. At the beginning or after a reset, $t^{\min}$ is set to $t_s + \mathrm{tol}_s^+$, where t_s is a timestamp retrieved from the local clock functionality, and L is set to the empty list.

Step 1. c). Upon receiving a message as above, the server s rejects if $t \notin [t_s^{\min} + 1, t_s + \mathrm{tol}_s^+]$ or if $(t', r, c') \in L$ for some t' and c', and otherwise proceeds

as follows: If L contains less than cap_s elements, it inserts (t, r, c) into L. Otherwise, the server deletes all tuples containing the oldest timestamp from L, until L contains less than cap_s tuples. Then it sets $t_s^{\min}$ to the timestamp contained in the last tuple deleted from L, and finally inserts (t, r, c) into L.

Step 2. b). When asked to send a payload p_s with message handle r, the server rejects if there is no triple $(t, r, c) \in L$ with $c \neq \varepsilon$. If it does not reject, it updates L by overwriting c with ε in the tuple (t, r, c) to ensure that the service cannot respond to the same message twice.

4.2 Implementation in the IITM Model

We will now describe the system of IITM's defined by

$$\mathcal{P}_{2\text{AMEX}-1} = !\mathcal{P}_{\text{C}} \mid !\mathcal{P}_{\text{S}} \mid !\underline{\underline{\mathcal{F}_{\text{Sig}}}} \mid !\underline{\underline{\mathcal{P}_{\text{SI}}}} \mid !\underline{\underline{\mathcal{F}_{\text{KS}}}} \mid !\underline{\underline{\mathcal{F}_{\text{LC}}}} \quad (2)$$

and illustrated in Figure 4, which implements the 2AMEX-1 protocol.

In (2), $\mathcal{P}_{\text{C}}$ is the client-side part of the protocol (defined in Figure 5), $\mathcal{P}_{\text{S}}$ is the server-side part of the protocol (defined in Appendix B), $\mathcal{F}_{\text{Sig}}$ is the signature functionality as defined in [18], $\mathcal{P}_{\text{SI}}$ is an interface which allows the adversary to access the signature functionality with few restrictions, $\mathcal{F}_{\text{KS}}$ is an ideal functionality of a trusted key store, and $\mathcal{F}_{\text{LC}}$ models a local clock which is controlled by the adversary, i. e. not synchronized with the clocks of other parties and not even monotone.

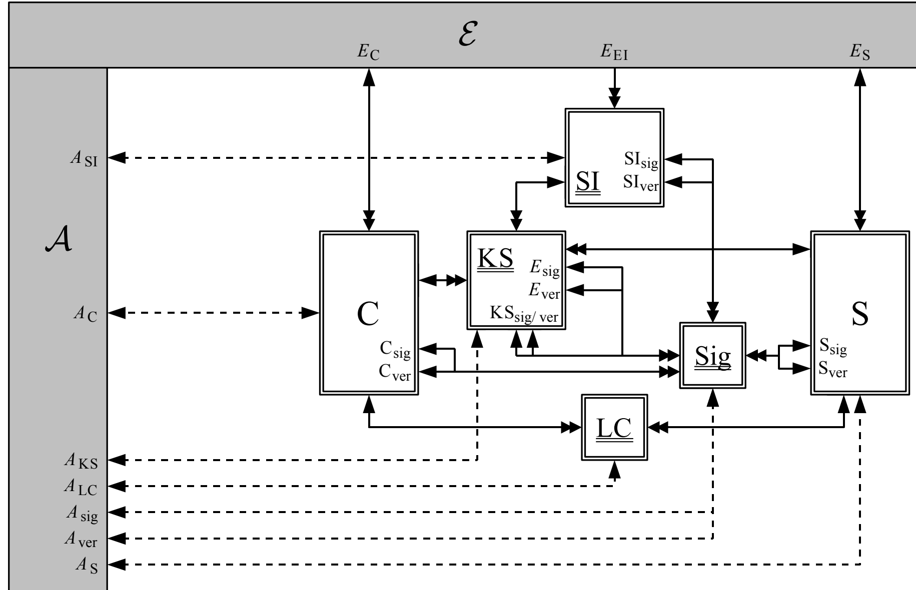

Fig. 4. Overview of 2AMEX-1 protocol implementation

4.3 Signatures and the Public Key Infrastructure

We model the digital signatures that 2AMEX-1 uses by the ideal functionality $\mathcal{F}_{\mathrm{Sig}}$ from [18], which was proved to be securely implementable using any existentially unforgeable signature scheme.

We give the adversary access to the signature scheme and allow him to sign any bit string that does not have the format of a 2AMEX-1 message. This models that our protocol does not have exclusive access to the keys used to sign the messages. For example, the same key can be used to sign a 2AMEX-1 message and parts of the payload contained in that message. This is realized by the signature interface functionality $\mathcal{P}_{\mathrm{SI}}$, which accepts requests from the adversary to (i) sign messages that do not have the format of 2AMEX-1 messages and (ii) verify arbitrary signatures. In $\mathcal{P}_{\mathrm{2AMEX-1}}$, the signature interface functionality is banged in the multi-user multi-session version, effectively meaning that the adversary has access to all keys used in the protocol.

As the signature interface needs resources from the environment to sign messages for the adversary, it has an enriching input tape E_{EI}. Its counterpart in the ideal system is a tape in the enriching input functionality EI.

To coordinate how different IITM's access a single instance of the signature functionality, we define the ideal functionality of a key store, $\mathcal{F}_{\mathrm{KS}}$, which allows clients, servers, and the signature interface functionality to retrieve trusted keys as well as the corruption status of that key. To be able to distribute the public key, $\mathcal{F}_{\mathrm{KS}}$ also initializes the instances of the signature functionality. The particular form of this functionality is due to the fact that we want to use $\mathcal{F}_{\mathrm{Sig}}$ from [18] as is. Nevertheless, one can implement $\mathcal{F}_{\mathrm{KS}}$ using standard techniques for building a public key infrastructure.

4.4 Client Implementation

The client protocol $\mathcal{P}_{\mathrm{C}}$ (see Figure 5) is a direct implementation of the ideal functionality $\mathcal{F}_{\mathrm{C}}$ with the following changes:

- The messages are transferred over the network (rather than exchanged directly between client and server). This is modeled by writing the messages on an external network tape.
- To secure the request message, the client signs it using a digital signature obtained from an instance of $\mathcal{F}_{\mathrm{Sig}}$ for this session. The server will be able to obtain the public key from the according key store and verify the signature.
- When receiving a response from the server, the signature of that message is verified by the client in the same way.
- The nonces are not generated by a centralized entity, but randomly chosen locally by each client. While this does not guarantee that the numbers are unique, the probability of a collision is negligible if the length of the nonces grows linearly with the security parameter.
- The request message is additionally secured by a timestamp. The client uses the local clock functionality $\mathcal{F}_{\mathrm{LC}}$ to obtain a timestamp.

Tapes: C $\leftarrow\!\!\!\rightarrow$ E_{C}, C $\leftarrow\!\!-\!\!-\!\!\rightarrow$ A_{C}, C $\longleftrightarrow$ KS, C $\longleftrightarrow$ LC, $\mathrm{C_{sig}}$ $\longleftrightarrow$ Sig, $\mathrm{C_{ver}}$ $\longleftrightarrow$ Sig
Initialization: $c = s = r = \varepsilon$, $n = 0$, $state = \mathsf{Init}$, $cor = \mathsf{false}$
Steps: `loop`
Send a request to the server:
`if` $(c', (\mathsf{Client}, s'), \mathsf{Init})$ received from E_{C}
If $state \neq \mathsf{Init}$, abort. Let $c = c'$ and $s = s'$.
Generate an η-bit nonce r randomly, where η is the security parameter.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Nonce}, r)$ to E_{C}.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Request}, p_c, 1^{n'})$ from E_{C}, let $n = n'$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{GetKey})$ to KS.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{PublicKey}, k_c)$ from KS.
Send $(c, (\mathsf{Client}, s, r), \mathsf{GetTime})$ to LC.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Time}, t)$ from LC.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Corrupted?})$ to KS.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Corrupted}, cor')$ from KS. If cor', abort.
Let $m_c = (\mathsf{From}\colon c, \mathsf{To}\colon s, \mathsf{MsgID}\colon r, \mathsf{Time}\colon t, \mathsf{Body}\colon p_c)$.
Send $(c, (\mathsf{Client}, s, r), \mathsf{Sign}, m_c)$ on $\mathrm{C_{sig}}$.
Recv $(c, (\mathsf{Client}, s, r), \mathsf{Signature}, \sigma_c)$ on $\mathrm{C_{sig}}$. Let $state = \mathsf{OK}$.
Send (m_c, σ_c) to A_{C}.
Receive and process a response from the server:
`if` (m_s, σ_s) received from A_{C} with $m_s = (\mathsf{From}\colon c, \mathsf{To}\colon s, \mathsf{Ref}\colon r, \mathsf{Body}\colon p_s)$
If $state \neq \mathsf{OK}$ or cor or $|p_s| > n$, abort.
Let $n = n - |p_s|$.
Send $(s, (\mathsf{Server}, c, r), \mathsf{GetKey})$ to KS.
Recv $(s, (\mathsf{Server}, c, r), \mathsf{PublicKey}, k_s)$ from KS.
Send $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Init})$ on $\mathrm{C_{ver}}$.
Recv $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Init})$ on $\mathrm{C_{ver}}$.
Send $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Corrupted?})$ to KS.
Recv $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Corrupted}, cor')$ from KS. If cor', abort.
Send $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Verify}, m_s, \sigma_s, k_s)$ on $\mathrm{C_{ver}}$.
Recv $(s, (\mathsf{Server}, c, r), \mathsf{Client}, \mathsf{Verified}, b)$ from on $\mathrm{C_{ver}}$, if $b \neq 1$, stop.
Let $state = \mathsf{Stopped}$ and send $(c, (\mathsf{Client}, s, r), \mathsf{Response}, p_s)$ to E_{C}.
Corruption: $\mathtt{Corr}(cor, \mathsf{true}, state \neq \mathsf{Init}, \varepsilon, A_{\mathrm{C}}, \{E_{\mathrm{C}}\}, E_{\mathrm{C}}, c, (\mathsf{Client}, s, r))$
CheckAddress: Check for c, s, and r as soon as each one has been set.

Fig. 5. The client protocol $\mathcal{P}_{\mathrm{C}}$

– Before using a signature functionality to sign or verify a message, the client checks if the signature or the verification functionality is corrupted. If either one is, the client aborts.

4.5 Server Implementation

The implementation $\mathcal{P}_{\mathrm{S}}$ of the server functionality (see Appendix B) is more complicated than the client. To be able to counteract replay attacks, one single IITM handles all sessions. That is, for each identity s all communication of that identity in the server role is handled by one single instance of $\mathcal{P}_{\mathrm{S}}$.

Therefore, the server maintains two lists: R stores resources passed by the environment (corresponding to the fact that in the ideal system, each session of the server is started by the environment), while L (corresponding to L described in Section 4.1) is used to store information from request messages received so far by this server. During initialization, i. e., when receiving the first message, the server asks the adversary to provide values for two parameters of the 2AMEX-1 protocol, namely the capacity cap_s and the tolerance tol_s^+.

When receiving a message from the client, the server (i) tries to retrieve the client's key, (ii) obtains the current time from $\mathcal{F}_{\mathrm{LC}}$ (and checks for monotonicity of the clock), (iii) verifies the signature, (iv) checks if a message with the

same nonce has already been accepted (i. e. the nonce is in L), (v) checks if the timestamp is in order (i. e. not too old and not too new), and (vi) forwards the message to the environment if everything is in order. If some step fails, the server simply drops the message.

When the environment wants to reply to a message, the server first checks if the nonce is valid (i. e. occurs in L), else it sends an error message to the environment. This is important as the nonce may have been removed from L due to capacity reasons without notification to the environment. Then, the server initializes its instance of the signature scheme for this session, signs the message, and writes it on an external network tape.

Note that during the steps to process a request or a response, the control may be passed to the adversary by some of the ideal functionalities the server uses. Hence, the execution of the steps when processing a request or response may be interrupted by the adversary (e. g., by sending another incoming message to this server). As soon as a message is received that is not related to processing the current message, the processing of the current message is aborted by the server and cannot be resumed later.

5 Results

Our result states that our protocol securely realizes the ideal functionality $\mathcal{F}_{2\mathrm{AM}}$. The formal statement of the theorem is as follows:

Theorem 1

$$!\mathcal{F}_{\mathrm{C}} \mid !\mathcal{F}_{\mathrm{S}} \mid \mathcal{F}_{\mathrm{NG}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{EI}}}} \geq^{\mathrm{BB}} !\mathcal{P}_{\mathrm{C}} \mid !\mathcal{P}_{\mathrm{S}} \mid \underline{\underline{!\mathcal{P}_{\mathrm{SI}}}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{KS}}}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{Sig}}}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{LC}}}} \quad (3)$$

$$\geq^{\mathrm{BB}} !\mathcal{P}_{\mathrm{C}} \mid !\mathcal{P}_{\mathrm{S}} \mid \underline{\underline{!\mathcal{P}_{\mathrm{SI}}}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{KS}}}} \mid !\mathcal{P}^{\mathrm{JS}}_{\mathrm{Sig}} \mid \underline{!\mathcal{F}_{\mathrm{Sig}}} \mid \underline{\underline{!\mathcal{F}_{\mathrm{LC}}}} \quad (4)$$

The first of these inequalities states that the IITM realization of our protocol, when using an ideal signature functionality, realizes the system consisting of the ideal functionalities for $\mathcal{F}_{2\mathrm{AM}}$. The main part of the proof is the construction of the simulator, which essentially simulates our concrete protocol. It keeps track of the internal state of all involved parties and "translates" problems that can occur in the real protocol (e. g., the server running out of memory and therefore being unable to respond to old messages) into attacks on the ideal functionalities (the simulator may, e. g., prevent the ideal server functionality from responding to an old message).

Due to the way in which the ideal signature functionality is used, the realization of the protocol as stated in the first inequality is unrealistic, because for each message sent a new key for the signature scheme is generated. This can be avoided by applying a joint-state theorem [12,18] allowing different sessions to use the same key. Essentially, a "wrapper" $\mathcal{P}^{\mathrm{JS}}_{\mathrm{Sig}}$ managing different sessions is used to access the signature functionalities. The second inequality in Theorem 1 (which follows directly from [18]) makes use of this wrapper, so that instead of

one key per party and per session ($!\underline{\underline{\mathcal{F}_{\text{Sig}}}}$), there is only a single key for each party ($!\underline{\mathcal{F}_{\text{Sig}}}$), as in a realistic public-key infrastructure.

Theorem 1 gives a security treatment of a complex protocol in a simulation-based security setting: Our protocol features a long-lived server role, uses timestamps to prevent replay attacks, and accesses a public-key infrastructure for digital signatures. It is easy to see that long-livedness and timestamps are required to realize our ideal functionality with bounded memory (see [10]). It is interesting to note that while our ideal server functionality is short-lived, a realization necessarily needs to be long-lived; this is a particular property of authenticated message exchange with only two rounds.

6 Discussion

Simulation-based security clearly has the advantage that it leads to an easier statement of security than an individual, trace-based definition, and moreover, allows to treat protocols for very different tasks in a single model. The security properties obtained by such an analysis are quite strong and hold (via composition) in an arbitrary context. The IITM framework (and related frameworks) is designed to support modular protocol analysis.

However, these advantages come with a price when considering a concrete complex protocol. In [10], we presented a customized model (based on the seminal work by Bellare and Rogaway [11]) for proving security of 2AMEX-1. A comparison between that work and the current paper gives insights into the advantages and disadvantages of both approaches.

The formulation of both ideal functionalities and concrete implementations for authenticated message exchange in the current paper is rather long and unintuitive (the latter are significantly more complex than their counterparts in [10]). Both feature unnatural communication (bit strings to provide computing resources, status and activation messages exchanged sent to and received from the adversary and the environment), which are necessary due to how resources and activation are handled. Intuitively, one would like the environment to only access the "service" provided by the functionalities, but in the IITM framework, the environment additionally needs to provide resources for the involved parties that allow them to process the input.

Furthermore, the handling of corruption in the IITM framework is more complex and seems less natural than in the Bellare-Rogaway based model. Also, for the analysis of our protocol, the modular approach provided by the IITM framework does not simplify the security analysis, compared to the proof in [10]. Finally, the use of the joint-state theorem to enable realistic treatment of signatures results in a slightly different protocol from the one originally stated in [10] and from a realistic implementation.

It would be very interesting to know whether the IITM framework can be adapted to remove the above-mentioned difficulties.

References

1. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS, pp. 136–145 (2001)
2. Pfitzmann, B., Waidner, M.: A model for asynchronous reactive systems and its application to secure message transmission. In: IEEE Symposium on Security and Privacy, pp. 184–201 (2001)
3. Backes, M., Pfitzmann, B., Waidner, M.: A general composition theorem for secure reactive systems. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 336–354. Springer, Heidelberg (2004)
4. Küsters, R.: Simulation-based security with inexhaustible interactive Turing machines. In: CSFW, pp. 309–320. IEEE Computer Society, Los Alamitos (2006)
5. Canetti, R., Krawczyk, H.: Universally composable notions of key exchange and secure channels. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 337–351. Springer, Heidelberg (2002)
6. Moran, T., Naor, M.: Receipt-free universally-verifiable voting with everlasting privacy. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 373–392. Springer, Heidelberg (2006)
7. Backes, M., Cervesato, I., Jaggard, A.D., Scedrov, A., Tsay, J.K.: Cryptographically sound security proofs for basic and public-key Kerberos. In: Gollmann, D., Meier, J., Sabelfeld, A. (eds.) ESORICS 2006. LNCS, vol. 4189, pp. 362–383. Springer, Heidelberg (2006)
8. Backes, M., Pfitzmann, B.: On the cryptographic key secrecy of the strengthened Yahalom protocol. In: Fischer-Hübner, S., Rannenberg, K., Yngström, L., Lindskog, S. (eds.) SEC. IFIP, vol. 201, pp. 233–245. Springer, Heidelberg (2006)
9. Gajek, S., Manulis, M., Pereira, O., Sadeghi, A.R., Schwenk, J.: Universally composable security analysis of TLS. In: Baek, J., Bao, F., Chen, K., Lai, X. (eds.) ProvSec 2008. LNCS, vol. 5324, pp. 313–327. Springer, Heidelberg (2008)
10. Kürtz, K.O., Schnoor, H., Wilke, T.: Computationally secure two-round authenticated message exchange. Cryptology ePrint Archive, Report 2009/262 (2009), http://eprint.iacr.org/
11. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)
12. Küsters, R., Tuengerthal, M.: Joint state theorems for public-key encryption and digital signature functionalities with local computation. In: CSF, pp. 270–284. IEEE Computer Society, Los Alamitos (2008)
13. Kürtz, K.O., Schnoor, H., Wilke, T.: A simulation-based treatment of authenticated message exchange. Cryptology ePrint Archive, Report 2009/368 (2009), http://eprint.iacr.org/
14. Mitra, N., Lafon, Y.: SOAP version 1.2 part 0: Primer (second edition). Technical report, W3C (2007), http://www.w3.org/TR/soap12-part0/
15. Liu, C.K., Booth, D.: Web services description language (WSDL) version 2.0 part 0: Primer. W3C recommendation, W3C (2007), http://www.w3.org/TR/wsdl20-primer
16. Sun Microsystems: RPC: Remote procedure call protocol specification version 2. IETF RFC 1057, Informational (1998)
17. Winer, D.: XML-RPC specification (1999), http://www.xmlrpc.com/spec
18. Küsters, R., Tuengerthal, M.: Joint state theorems for public-key encryption and digital signature functionalities with local computation. Cryptology ePrint Archive, Report 2008/006 (2008), http://eprint.iacr.org/

A Corruption

Both in the ideal functionality $\mathcal{F}_{2\mathrm{AM}}$ and in the implementation $\mathcal{P}_{2\mathrm{AMEX}-1}$ we model corruption by using the corruption macro from [12] in a slightly modified variant, in which we add parameters for an addressing mechanism. Details can be found in [13].

Using the corruption macro we allow the adversary to corrupt our clients and servers, while the environment can check the corruption status of each instance and provide resources for corrupted machines. Once corrupted, clients and servers abort their normal execution and only forward messages from and to the adversary as defined in the macro.

While the adversary can corrupt single client instances, the situation on the server side is different: If the adversary sends a corruption request to one instance of $\mathcal{F}_{\mathrm{S}}$ running under identity s, this instance will accept all messages which are directed to any instance running under identity s. This reflects that in the implementation $\mathcal{P}_{2\mathrm{AMEX}-1}$ only one (long-lived) instance of $\mathcal{P}_{\mathrm{S}}$ is running per identity.

Note that the signature and verification functionality $\mathcal{F}_{\mathrm{Sig}}$ used in $\mathcal{P}_{2\mathrm{AMEX}-1}$ also allows corruption. But if the adversary would corrupt, e.g., a verification instance, it would have no advantage against our protocol as long as it does not also corrupt the server or client using that particular instance of the verifier. In addition, in $\mathcal{P}_{2\mathrm{AMEX}-1}$ the environment would have to pass resources to that verification instance, while in $\mathcal{F}_{2\mathrm{AM}}$ no signature scheme is available to receive the resources—but adding a mechanism to $\mathcal{F}_{2\mathrm{AM}}$ which receives the resources and passes them on to the simulator would result in a rather unnatural ideal functionality.

Therefore, even though we technically allow the adversary to corrupt instances of the signature scheme (or its verifiers) in $\mathcal{P}_{2\mathrm{AMEX}-1}$, we make it rather useless: Before $\mathcal{P}_{\mathrm{C}}$ and $\mathcal{P}_{\mathrm{S}}$ use any signature or verification functionality, they check the functionalities' corruption status and abort if it is corrupted. Note that the adversary may still get complete control over the input and output of a client or server by simply corrupting that client or server instance.

B The Server Protocol $\mathcal{P}_{\mathrm{S}}$

Tapes: $\mathrm{S} \twoheadleftarrow\!\!\longrightarrow E_{\mathrm{S}}$, $\mathrm{S} \dashleftarrow\dashrightarrow A_{\mathrm{S}}$, $\mathrm{S} \longleftrightarrow\!\!\!\!\twoheadrightarrow \mathrm{KS}$, $\mathrm{S} \longleftrightarrow\!\!\!\!\twoheadrightarrow \mathrm{LC}$, $\mathrm{S_{sig}} \longleftrightarrow\!\!\!\!\twoheadrightarrow \mathrm{Sig}$, $\mathrm{S_{ver}} \longleftrightarrow\!\!\!\!\twoheadrightarrow \mathrm{Sig}$
Initialization: $s = \mathrm{cap}_s = \mathrm{tol}_s^+ = m_c = \sigma_c = k_c = \varepsilon$, $R = L = []$, $t_s = t^{\min} = 0$, $state = \mathsf{Init}$,
$cor = \mathsf{false}$
Steps: `loop`
Initialize a new buffer:
`if` $(s', (\mathsf{Server}), \mathsf{Init}, 1^n)$ received from E_{S}
If $state = \mathsf{Init}$,
Send $(s', (\mathsf{Server}), \mathsf{GetParameters})$ to A_{S}.
Recv $(s', (\mathsf{Server}), \mathsf{Parameters}, \mathrm{cap}, \mathrm{tol}^+)$ from A_{S}.
Let $s = s'$. If $\mathrm{cap} \leq 0$ or $\mathrm{tol}^+ \leq 0$, abort.
Send $(s, (\mathsf{Server}, c, r), \mathsf{GetTime})$ to LC.
Recv $(s, (\mathsf{Server}, c, r), \mathsf{Time}, t)$ from LC.
Let $state = \mathsf{OK}$, $\mathrm{cap}_s = \mathrm{cap}$, $\mathrm{tol}_s^+ = \mathrm{tol}^+$, $t_s = t$, $t^{\min} = t_s + \mathrm{tol}_s^+$.
Append n to R.

Receive and process a request: Request the client's key:
`if` (m, σ) received from A_S with $m = (\mathsf{From}\colon c, \mathsf{To}\colon s, \mathsf{MsgID}\colon r, \mathsf{Time}\colon t, \mathsf{Body}\colon p_c)$
  If $state = \mathsf{Init}$ or R is empty or cor, abort.
  Let n be the first item of R. If $|p_c| > n$, abort. Remove n from R.
  Let $state = \mathsf{WaitingForKey}_c$, $m_c = m$, and $\sigma_c = \sigma$.
  Send $(c, (\mathsf{Client}, s, r), \mathsf{GetKey})$ to KS.
Receive and process a request: Receive the key, request time:
`if` $(c, (\mathsf{Client}, s, r), \mathsf{PublicKey}, k)$ received from KS
  If $state \neq \mathsf{WaitingForKey}_c$ or cor, abort. Let $state = \mathsf{WaitingForTime}$ and $k_c = k$.
  Send $(s, (\mathsf{Server}, c, r), \mathsf{GetTime})$ to LC.
Receive and process a request: Receive time, initialize the verifier:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Time}, t)$ received from LC
  If $state \neq \mathsf{WaitingForTime}$ or cor, abort.
  If $t \geq t_s$, let $t_s = t$. Let $state = \mathsf{WaitingForVerifier}$.
  Send $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Init})$ on $\mathrm{S_{ver}}$.
Receive and process a request: Execute 2AMEX-1 protocol steps, relay request:
`if` $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Init})$ received on $\mathrm{S_{ver}}$
  If $state \neq \mathsf{WaitingForVerifier}$ or cor, abort. Let $state = \mathsf{OK}$.
  Send $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Corrupted?})$ to KS.
  Recv $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Corrupted}, cor')$ from KS. If cor', abort.
  Send $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Verify}, m_c, \sigma_c, k_c)$ on $\mathrm{S_{ver}}$.
  Recv $(c, (\mathsf{Client}, s, r), \mathsf{Server}, \mathsf{Verified}, b)$ on $\mathrm{S_{ver}}$.
  If $b \neq 1$, $t \leq t^{\min}$ or $t > t_s + \mathrm{tol}_s^+$, or $(t', r, c') \in L$ for some t', c', abort.
  While $|L| \geq \mathrm{cap}_s$:
    Let $t^{\min} = \min\{t' \mid (t', r', c') \in L\}$ and $L = \{(t', r', c') \in L \mid t' > t^{\min}\}$.
  Insert (t, r, c) into L and send $(s, (\mathsf{Server}, c, r), \mathsf{Request}, p_c)$ to E_S.
Receive and process a response: Receive response payload, request key:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{Response}, p_s)$ received from E_S
  If $state = \mathsf{Init}$ or cor, abort.
  If $(t', r, c) \notin L$ for any t':
    Let $state = \mathsf{OK}$, send $(s, (\mathsf{Server}, c, r), \mathsf{Response}, \mathsf{Error})$ to E_S, and abort.
  Let $state = \mathsf{WaitingForKey}_s$ and send $(s, (\mathsf{Server}, c, r), \mathsf{GetKey})$ to KS.
Receive and process a response: Construct, sign, and send response message:
`if` $(s, (\mathsf{Server}, c, r), \mathsf{PublicKey}, k)$ received from KS
  If $state \neq \mathsf{WaitingForKey}_s$ or cor, abort. Let $state = \mathsf{OK}$.
  Send $(s, (\mathsf{Server}, c, r), \mathsf{Corrupted?})$ to KS.
  Recv $(s, (\mathsf{Server}, c, r), \mathsf{Corrupted}, cor)$ from KS. If cor', abort.
  Let $m_s = (\mathsf{From}\colon c, \mathsf{To}\colon s, \mathsf{Ref}\colon r, \mathsf{Body}\colon p_s)$.
  Send $(s, (\mathsf{Server}, c, r), \mathsf{Sign}, m_s)$ on $\mathrm{S_{sig}}$.
  Recv $(s, (\mathsf{Server}, c, r), \mathsf{Signature}, \sigma_s)$ on $\mathrm{S_{sig}}$.
  Update (t, r, c) to $(t, r, *)$ in L and send (m_s, σ_s) to A_S.
Reset the server:
`if` (s, Reset) received from A_S
  If $state = \mathsf{Init}$ or cor, abort.
  Send $(s, \mathsf{Server}, \mathsf{GetTime})$ to LC.
  Recv $(s, \mathsf{Server}, \mathsf{Time}, t)$ from LC.
  If $t \geq t_s$, let $t_s = t$.
  Let $t^{\min} = t_s + \mathrm{tol}_s^+$, $R = L = [\,]$, and $state = \mathsf{OK}$.
Corruption: $\mathtt{Corr}(cor, \mathsf{true}, state \neq \mathsf{Init}, \varepsilon, A_\mathrm{S}, \{E_\mathrm{S}\}, E_\mathrm{S}, s)$
CheckAddress: Check for s as soon as it has been set.