

Universally Composable Time-Stamping Schemes with Audit ^{*}

Ahto Buldas^{1,2,3,**}, Peeter Laud^{1,2,***}, Märt Saarepera, and Jan Willemson^{1,4}

¹ University of Tartu, Liivi 2, 50409 Tartu, Estonia.

² Cybernetica, Akadeemia tee 21, 12618 Tallinn, Estonia.

³ Tallinn University of Technology, Raja 15, 12618 Tallinn, Estonia.

⁴ Playtech Estonia.

Abstract. We present a Universally Composable (UC) time-stamping scheme based on universal one-way hash functions. The model we use contains an ideal auditing functionality, the task of which is to check that the rounds' digests are correctly computed. Our scheme uses hash-trees and is just a slight modification of the known schemes of Haber-Stornetta and Benaloh-de Mare, but both the modifications and the audit functionality are crucial for provable security. We prove that our scheme is nearly optimal – in every UC time-stamping scheme, almost all time stamp requests must be communicated to the auditor.

1 Introduction

Time-stamping is an important data integrity protection mechanism the main objective of which is to prove that electronic records existed at certain time. The scope of applications of time-stamping is very large and the combined risks related to time stamps are potentially unbounded. Hence, the standard of security for time-stamping schemes must be very high. It is highly unlikely that currently popular trusted third party solutions are sufficient for all needs, since the practice has shown that insider threats by far exceed the outside ones. This motivates the development of time-stamping schemes that are provably secure even against malicious insiders.

Several constructions of potentially insider-resistant time-stamping schemes have been proposed [6, 15, 16, 7, 20] based on collision-resistant hash functions, but only few analytical arguments confirm the security of these schemes. Two early attempts to sketch a security proof [6, 16] were recently shown to be flawed [9]. Presently, there are two schemes with correct security proofs: a non-interactive time-stamping scheme in the bounded storage model [20] and a bounded hash-chain scheme in the standard model [9]. However, the schemes in use (like [26–28]) still have no security proofs.

The formal security conditions for time-stamping schemes are still a subject under discussion. The early works [6, 15, 16] focused on the *consistency of databases* maintained by time-stamping service providers. It was required to be hard to change the database without compromising its consistency with a digest published in a secure

^{*} This paper is an extended abstract. Proofs of the results are presented in the full version [8].

^{**} Supported by Estonian SF grant no. 5870.

^{***} Supported by Estonian SF grant no. 6095.

repository. In [9] it was pointed out that one of the implicit assumptions of the consistency condition – the adversary knows at least one pre-image of a published digest – may be unjustified for malicious service providers. An independent security condition was proposed [9] in which the stream of time stamp requests is modeled as a high-entropy distribution. Considering the wide range of time-stamping applications, it cannot be taken for granted that these two conditions are sufficient. *Universal Composability (UC) framework* [1–4, 10, 22–24] provides a more general approach to security – rather than studying *ad hoc* behavior of adversaries, it is proved that real security primitives faithfully implement a certain *ideal primitive* the security of which is evident. Therefore, *all* security features of the ideal primitive (including the *ad hoc* ones mentioned above) are transferred to the real primitive.

In this paper, we construct universally composable time-stamping schemes under an assumption that they contain a third party auditing functionality. The idea of third-party audit in time-stamping schemes is natural and certainly not new. It has been proposed as one of the additional security measures in commercial time-stamping schemes [26]. Still, the formal security conditions presented thus far do not include the audit explicitly. We include audit functions into a general time-stamping scheme and present new security conditions that reflect two different types of auditability – audit-supported publishing and multi-round audit. We present a practical construction of an auditable time-stamping scheme that uses slightly modified Merkle trees [19] and collision resistant (or universal one-way) hash functions. We prove that the scheme is secure in the sense of conventional security conditions, assuming that the underlying hash function is collision-resistant. The auditor is crucial in the scheme – the negative results in [9] imply that the ordinary reduction techniques are insufficient for such proofs in case no additional functionalities are added to the time-stamping scheme.

We also prove that our construction of a time-stamping scheme with audit-supported publishing is universally composable, if the hash function used is universally one-way. Our construction turns out to be nearly optimal in the sense of communication between the time-stamping service and the auditor.

In Section 2, we present notations and definitions. In Section 3, we define auditable time-stamping schemes and the corresponding security notions. In Section 4, we construct an auditable time-stamping scheme based on collision-resistant hash functions. In Section 5, we outline a proof that our construction gives a universally composable time-stamping scheme with audit-supported publishing. In Section 6, we prove that our construction is nearly optimal.

2 Notation and Definitions

By $x \leftarrow \mathcal{D}$ we mean that x is chosen randomly according to a distribution \mathcal{D} . If A is a probabilistic function or a Turing machine, then $x \leftarrow A(y)$ means that x is chosen according to the output distribution of A on an input y . By \mathcal{U}_n we denote the uniform distribution on $\{0, 1\}^n$. If $\mathcal{D}_1, \dots, \mathcal{D}_m$ are distributions and $F(x_1, \dots, x_m)$ is a predicate, then $\Pr[x_1 \leftarrow \mathcal{D}_1, \dots, x_m \leftarrow \mathcal{D}_m: F(x_1, \dots, x_m)]$ denotes the probability that $F(x_1, \dots, x_m)$ is true after the ordered assignment of x_1, \dots, x_m . For functions $f, g: \mathbb{N} \rightarrow \mathbb{R}$, we write $f(k) = O(g(k))$ if there are $c, k_0 \in \mathbb{R}$, so that $f(k) \leq cg(k)$

($\forall k > k_0$). We write $f(k) = \omega(g(k))$ if $\lim_{k \rightarrow \infty} \frac{g(k)}{f(k)} = 0$. If $f(k) = k^{-\omega(1)}$, then f is *negligible*. A Turing machine M is *polynomial-time* (*poly-time*) if it runs in time $k^{O(1)}$, where k denotes the input size. Let FP be the class of all functions $f: \{0, 1\}^* \rightarrow \{0, 1\}^*$ computable by a poly-time M . A distribution \mathcal{D} on $\{0, 1\}^*$ is *poly-sampleable* if it is an output distribution of a poly-time Turing machine. A poly-sampleable distribution \mathcal{D} is *unpredictable* if $\Pr[L \leftarrow \Pi(1^k), x \leftarrow \mathcal{D}: x \in L] = k^{-\omega(1)}$ for every predictor $\Pi \in \text{FP}$. We say that \mathcal{D}_1 and \mathcal{D}_2 are *indistinguishable* (and write $\mathcal{D}_1 \approx \mathcal{D}_2$) if for every distinguisher $\Delta \in \text{FP}$: $|\Pr[x \leftarrow \mathcal{D}_1: \Delta(1^k, x) = 1] - \Pr[x \leftarrow \mathcal{D}_2: \Delta(1^k, x) = 1]| = k^{-\omega(1)}$.

A *collision-resistant hash function* is a family $h = \{h_k: \{0, 1\}^* \rightarrow \{0, 1\}^k\}_{k \in \mathbb{N}}$, such that $\delta(k) = \Pr[h_k \leftarrow \mathfrak{F}, (x, x') \leftarrow A(1^k, h_k): x \neq x', h_k(x) = h_k(x')] = k^{-\omega(1)}$ for every $A \in \text{FP}$. Here, \mathfrak{F} is a poly-sampleable distribution on F^* . A family h , such that $\Pr[x \leftarrow A_1(1^k), h_k \leftarrow \mathfrak{F}, x' \leftarrow A_2(h_k, x): x \neq x', h_k(x) = h_k(x')] = k^{-\omega(1)}$ for every $(A_1, A_2) \in \text{FP}$, is called a *Universal One-Way Hash Function (UOWHF)*. UOWHFs can be built from one-way functions [21]. We write $h(x)$ instead of $h_k(x)$.

3 Auditable Time-Stamping Schemes

3.1 General Definition of a Time-Stamping Scheme

A time-stamping scheme TS is capable of: (1) assigning a time-value $t \in \mathbb{N}$ to each request $x \in \{0, 1\}^k$, and (2) verifying whether x was time-stamped during the t -th time unit (hour, day, week, etc.). Time-stamping schemes consist of the following processes:

- *Repository* – a write only database that receives k -bit digests, adds them to a list \mathcal{D} . Repository also receives queries $\tau \in \mathbb{N}$ and returns $\mathcal{D}[\tau]$ if $\tau \leq |\mathcal{D}|$. Otherwise, Repository returns NIL. We assume that the repository is updated in a regular way (say daily), and the update time/date is known to the users of the system. This is a link between the real time and the modeled time value $t = |\mathcal{D}|$. Practical schemes [26] use newspaper-publishing as the Repository. Hence, it is reasonable to assume that Repository is costly and to keep the number of stored bits as small as possible.
- *Stamper* – operates in discrete time intervals called *rounds*. During the t -th round, Stamper receives requests x and returns pairs (x, t) . We assume that Stamper “knows” how many digests have been stored to Repository. Let L_t be the list of all requests received during the t -th round. In the end of the round, Stamper creates a *certificate* $c = \text{Stp}(x; L_t, L_{t-1}, \dots, L_1)$ for each request $x \in L_t$. Besides, Stamper computes a digest $d_t = \text{Pub}(L_t, \dots, L_1)$ and sends d_t to Repository.
- *Verifier* – a computing environment for verifying time stamps. In practice, each user may have its own Verifier but for the security analysis, it is sufficient to have only one. It is assumed that Verifier has a tamperproof access to Repository. On input (x, t) , Verifier obtains a certificate c from Stamper, and a digest $d = \mathcal{D}[t]$ from Repository, and returns $\text{Ver}(x, c, d) \in \{\text{yes}, \text{no}\}$. It is not specified how c is transmitted from Stamper to Verifier. In practice, c can be stored together with x . Hence, the size of c should be reasonable. Note that x can be verified only after the digest d_t is sent to Repository. This is acceptable, because in the applications we address, x is verified long after stamping.

- Client – any application-environment that uses Stamper and Verifier.

Definition 1 (Correctness). A triple $TS = (\text{Stp}, \text{Pub}, \text{Ver})$ is a time-stamping scheme if $\text{Ver}(x_n, \text{Stp}(x, \mathcal{L}), \text{Pub}(\mathcal{L})) = \text{yes}$ for every $\mathcal{L} = (L_t, \dots, L_1)$, and $x \in L_t$.

3.2 Security Conditions

It is assumed that an adversary A is able to corrupt Stamper, some instances of Client and some instances of Verifier. The Repository is assumed to be non-corrupting. After closing the t -th round (i.e. after publishing d_t) it should be impossible to add a new request x to the set L_t of requests and prove to a Verifier that $x \in L_t$ by finding a suitable certificate c . This suggests the following security condition:

Definition 2 (Consistency). A time-stamping scheme is consistent if for every $A \in \text{FP}$:

$$\Pr[(L_t, \dots, L_1, c, x) \leftarrow A(1^k): x \notin L_t, \text{Ver}(x, c, \text{Pub}(L_t, \dots, L_1)) = \text{yes}] = k^{-\omega(1)} . \quad (1)$$

The condition (1) is not completely satisfactory because the adversary has to explicitly construct the lists L_t, \dots, L_1 of time-stamped requests. Back-dating attacks can be possible without A creating these lists. For example, A may publish a value d which is not necessarily computed by using the Pub function and then, after obtaining a new (random) x , to find c so that $\text{Ver}(x, c, d) = \text{yes}$. This suggests a different condition [9]:

Definition 3 (Security against random back-dating). A time-stamping scheme is secure if for every unpredictable distribution \mathcal{D} on $\{0, 1\}^k$ and $(A_1, A_2) \in \text{FP}$:

$$\Pr[(d, a) \leftarrow A_1(1^k), x \leftarrow \mathcal{D}, c \leftarrow A_2(x, a): \text{Ver}(x, c, d) = \text{yes}] = k^{-\omega(1)} . \quad (2)$$

In some applications, additional security features (like confidentiality of messages, availability etc.) of time-stamping schemes are required. We do not study these features.

3.3 Time-Stamping Schemes with Audit

It is essential for the security of time-stamping that a corrupted Stamper is not able to publish a value d in Repository without actually knowing a database (x_1, \dots, x_n) such that $\text{Pub}(x_1, \dots, x_n) = d$. Otherwise, it could be difficult (or even impossible) to find a security proof [9]. The easiest way to prove such knowledge is sending the requests x_1, \dots, x_n to a trusted Auditor who checks if $\text{Pub}(x_1, \dots, x_n) = d$. Audit can be performed before or after publishing. We observe two different audit models:

- *Audit-Supported Publishing.* In this model, the roles of Repository and Auditor are merged. If the t -th round is closed, the Auditor/Repository receives a list L_t of bit-strings and an audit report from Stamper and checks their correctness. The digest is not published if the audit report is incorrect.
- *Multi-Round Audit.* In this model, audit reports are checked long after publishing, which is much more close to the real-life (yearly) audit.

We define two additional functions: Rep for creating a report $r_t = \text{Rep}(L_t, \dots, L_1)$, and Aud for checking the consistency of r_t and $d_t = \text{Pub}(L_t, \dots, L_1)$.

Definition 4. A 5-tuple $ATS = (\text{Stp}, \text{Pub}, \text{Ver}, \text{Rep}, \text{Aud})$ is an auditable time-stamping scheme if $\text{Aud}(\text{Rep}(\mathcal{L}), \text{Pub}(\mathcal{L})) = \text{yes}$, for any $\mathcal{L} = (L_t, \dots, L_1)$ (properly created audit reports verify successfully), and $(\text{Stp}, \text{Pub}, \text{Ver})$ is a time-stamping scheme.

In this paper, we assume that $\text{Aud}(L_t, \dots, L_1)$ depends only on the first argument L_t . The results we obtain for such schemes can be easily generalized.

Schemes with Audit-Supported Publishing. The audit is performed during (or before) publishing. The auditor is a trusted middle-man between Stamper and Publisher. After the t -th round, Stamper computes a digest $d_t = \text{Pub}(L_t, \dots, L_1)$ and an audit report $r_t = \text{Rep}(L_t, \dots, L_1)$. Having sent a pair (d, r) , the auditor checks whether $\text{Aud}(r, d) = \text{yes}$ and sends d to Repository. Hence, a successful publishing is possible only if a correct audit report is sent to the auditor. A time-stamping scheme with audit-supported publishing is secure against random back-dating if for every unpredictable \mathcal{D} and for every $(A_1, A_2) \in \text{FP}$:

$$\Pr[(d, r, a) \leftarrow A_1(1^k), x \leftarrow \mathcal{D}, c \leftarrow A_2(x, a): \text{Ver}(x, c, d) = \text{yes} = \text{Aud}(r, d)] = k^{-\omega(1)} . \quad (3)$$

Schemes with Multi-Round Audit. Publishing is done like in the schemes without audit. The audit function is performed after publishing. If N rounds are passed, Stamper computes audit reports $r_1 = \text{Rep}(L_1), \dots, r_N = \text{Rep}(L_N, \dots, L_1)$ and sends (r_1, \dots, r_N) to the auditor. For $t = 1 \dots N$, the auditor obtains d_t from Repository and computes $\text{Aud}(r_t, d_t)$. If for some t the result is no then all users are informed. A time-stamping scheme with multi-round audit is secure against random back-dating if for every unpredictable \mathcal{D} and for every $(A_1, A_2) \in \text{FP}$:

$$\Pr[(d, a) \leftarrow A_1(1^k), x \leftarrow \mathcal{D}, (c, r) \leftarrow A_2(x, a): \text{Ver}(x, c, d) = \text{yes} = \text{Aud}(r, d)] = k^{-\omega(1)} . \quad (4)$$

3.4 Records of Arbitrary Length

The definitions above assume that all time stamp requests are k bits long. To time-stamp longer records, practical schemes use collision-resistant hash functions (at the client side) to make requests shorter. Since these hash functions have influence on security, they have to show up in the security conditions.

Definition 5. A time-stamping scheme with audit-supported publishing is secure relative to a client side hash function $h: \{0, 1\}^* \rightarrow \{0, 1\}^k$ if for every unpredictable \mathcal{D} on $\{0, 1\}^*$ and for every $(A_1, A_2) \in \text{FP}$:

$$\Pr[(d, r, a) \leftarrow A_1(1^k), X \leftarrow \mathcal{D}, c \leftarrow A_2(X, a): \text{Ver}(h(X), c, d) = \text{yes} = \text{Aud}(r, d)] = k^{-\omega(1)} . \quad (5)$$

A time-stamping scheme with multi-round audit is said to be secure relative to a client side hash function h if for every unpredictable \mathcal{D} and for every $(A_1, A_2) \in \text{FP}$:

$$\Pr[(d, a) \leftarrow A_1(1^k), X \leftarrow \mathcal{D}, (c, r) \leftarrow A_2(X, a): \text{Ver}(h(X), c, d) = \text{yes} = \text{Aud}(r, d)] = k^{-\omega(1)} . \quad (6)$$

Lemma 1. If \mathcal{D} is unpredictable and h is collision-resistant then $h(\mathcal{D})$ is unpredictable.

In spite of Lemma 1, a secure auditable scheme in the sense of (3) or (4) is not necessarily secure relative to every collision-resistant hash function ((5),(6)) because, in (5) and (6), A_2 has more information (an h -pre-image X of x) than in (3) and (4).

4 Construction of an Auditable Time-Stamping Scheme

Let h be a collision-resistant hash function, or a UOWHF chosen by Repository. We define $\text{ATS}^h = (\text{Pub}^h, \text{Stp}^h, \text{Ver}^h, \text{Rep}^h, \text{Aud}^h)$ and prove that this 5-tuple of functions form a secure time-stamping scheme with audit. Let $L = (x_0, \dots, x_{m-1})$ be all requests received during the t -th round. For simplicity, we assume that $m = 2^\ell$.

The publishing function $\text{Pub}^h(L)$ builds a complete binary tree of height ℓ each vertex v of which has a $(k+1)$ -bit label $\Lambda[v] = b\|H[v]$, where $b \in \{0, 1\}$ indicates whether v is a leaf ($b = 0$ iff v is a leaf) and $H[v] \in \{0, 1\}^k$ is a hash value computed by the following (inductive) scheme. For the n -th leaf v , we define $H[v] = x_n$, and $H[v] = h(\Lambda[v_L]\|\Lambda[v_R])$ for any non-leaf v , where v_L and v_R are the left- and the right child of v , respectively. As a result, $\text{Pub}^h(L)$ returns a $(k+1)$ -bit root label of the tree.

The stamping function $\text{Stp}^h(L, n)$ builds the same tree as above. Let v be the n -th leaf and $v = v_0, v_1, \dots, v_{\ell-1}, v_\ell$ be the unique path from v to the root vertex (v_ℓ), i.e. v_i is a child of v_{i+1} for every $i \in \{0, \dots, \ell-1\}$. Let $v'_0, v'_1, \dots, v'_{\ell-1}$ denote the siblings of $v_0, v_1, \dots, v_{\ell-1}$, respectively. Let $z_i = \Lambda[v'_i]$ for every $i \in \{0, \dots, \ell-1\}$ and $z = (z_0, \dots, z_{\ell-1})$. The certificate is $c = \text{Stp}^h(L, n) = (n, z)$.

The verification function $\text{Ver}^h(x, (n, z), d)$ recomputes d (based on x and (n, z)) and compares the results. Let $n = n_{\ell'-1}n_{\ell'-2} \dots n_0$ be the binary representation of n and $z = (z_0, z_1, \dots, z_{\ell'-1})$. The function Ver^h computes sequences $\lambda = (\lambda_0, \lambda_1, \dots, \lambda_{\ell'}) \in (\{0, 1\}^{k+1})^{\ell'}$ and $\chi = (\chi_0, \chi_1, \dots, \chi_{\ell'}) \in (\{0, 1\}^k)^{\ell'}$ inductively, so that $\chi_0 = x$, $\lambda_0 := 0\|x$, and for every $i > 0$, $\lambda_i = 1\|\chi_i$, where

$$\chi_i := \begin{cases} h(z_i\|\lambda_{i-1}) & \text{if } n_{i-1} = 1 \\ h(\lambda_{i-1}\|z_i) & \text{if } n_{i-1} = 0 \end{cases} . \quad (7)$$

The verification procedure outputs yes, iff $\lambda_{\ell'} = d$.

The report function is trivial, i.e. $\text{Rep}^h(L) = L$ for every list L . *The audit function* $\text{Aud}^h(L, d)$ computes $d' = \text{Pub}^h(L)$ and returns yes iff $d' = d$.

Lemma 2. (A) *If $x \notin L$, and $\text{Aud}^h(L, d) = \text{Ver}^h(x, c, d) = \text{yes}$ then the h -calls of Ver^h and Pub^h contain a collision for h .* **(B)** *If $L \neq L'$ and $\text{Pub}^h(L) = \text{Pub}^h(L')$ then the h -calls performed by Pub^h contain a collision for h .*

Theorem 1. *If h is collision resistant then a time-stamping scheme ATS^h with audit-supported publishing is secure relative to a client-side hash function h .*

Theorem 2. *If h is collision resistant then a time-stamping scheme ATS^h with multi-round audit is secure relative to a client-side hash function h .*

Proofs of these results are presented in the full version [8].

Lemma 2 directly implies the consistency condition (1) for ATS^h . Hence, we have proved that our construction is secure in the conventional sense. Note also that it is

probably *not possible* to prove that the scheme TS^h without audit is secure against random back-dating (2), based on the collision-resistance of h . The reason is that one can find an oracle \mathcal{O} and choose a hash function h (that uses \mathcal{O}) so that h is collision-resistant but TS^h is still insecure [9]. As the ordinary reduction techniques relativize, the security of TS^h cannot be proved (in ordinary way) in the real world either. In this sense, the audit functionality is crucial for provable security.

5 Universally Composable Time-Stamping Schemes

5.1 Universal Composability Framework

To prove that a cryptographic primitive is secure in *every reasonable application* the *universal composability* (UC) paradigm is used [1–4, 10–12, 22–24]. If the reader is not familiar with the UC paradigm, we recommend to study the seminal works by Canetti [11, 12] and the monograph on composability by Lindell [17]. Rather than using *ad hoc* behavior of adversaries, the UC paradigm defines an *ideal primitive* which is “obviously secure” and then proves that if $A \in \text{FP}$ is an adversary for an application of the real primitive then there is another adversary $A' \in \text{FP}$ for the same application in which the real primitive is replaced with the ideal primitive. Loosely speaking, no security incident in any application of the primitive is caused by the difference between the real and the ideal primitives – *the real functionality faithfully implements the ideal functionality*.

We use the language of Finite State Machine (FSM) theory borrowed from Pfitzmann [24] when describing the UC formalism. Every component of the system (for a fixed value of k) is a (probabilistic) FSM with input- and output ports. Each port has a name and a type (in or out). By a *composition* $\langle M_1, M_2 \rangle$ of two machines M_1 and M_2 we mean a network of machines obtained by connecting the input and output port pairs in a certain (pre-defined) way. For example, pairs with identical names can be connected. The precise formalism for describing the connections is unimportant in this paper, because the networks we use are very simple. We assume that each input port is buffered, whereas the length of the buffer is unlimited. When analyzing a particular machine, we use the following abbreviations. By $\text{in}_\nu \rightarrow x$ we denote the event that the machine has input x in the port in_ν . By $y \rightarrow \text{out}_\nu$ we mean that y is sent to the output port out_ν . To overcome the difficulties related to the asynchronous behavior of the network, it is assumed that no two machines run at the same time. Technically, this is achieved by introducing the clock-ports to the system. Each machine, after finishing its work, can clock (give the token to) only one machine. In this paper, we use clocked output signals. By $x \xrightarrow{\text{clk}} \text{out}_\nu$ we mean that x is sent to the output port named ν and the token is given to the machine with input port in_ν . By the *view* of M_i in a composition $M = \langle M_1, \dots, M_n \rangle$ we mean the sequence of all input/output signals of M_i in a particular run of M . The view is denoted by $\text{VIEW}_{M_i} \langle M_1, \dots, M_n \rangle$. In general, the view is a probability space.

In the UC framework, we have an ideal time-stamping scheme TS_I , a real scheme TS_R , and an environment Client. A composition $\langle \text{Client}, TS_R \rangle$ is called a *real application*, while $\langle \text{Client}, TS_I \rangle$ is called an *ideal application*. Each machine has special input/output ports for an adversary A .

Definition 6 (Universal Composability). TS_R is universally composable, if for every $A \in \text{FP}$ there is a $A' \in \text{FP}$, so that for every $\text{Client} \in \text{FP}$: $\text{VIEW}_{\text{Client}}(\text{Client}, TS_R, A) \approx \text{VIEW}_{\text{Client}}(\text{Client}, TS_I, A')$.

Informally, this condition means that anything that may happen to the real application $\langle \text{Client}, TS_R \rangle$ may also happen to the ideal application $\langle \text{Client}, TS_I \rangle$.

In the proofs of UC, a simulator S is constructed that uses A as a black-box, i.e. $A' = \langle S, A \rangle$. It is then proved that $\langle TS_I, S \rangle$ and TS_R behave identically, except when certain cryptographic primitives (used by TS_R) are broken. Hence, if the primitives are believed to be secure, this implies the indistinguishability of views and also the security of TS_R in the strongest possible sense. To prove the identical behavior of $\langle TS_I, S \rangle$ and TS_R , a *bisimulation* between these two machines is constructed.

5.2 On the Model

Some primitives are hard to cast in terms of the UC framework. The *commitment problem* occurs, meaning that a simulator that acts as an intermediary between the real-world adversary and the ideal functionality has to fix the value of a certain data item without knowing all the components it was created from, and also without the ability to present instead of this data item something that is and remains indistinguishable from it. Canetti and Fishlin proved [13] that UC bit-commitment is impossible in the "plain model" (i.e. a model without ideal functionalities) but it becomes possible in the *Common Reference String (CRS)* model, where a common (and accessible) random string is added to the system as an ideal functionality. Similar problems occur when trying to define universally composable time-stamping schemes, but fortunately, the problems disappear if an *ideal audit functionality* (represented in our model by Repository that is merged with Auditor) is added to the system. The universal composability can be proved based on the *universal one-wayness* of a hash function h , assuming that a new random instance of h is generated (by Repository) during each round. The reduction we obtain is linear-preserving and gives good practical security guarantees.

Hence, our UC Time-Stamping scheme construction is not in the plain model. However, adding the trusted Repository to the system is reasonable because: (1) there are real-life systems that behave in a similar way (e.g. newspapers); (2) it is possible to implement similar functionalities in the CRS model by using public-key cryptography.

5.3 Ideal Time-Stamping Scheme

The ideal scheme is a secure host that stores for each round number t a set L_t of all bit-strings that were stamped during the t -th round. The value of t is initially 0 and is incremented each time the round is closed. In our real scheme, we allow the stamping functionality to be corrupted. This is reflected in the ideal scheme by giving the adversary complete control on which bit-strings will be considered stamped during the current round. As we shall see, at the end of the round t the adversary sends the contents of L_t to the secure host. Hence no availability is guaranteed. The important property is, however, that once the t -th round has ended, no more bit-strings can be added to L_t — *back-dating* is not possible.

In the real world, the verification of a time-stamp may fail for a number of reasons that are under the control of the adversary. For example, the repository may be currently unavailable or it may be available but not yet contain the digest of the round we are interested in. In this case we cannot rely on the time-stamp and must behave as if it was invalid. In the ideal world we model this situation by allowing the adversary to declare any verification attempt unsuccessful. However, the adversary is unable to declare a time-stamp valid if it really was not.

The internal state of the ideal time-stamping scheme TS_I consists of an indexed list \mathcal{L}_I each element $\mathcal{L}_I[t]$ of which is a set of k -bit strings. Initially, $\mathcal{L}_I = \llbracket$. The ideal scheme TS_I (Fig. 1, left) offers service on ports in_{req} , out_{st} , in_{ver} , and out_{res} . The other ports (out_{req} , in_{st} , in_{aud} , out_{ver} , and in_{res}) are intended for communication with an adversary A' . In the following, we describe the behavior of TS_I by defining its reaction to any possible input.

- If $in_{req} \rightarrow x$ then $x \rightarrow out_{req}$.
- If $in_{st} \rightarrow (x, t)$ then $(x, t) \rightarrow out_{req}$.
- If $in_{aud} \rightarrow L$ then $\mathcal{L}_I := \mathcal{L}_I \parallel L$.
- If $in_{ver} \rightarrow (x, \tau)$ then $(x, \tau) \rightarrow out_{ver}$.
- If $in_{res} \rightarrow (x, \bar{b}, \tau)$ then $b := \bar{b} \& \text{True}(x \in \mathcal{L}_I[\tau])$ and $(x, b, \tau) \rightarrow out_{res}$.

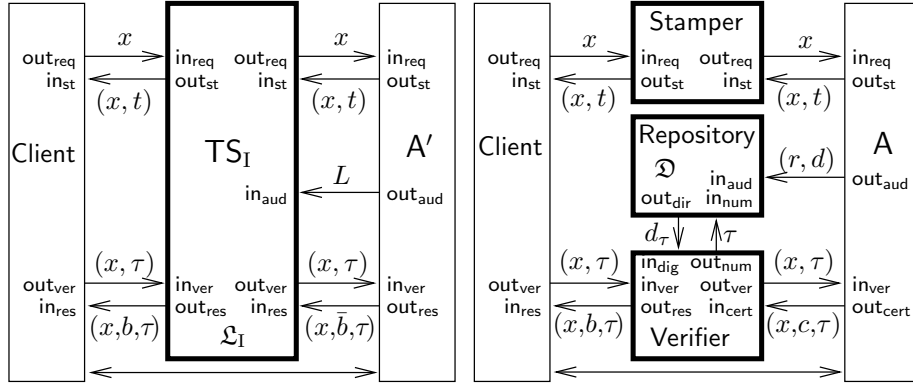


Fig. 1. The ideal scheme TS_I and the real scheme $TS_R = \langle \text{Stamper}, \text{Repository}, \text{Verifier} \rangle$.

5.4 Real Scheme

In the real scheme, the trusted host is replaced by a number of Verifier hosts. Some of them may be corrupted but we observe only one non-corrupted Verifier. This is allowed because in the standard time-stamping setting, there is no communication between verifiers. We assume that the channel between Repository and (non-corrupted) Verifier is

tamperproof. It is a reasonable practical assumption because channels with similar security properties (e.g. newspapers) exist in the real life.

Having obtained a verification request (x, t) (which reads "Was x time-stamped during the t -th round?"), Verifier obtains the corresponding r_t from Repository and applies the Ver^h procedure. However, Verifier needs a certificate c for verification. We take into account possible (malicious) modification of the certificate before verification. Therefore, it is natural to assume that the certificate is entirely provided by the adversary A. The real scheme TS_R (Fig. 1, right) consists of three components:

- Stamper – a prototype for a server that receives time stamp requests and returns time stamps to Client. As we assume that the adversary A has full control over Stamper, we define Stamper as a stateless intermediary between Client and A. Stamper offers service on ports in_{req} and out_{st} . Two other ports (out_{req} and in_{st}) are for the communication with A. The behavior of Stamper is defined as follows:
 - If $\text{in}_{\text{req}} \rightarrow x$ then $x \rightarrow \text{out}_{\text{req}}$.
 - If $\text{in}_{\text{st}} \rightarrow (x, t)$ then $(x, t) \rightarrow \text{out}_{\text{st}}$.
- Repository – a prototype for a secure repository that publishes the digests of rounds. The internal state of Repository consists of a (initially empty) list \mathcal{D} of k -bit strings. Repository offers service on ports in_{num} and out_{dir} . The third port in_{aud} is for the communication with A. The behavior of Repository is defined as follows:
 - If $\text{in}_{\text{num}} \rightarrow \tau$ and $\tau < |\mathcal{D}|$, then $d_\tau := \mathcal{D}[\tau]$ and $d_\tau \rightarrow \text{out}_{\text{dig}}$.
 - If $\text{in}_{\text{num}} \rightarrow \tau$ and $\tau \geq |\mathcal{D}|$, then $\text{NIL} \rightarrow \text{out}_{\text{dig}}$.
 - If $\text{in}_{\text{aud}} \rightarrow (r, d)$ and $\text{Aud}(r, d) = \text{yes}$ then $\mathcal{D} := \mathcal{D} \| d$.
- Verifier – a prototype for a real verification environment, which typically is a trusted client computer. Verifier receives verification requests and answers with a verification result. It is assumed that Verifier is able to obtain the digests d_τ from Repository in a tamperproof way. The internal state of Verifier consists of a bit-string variable \mathfrak{r} . Initially, $\mathfrak{r} = \perp$. Verifier offers service on ports in_{ver} and out_{res} . Two ports – out_{num} and in_{dig} – are for requesting the digests from Repository, and two last ports (out_{ver} and in_{dig}) are for the communication with A. Let $y \xrightarrow{\Delta} \text{out}_c$ denote the event that y is sent to the output channel out_c and the corresponding connection is clocked. The behavior of Verifier is defined as follows:
 - If $\text{in}_{\text{ver}} \rightarrow (x, \tau)$ then $(x, \tau) \rightarrow \text{out}_{\text{ver}}$.
 - If $\text{in}_{\text{cert}} \rightarrow (x, c, \tau)$ then $\mathfrak{r} := (x, c, \tau)$ and $\tau \xrightarrow{\Delta} \text{out}_{\text{num}}$.
 - If $\text{in}_{\text{dig}} \rightarrow d_\tau \in \{0, 1\}^k$ and $\mathfrak{r} = (x, c, \tau)$, then $b := \text{Ver}(x, c, d_\tau)$ and $(x, b, \tau) \rightarrow \text{out}_{\text{res}}$.
 - If $\text{in}_{\text{dig}} \rightarrow \text{NIL}$ and $\mathfrak{r} = (x, c, \tau)$, then $(x, \text{no}, \tau) \rightarrow \text{out}_{\text{res}}$.

For completing the description of the real scheme, it is sufficient to give efficient constructions for Pub, Aud, and Ver, i.e. exactly the components of an auditable time-scheme that appear in the security conditions (3), (4), (5), and (6). Hence, for any auditable time-stamping scheme it is reasonable to speak about *universal composability as a security condition*.

5.5 Simulator for ATS^h

We define a simulator for the scheme ATS^h , we presented in Section 4. The internal state of the simulator S (Fig. 2, right) consists of two lists $(\mathcal{D}_I, \mathcal{C}_I)$ and a bit-string τ_I . The elements of \mathcal{D}_I are k -bit strings, while the elements of \mathcal{C}_I are sets of k -bit strings. Initially, $\mathcal{D}_I = []$, $\mathcal{C}_I = (\emptyset, \emptyset, \dots)$, and $\tau_I = []$. The simulator has five ports (in_{req} , out_{st} , out_{aud} , in_{ver} , and out_{res}) for communicating with TS_I and five ports (out_{req} , in_{st} , in_{aud} , out_{ver} , and in_{cert}) for communicating with A . The behavior of S is defined as follows:

- If $in_{req} \rightarrow x$ then $x \rightarrow out_{req}$.
- If $in_{st} \rightarrow (x, t)$ then $(x, t) \xrightarrow{P} out_{st}$.
- If $in_{aud} \rightarrow (L, d)$ and $d = \text{Pub}(L)$, then $\mathcal{D}_I := \mathcal{D}_I || d$, and $L \xrightarrow{P} out_{aud}$.
- If $in_{ver} \rightarrow (x, \tau)$ then $(x, \tau) \rightarrow out_{ver}$.
- If $in_{cert} \rightarrow (x, c, \tau)$ then $\tau_I := (x, c, \tau)$, $\bar{b} := \tau < |\mathcal{D}_I| \ \& \ \text{Ver}(x, c, \mathcal{D}_I[\tau])$, and $(x, \bar{b}, \tau) \xrightarrow{P} out_{res}$. If $\bar{b} = \text{yes}$ then $\mathcal{C}_I[\tau] := \mathcal{C}_I[\tau] \cup \{x\}$.




Fig. 2. The real machine M_R , the simulator S , and the ideal machine $M_I = \langle TS_I, S \rangle$.

5.6 Bisimilarity of the Real- and the Ideal Machines

We start the proof by augmenting the state of the components Repository and Verifier of the real functionality. From the following description it is obvious that this extra state has no influence on the behavior of these components as the existing parts make no use of the new state. We add an initially empty list \mathcal{L} of sets of k -bit strings to the state of Repository. We also replace the third item in the description of its behavior by

- If $in_{aud} \rightarrow (r, d)$ and $\text{Aud}(r, d) = \text{yes}$ then $\mathcal{D} := \mathcal{D} || d$, and $\mathcal{L} := \mathcal{L} || r$.

We add a list \mathfrak{C} of sets of k -bit strings to the state of Verifier. Initially, $\mathfrak{C} = (\emptyset, \emptyset, \dots)$. We replace the third item in the description of the behavior of Verifier by

- If $\text{in}_{\text{dig}} \rightarrow d_\tau \in \{0, 1\}^k$ and $\mathfrak{r} = (x, c, \tau)$, then $b := \text{Ver}(x, c, d_\tau)$ and $(x, b, \tau) \rightarrow \text{out}_{\text{res}}$. If $b = \text{yes}$ then $\mathfrak{C}[\tau] := \mathfrak{C}[\tau] \cup \{x\}$.

Let $M_R = \text{TS}_R$ be the real machine and $M_I = \langle \text{TS}_I, S \rangle$ be the ideal machine (Fig. 2). A state $s = (\mathfrak{L}, \mathfrak{D}, \mathfrak{C}, \mathfrak{r})$ is said to be *faulty* if $\exists \tau: \mathfrak{C}[\tau] \not\subseteq \mathfrak{L}[\tau]$. Let S_R and S_I be the sets of states of M_R and M_I , respectively. Let F_R and F_I be the corresponding sets of faulty states. Let I and O be the sets of inputs and outputs (common for M_R and M_I). Let $\delta_R: I \times S_R \rightarrow S_R$ be the next-state function of M_R and $\lambda_R: I \times S_R \rightarrow O$ be the output function of M_R . We define δ_I and λ_I analogously for M_I . Let s_R^0 and s_I^0 be the initial states of the corresponding machines. By Lemma 2, if one of the machines reaches a faulty state then the h -calls performed so far comprise a collision for h .

Definition 7. Two machines M_R and M_I are said to be bisimilar with error (F_R, F_I) , if there is a binary relation (called a bisimulation) $\beta \subseteq S_R \times S_I$ such that $(s_R^0, s_I^0) \in \beta$ and for any pair of states $(s_R, s_I) \in \beta$ and for any input $i \in I$, at least one of the following three conditions holds: **(1)** $\delta_R(i, s_R) \in F_R$, **(2)** $\delta_I(i, s_I) \in F_I$, or **(3)** $(\delta_R(i, s_R), \delta_I(i, s_I)) \in \beta$ and $\lambda_R(i, s_R) = \lambda_I(i, s_I)$.

Theorem 3. The machines $M_R = \text{TS}_R$ and $M_I = \langle \text{TS}_I, S \rangle$ are bisimilar with error (F_R, F_I) , whereas the bisimulation β is defined as follows:

$$(\mathfrak{L}, \mathfrak{D}, \mathfrak{C}, \mathfrak{r}) \beta (\mathfrak{L}_I, \mathfrak{D}_I, \mathfrak{C}_I, \mathfrak{r}_I) \quad \equiv \quad (\mathfrak{L} = \mathfrak{L}_I) \& (\mathfrak{D} = \mathfrak{D}_I) \& (\mathfrak{C} = \mathfrak{C}_I) \& (\mathfrak{r} = \mathfrak{r}_I) .$$

Corollary 1. If h is a collision-resistant (or universal one-way) hash function then ATS^h is a UC time-stamping scheme with audit-supported publishing.

Proofs of these results are presented in the full version [8].

6 Size of the Audit Report

In ATS^h , the Rep function is not length-decreasing which means that the network load (and the computations) are doubled, compared to the schemes without audit. It is natural to ask whether the length of the report can be reduced. The answer turns out to be negative: in every UC time-stamping scheme with audit-supported publishing $|r_t| \approx |L_t|$.

We construct a Client and an adversary A (for $\langle \text{Client}, \text{TS}_R \rangle$) so that no efficient adversary A' (for $\langle \text{Client}, \text{TS}_I \rangle$) can simulate A unless $|\text{Rep}(r_t)| \approx |L_t|$. Our construction exploits the commitment problem – the adversary A' (or a simulator) knows only d_t but has to send L_t to TS_I , and hence L_t should be efficiently computable from d_t .

The internal state of Client consists of a $p(k)$ -element array $L = (x_1, \dots, x_{p(k)}) \in \{0, 1\}^{k \times p(k)}$ (where $p(k) = k^{O(1)}$), a k -bit string z (initially 0), and a boolean value RoundOver that is initially false. Client reacts to the input events as follows:

- If $\text{in}_A \rightarrow \text{init}$ then the Client generates $x_1, \dots, x_{p(k)}$ independently at random, computes $r = \text{Rep}(L)$, $d = \text{Pub}(L)$, and outputs $(r, d) \rightarrow \text{out}_A$.

- If $\text{in}_A \rightarrow \text{round}$ then the Client outputs $(0^k, 1) \xrightarrow{\Delta} \text{out}_{\text{ver}}$.
- If $\text{in}_{\text{ver}} \rightarrow (0^k, \text{yes}, 1)$ (a confirmation that the round is closed) then the Client sets $\text{RoundOver} := \text{true}$ and outputs $L \rightarrow \text{out}_A$ (reveals L to the adversary).
- If $\text{in}_A \rightarrow \text{verify}$ then Client generates $i \leftarrow \{1, \dots, p(k)\}$ uniformly at random, sets $z := x_i$ and outputs $(z, 0) \xrightarrow{\Delta} \text{out}_{\text{ver}}$.
- If $\text{in}_{\text{ver}} \rightarrow (z, \text{yes}, 0)$ and $\text{RoundOver} = \text{yes}$ then Client outputs $\text{yes} \rightarrow \text{out}_A$.

The adversary A is defined as follows. The internal state of A consists of a $p(k)$ -element array $L_A = (a_1, \dots, a_{p(k)}) \in \{0, 1\}^{k \times p(k)}$ (where $p(k) = k^{O(1)}$). First of all, the adversary A outputs $\text{init} \rightarrow \text{out}_{\text{Client}}$ and then reacts to the input events as follows:

- If $\text{in}_{\text{Client}} \rightarrow (r, d)$ then A outputs $(r, d) \xrightarrow{\Delta} \text{out}_{\text{aud}}$. After getting control again, the adversary A outputs $(\text{Rep}(0^k), \text{Pub}(0^k)) \xrightarrow{\Delta} \text{out}_{\text{aud}}$. Finally, A outputs $\text{round} \xrightarrow{\Delta} \text{out}_{\text{Client}}$.
- If $\text{in}_{\text{ver}} \rightarrow (0^k, 1)$ then A outputs $(0^k, \text{Stp}(0^k, 1), 0) \rightarrow \text{out}_{\text{cert}}$.
- If $\text{in}_{\text{Client}} \rightarrow L$ then A sets $L_A := L$ and outputs $\text{verify} \xrightarrow{\Delta} \text{out}_{\text{Client}}$.
- If $\text{in}_{\text{ver}} \rightarrow (z, 0)$ then A finds an i , such that $L_A = z$, computes $c := \text{Stp}(i, L_A)$, and outputs $(z, c, 0) \rightarrow \text{out}_{\text{cert}}$. The adversary halts if there is no such i .

With probability one, $\text{VIEW}_{\text{Client}}(\text{Client}, \text{TS}_R, A)$ contains the output yes from Client. Let $A' \in \text{FP}$ and $\text{VIEW}_{\text{Client}}(\text{Client}, \text{TS}_R, A) \approx \text{VIEW}_{\text{Client}}(\text{Client}, \text{TS}_I, A')$. Due to the indistinguishability, with probability $1 - k^{-\omega(1)}$ the view $\text{VIEW}_{\text{Client}}(\text{Client}, \text{TS}_I, A')$ contains the output yes from Client. From the description of TS_I , it follows that with probability $1 - k^{-\omega(1)}$ the adversary A' (based on partial information $(\text{Rep}(L), \text{Pub}(L))$) is capable of finding L_A such that $x_i \in L_A$. Lemma 3 below shows that such A' is possible only if the bit-length of $(\text{Rep}(L), \text{Pub}(L))$ is $\approx k \cdot p(k)$.

Lemma 3. *Let $X = (X_1, \dots, X_{p(k)}) \in \{0, 1\}^{k \times p(k)}$ (where $p(k) = k^{O(1)}$) and $\mathfrak{S} \leftarrow \{1, \dots, p(k)\}$ be independent and uniformly distributed random variables. Let $f: \{0, 1\}^{k \times p(k)} \rightarrow \{0, 1\}^{\ell(k)}$, $A: \{0, 1\}^{\ell(k)} \rightarrow \{0, 1\}^{k \times m(k)}$, $m(k) = k^{O(1)}$, and*

$$\delta = \Pr[X \leftarrow \{0, 1\}^k, L \leftarrow A(f(X)), \mathfrak{S} \leftarrow \{1, \dots, p(k)\}: X_{\mathfrak{S}} \in L] = 1 - k^{-\omega(1)} .$$

Then $\ell(k) = k \cdot p(k) - O(\log k)$.

Proof. A $p(k)$ -tuple $(x_1, \dots, x_{p(k)})$ is *good* if $x_i \in A(f(x_1, \dots, x_{p(k)}))$ for all $i \in \{1, \dots, p(k)\}$. Other tuples are *bad*. As for any bad tuple $(x_1, \dots, x_{p(k)})$ the probability of error $1 - \delta \geq \frac{1}{p(k)} \neq k^{-\omega(1)}$, the number of good tuples should be $(1 - k^{-\omega(1)}) \cdot 2^{k \cdot p(k)}$. On the other hand, the number of good tuples cannot exceed $2^{\ell(k)} \cdot m(k)^{p(k)}$ and hence $2^{\ell(k)} \cdot m(k)^{p(k)} = (1 - k^{-\omega(1)}) \cdot 2^{k \cdot p(k)}$, which gives (by taking logarithm from both sides) $\ell(k) = k \cdot p(k) - O(\log k)$. \square

Corollary 2. *In every UC secure time-stamping scheme with audit-supported publishing $\text{ATS} = (\text{Pub}, \text{Stp}, \text{Ver}, \text{Rep}, \text{Aud})$, where the report and the publishing functions have types: $\text{Rep}: \{0, 1\}^{k \times p(k)} \rightarrow \{0, 1\}^{r(k)}$ and $\text{Pub}: \{0, 1\}^{k \times p(k)} \rightarrow \{0, 1\}^{d(k)}$:*

$$r(k) + d(k) \geq k \cdot p(k) - O(\log k) ,$$

i.e. the amount of data sent to the auditor is comparable to the list of all time stamps.

Actually, the last corollary holds for a weaker security notion – *simulatability*:

Definition 8 (Simulatability). TS_R is simulatable, if for every $\text{Client}, A \in \text{FP}$ there is $A' \in \text{FP}$, so that $\text{VIEW}_{\text{Client}}\langle \text{Client}, TS_R, A \rangle \approx \text{VIEW}_{\text{Client}}\langle \text{Client}, TS_I, A' \rangle$.

Like the Universal Composability, also the Simulatability implies both the Consistency (1) and the security against Random Back-Dating (2) but not the other way round. So it is still possible that one can compress the published information and still have a provably secure auditable time-stamping scheme in the sense of (1) and (2). One such construction is presented in [9] but their security reduction is very inefficient.

The Simulatability (and the Universal Composability) conditions depend on the definitions of TS_I and TS_R . It is not completely excluded that these definitions can be relaxed (in a reasonable way) so that the compression of the published information becomes possible. However, we cannot even imagine how this could be done.

7 Discussion on Practical Implementation

As in the schemes with audit, all time stamp requests are sent from Stamper to Auditor who then performs the same hash computations. Hence, if there are m stampers in the scheme and each stamper performs p hash operations per round, then the auditor must perform $m \cdot p$ hash operations per round. Hence, the cost of the service increases by a constant factor, no matter how many users there are.

In the schemes described above, we have only one trusted Auditor. As one of our main goal was to develop measures against insider attacks, it is reasonable to assume that also the Auditor can be malicious. A natural approach would be to replace a trusted Auditor with a list $\text{Auditor}_1, \dots, \text{Auditor}_n$ of auditors and use the secure multi-party computation. A simplified approach would be that a Stamper sends the digest d and the report r to all auditors in the list. The auditors check whether $\text{Aud}(r, d) = 1$ and send r , d and the result (of the check) to Repository who then decides by clear majority which value to publish. This works if the Repository and $\frac{n+1}{2}$ auditors are honest.

Acknowledgements. The authors are grateful to Estonian Science Foundation for supporting the study.

References

1. Michael Backes. *Cryptographically Sound Analysis of Security Protocols*. PhD thesis, Universität des Saarlandes, 2002.
2. Michael Backes and Birgit Pfitzmann. Symmetric Encryption in a Simulatable Dolev-Yao Style Cryptographic Library. In *17th IEEE Computer Security Foundations Workshop*, Pacific Grove, CA, June 2004.
3. Michael Backes, Birgit Pfitzmann, and Michael Waidner. Symmetric authentication within a simulatable cryptographic library. In Einar Snekkenes and Dieter Gollmann, editors, *Computer Security - ESORICS 2003, 8th European Symposium on Research in Computer Security*, volume 2808 of *LNCS*, pages 271–290. Gjøvik, Norway, October 2003. Springer-Verlag.
4. Michael Backes, Birgit Pfitzmann, and Michael Waidner. A Universally Composable Cryptographic Library. In *Proceedings of the 10th ACM Conference on Computer and Communications Security*, Washington, DC, October 2003. ACM Press.

5. Dave Bayer, Stuart Haber, and W.-Scott Stornetta. Improving the efficiency and reliability of digital time-stamping. In *Sequences II: Methods in Communication, Security, and Computer Science*, pp.329-334, Springer-Verlag, New York 1993.
6. Josh Benaloh and Michael de Mare. Efficient broadcast time-stamping. Tech. report 1, Clarkson Univ. Dep. of Mathematics and Computer Science, August 1991.
7. Ahto Buldas, Peeter Laud, Helger Lipmaa, and Jan Willemson. Time-Stamping with Binary Linking Schemes. In *Advances in Cryptology – CRYPTO’98, LNCS 1462*, pp. 486-501, 1998.
8. Ahto Buldas, Peeter Laud, Märt Saarepera, and Jan Willemson. Universally Composable Time-Stamping Schemes with Audit. IACR ePrint Archive, 2005.
9. Ahto Buldas and Märt Saarepera. On provably secure time-stamping schemes. In *Advances in Cryptology – ASIACRYPT 2004, LNCS 3329*, pp.500–514, 2004.
10. Ran Canetti. A unified framework for analyzing security of protocols. *Electronic Colloquium on Computational Complexity (ECCC)*, 8(16), 2001.
11. Ran Canetti. Security and composition of multi-party cryptographic protocols. *Journal of Cryptology*, 13(1):143-202, 2000.
12. Ran Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*, pp. 136–145. 2001.
13. Ran Canetti and Marc Fischlin. Universally Composable Commitments. In *CRYPTO’01, LNCS 2139*, pp. 19–40. 2001.
14. D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
15. Stuart Haber and W.-Scott Stornetta. How to time-stamp a digital document. *Journal of Cryptology*, Vol. 3, No. 2, pp. 99-111 (1991).
16. Stuart Haber and W.-Scott Stornetta. Secure Names for Bit-Strings. In *ACM Conference on Computer and Communications Security*, pp. 28–35, 1997.
17. Yehuda Lindell. *Composition of Secure Multi-Party Protocols. A Comprehensive Study.* LNCS 2815. 2003.
18. Michael Luby. *Pseudorandomness and cryptographic applications.* Princeton University Press, 1996.
19. Ralph C. Merkle. Protocols for public-key cryptosystems. *Proceedings of the 1980 IEEE Symposium on Security and Privacy*, pp.122-134, 1980.
20. Tal Moran, Ronen Shaltiel and Amnon Ta-Shma. Non-interactive timestamping in the bounded storage model. In *Advances in Cryptology – CRYPTO 2004, LNCS 3152*, 2004.
21. Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. Proceedings of the *Twenty First Annual ACM Symposium on Theory of Computing*. May 15–17 1989: Seattle, ACM Press, pp. 33–43, 1989.
22. Birgit Pfitzmann, Matthias Schunter, and Michael Waidner. Cryptographic Security of Reactive Systems. In Steve Schneider and Peter Ryan, editors, *Workshop on Secure Architectures and Information Flow*, volume 32 of *Electronic Notes in Theoretical Computer Science*, Royal Holloway, University of London, 2000. Elsevier Science.
23. Birgit Pfitzmann and Michael Waidner. Composition and integrity preservation of secure reactive systems. In *CCS 2000, Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 245–254, Athens, Greece, November 2000. ACM Press.
24. Birgit Pfitzmann and Michael Waidner. A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission. In *2001 IEEE Symposium on Security and Privacy*, pages 184–200, Oakland, California, May 2001. IEEE Computer Society Press.
25. Alexander Russell. Necessary and sufficient conditions for collision-free hashing. *Journal of Cryptology* (1995) 8: 87–99.
26. www.surety.com
27. www.authentidate.com
28. www.digistamp.com