

One-Way Signature Chaining - A New Paradigm For Group Cryptosystems And E-Commerce

Amitabh Saxena and Ben Soh
Dept. of Computer Science and Computer Engineering
La Trobe University, Bundoora, VIC, Australia 3086

November 28, 2006

Abstract

In this paper, we describe a new cryptographic primitive called (*One-Way Signature Chaining*). Signature chaining is essentially a method of generating a chain of signatures on the same message by different users. Each signature acts as a “link” of the chain. The *one-way*-ness implies that the chaining process is one-way in the sense that more links can be easily added to the chain. However, it is computationally infeasible to remove any intermediate links without removing all the links. The signatures so created are called chain signatures. We give precise definitions of chain signatures and discuss some applications in trust transfer. We also present a practical construction of a CS scheme that is secure under the Computational Diffie-Hellman (CDH) assumption in bilinear maps.

1 Introduction

Over recent years, a lot of research in e-commerce systems has been on the problem of *trust transfer*. We will attempt to give a formal treatment of trust transfer using the notion of *proxy* signatures in which the original signer delegates the signing authority to a proxy signer [1]. In other words, the original signer “transfers” his trust to the proxy.

Roughly speaking, trust transfer is the act of transferring the trust placed on the original user (the *trusted*) to a proxy user (the *trustee*) such that some other user (the *truster*) can delegate the same responsibilities to the trustee that he would have delegated to the trusted in some *trust context* (for instance, a electronic transaction). Trust transfer makes sense only in a non-interactive environment where the original user is no longer available for interaction once the trust delegation is complete (in other words, the truster can only interact with the trustee).

Once the act of delegation is complete, the trusted and trustee are said to be connected in a *trust relationship*. It may be the case that a truster is a trustee as well. For instance, in an e-commerce scenario, there may be multiple entities involved in a transaction such that interaction is not permitted (or feasible) between many of them.¹ Thus, in this scenario a truster will act as trustee for the next entity in the chain such that trust is transferred from the original user to the last user via the chain of trustees. More generally, where the participants of a distributed protocol (i.e. involving more than two entities) are required to authenticate to each other and non-repudiation is required, the number of interactions between participants need to be minimized.

In this paper, we propose the concept of *one-way chaining* to demonstrate how a chain of trustees can be constructed by forming individual trust relationships between the links of the chain without having to interact with any other links. Intuitively CS are similar to transitive signatures [4] in that they allow trust to be transferred between multiple entities. The difference is that transitive signatures hide the intermediate nodes of trust transfer, while CS try to ensure that intermediate nodes cannot be removed.

¹As an example, in a mobile agent scenario, each sending platform must convince the receiving platform about the trustworthiness of the agent code [2, 3].

In many ways, chain signatures exhibit properties similar to Verifiably Encrypted Signatures (VES) (eg, of [5]). However, there are some inherent differences: A VES is created entirely by one entity, while a chain signature is defined only when two or more entities are involved in creating the signature.

The rest of this paper is organized as follows. We give the motivation for chain signatures in Section 2. We then formalize the intuition of chain signatures in Section 3.2. Finally, in Section 4 we present the scheme and prove its security in Section 4.3. We then discuss some applications of chain signatures in section 5.

2 Motivation

To see the motivation behind chain signatures consider a scenario with three users Alice, Bob and Carol. Alice signs some message m and sends the signature σ_A to Bob, after which she is not available for interaction. Now Bob wants to convince Carol that Alice indeed signed the message m . However if Carol later wants to convince a third party (using Bob’s proof) the statement “**Alice signed the message m** ” then she must only be able to do so by proving that “**Bob knows about it too**”.

2.1 Intuition Behind Chain Signatures

Assume that users A, B compute signatures σ_A, σ_B (on the same message) using private keys SK_A, SK_B respectively. Define the following properties:

1. **Aggregation:** Given signatures σ_A, σ_B it is easy to compute a combined signature $\sigma_{\{A,B\}}$ that can be verified using public keys PK_A, PK_B .
2. **Delete Protection:** Given $\{\sigma_{\{A,B\}}, PK_A, PK_B\}$, it must be infeasible to compute σ_A or σ_B
3. **Strong Delete Protection:** This is a stronger variant of the previous property.
 - Given $\{\sigma_{\{A,B\}}, PK_A, PK_B, SK_A\}$, it must be infeasible to compute σ_B .
 - Given $\{\sigma_{\{A,B\}}, PK_A, PK_B, SK_B\}$, it must be infeasible to compute σ_A .

A signature scheme that satisfies the aggregation and delete protection conditions is called a *Chain Signature* (CS) scheme. A CS scheme that additionally satisfies the strong delete protection condition is called a *Strong Chain Signature* (SCS) scheme. The above idea can be extended to an arbitrary number of users.

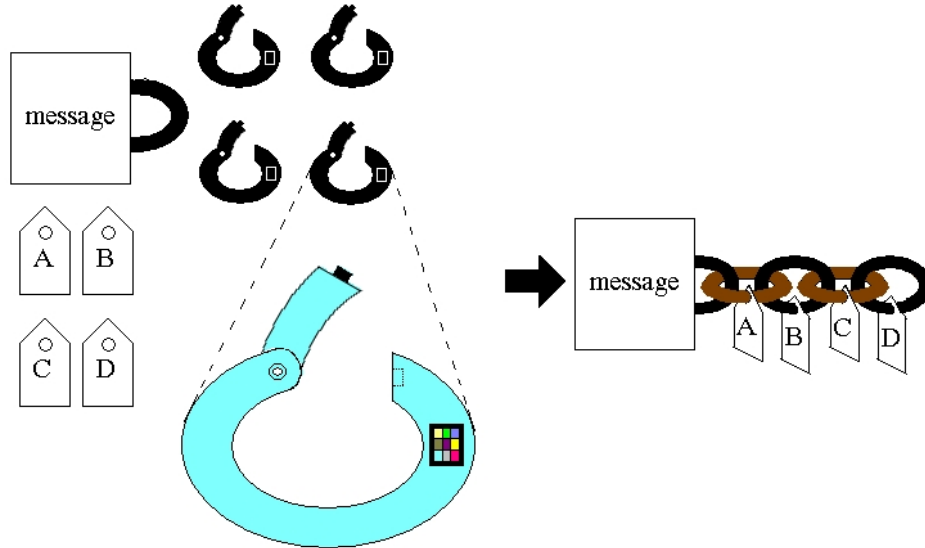
It is fairly trivial to extend the previous argument to arbitrary number of distinct users. Assume that users $1, 2, \dots, n$ compute signatures $\sigma_1, \sigma_2, \dots, \sigma_n$ (on the same message) using private keys SK_1, SK_2, \dots, SK_n respectively. Then we can similarly define:

1. **Aggregation:** Given signatures $\sigma_1, \sigma_2, \dots, \sigma_n$, it is easy to compute a combined signature $\sigma_{\{1,2,\dots,n\}}$ that can be verified using public keys PK_1, PK_2, \dots, PK_n .
2. **Delete Protection:** Given $\{\sigma_{\{1,2,\dots,n\}}, PK_1, PK_2, \dots, PK_n\}$, it must be infeasible to compute σ_α for any $\alpha \subsetneq \{1, 2, \dots, n\}$
3. **Strong Delete Protection:** Given $\{\sigma_{\{1,2,\dots,n\}}, PK_1, PK_2, \dots, PK_n, SK_{\beta_1}, SK_{\beta_2}, \dots, SK_{\beta_i}\}$ for $i < n$ and $\{\beta_1, \beta_2, \dots, \beta_i\} \subsetneq \{1, 2, \dots, n\}$, it must be infeasible to compute σ_α for $\alpha = \{1, 2, \dots, n\} \setminus \{\beta_1, \beta_2, \dots, \beta_i\}$.

Although in the above (informal) description we assumed that the combined signature is “unordered”, in our formal definition we will also take into account the order in which the users contribute. Our model of CS will not provide “strong delete protection”. However, it will provide “delete protection”.

2.2 Physical Analogue Of Chain Signatures

Chain signatures can be intuitively visualized by considering a box with a link and a set of “intermediate” links with an *asymmetric* combination lock, as shown in the figure. In an asymmetric combination lock, the opening combination is different from the closing combination and cannot be derived from it. The initiator sends the box along with several open links and their closing combination(s). The opening combination(s) are kept secret. Each user may then add a tag to the “message”, which is the equivalent of authentication.



3 Formal Definition - Chain Signatures

In this section, we will formalize the above intuition of chain signatures. Since chain signatures inherently deal with ordered elements (i.e. the public keys), we first develop some notation to deal with ordered elements, which we call sequences.

1. A *sequence* is similar to a set except that the order of its elements matters. We require that the elements of a sequence must be distinct. The elements of a sequence are written in order and enclosed with $\langle \rangle$ symbols. For instance, $\langle y_1, y_2, y_3 \rangle$.
2. The symbol θ denotes the empty sequence with zero elements. The symbol ϵ denotes the empty string of zero length.
3. Let $L_a = \langle y_1, y_2, \dots, y_k \rangle$ be some non-empty sequence. For any other sequence L_b , we say that $L_b \prec L_a$ if and only if $L_b = \langle y_1, y_2, \dots, y_i \rangle$ and $0 \leq i \leq k$.
4. We say that two sequences $\{L_a, L_b\}$ **overlap** if there exists a non-empty sequence L' such that $L' \prec L_a$ and $L' \prec L_b$. For instance, $\{\langle y_1, y_2 \rangle, \langle y_1 \rangle\}$ overlap, while $\{\langle y_1, y_2 \rangle, \langle y_2 \rangle\}$ do not.
5. For any two sequences L_a, L_b , the symbol $L_a \cup L_b$ denotes the **set** of elements that belong to at least one of $\{L_a, L_b\}$. Similarly $L_a \cap L_b$ denotes the **set** of elements that belong to both L_a and L_b . We denote by $L_a \odot L_b$ to be the **set** of elements from the largest sequence L' such that $L' \prec L_a$ and $L' \prec L_b$. Clearly, for two overlapping sequences $\{L_a, L_b\}$, we have that $L_a \odot L_b \neq \emptyset$.

3.1 Algorithms

A chain signature scheme is defined using three PPT algorithms **KeyGen**, **ChainSign**, **ChainVerify** as follows. (It is more convenient to describe **ChainVerify** before **ChainSign**.)

KeyGen (Key Generation) This *randomized* algorithm takes as input a security parameter τ and outputs a randomly selected key-pair (x, y) such that x is the private key and y is the public key. We say that $(x_i, y_i) \stackrel{R}{\leftarrow} \text{KeyGen}$ on the i^{th} run of this algorithm.

ChainVerify (Verification) This algorithm takes as input a tuple (m, σ_i, L_i) . Here $L_i = \langle y_1, y_2, \dots, y_i \rangle$ is some sequence of i public keys and the pair (σ_i, L_i) is a purported chain signature on message m . The algorithm works as follows:

1. If $L_i = \theta$ and $\sigma_i = \epsilon$ the algorithm outputs **VALID** and terminates.
2. If $L_i = \theta$ and $\sigma_i \neq \epsilon$ the algorithm outputs **INVALID** and terminates.
3. If this step is executed then $L_i \neq \theta$. The algorithm invokes a deterministic poly-time procedure after which it outputs either **VALID** or **INVALID** and terminates.

ChainSign (Signing) The **ChainSign** algorithm takes as input a tuple $(x_i, y_i, m, \sigma_j, L_j)$. Here (x_i, y_i) is a valid private-public key-pair (generated using the **KeyGen** algorithm), the pair (σ_j, L_j) is a purported chain signature on message m , and $L_j = \langle y_1, y_2, \dots, y_j \rangle$ is some sequence of j public keys such that $y_i \notin \{y_1, y_2, \dots, y_j\}$. The algorithm works as follows:

1. If any of the input conditions (as described above) are violated, the algorithm outputs **ERROR** and terminates.
2. The algorithm invokes **ChainVerify** with (m, σ_j, L_j) as input (i.e. it checks whether (σ_j, L_j) is a valid chain signature on m or not). If (σ_j, L_j) is not a valid chain signature on message m , the algorithm outputs **ERROR** and terminates.
3. If this step is executed then no input conditions are violated and (σ_j, L_j) is a valid chain signature on m . In this case this algorithm uses the private key x_i to compute a new valid chain signature (σ_i, L_i) on message m such that $L_i = \langle y_1, y_2, \dots, y_j, y_i \rangle$. It outputs (σ_i, L_i) and terminates.

The **ChainVerify** and **ChainSign** algorithms must satisfy the standard consistency constraint of signatures. That is, if the input (m, σ_i, L_i) to the **ChainVerify** is the output of the **ChainSign** algorithm then the **ChainVerify** algorithm must output **VALID**. Note that **ChainSign** can be initialized by setting $\sigma_j = \epsilon$ and $L_j = \theta$.

3.2 Security Model

We define adaptive security of chain signatures using Game 1 below. For simplicity, we assume that the adversary is not allowed to use a chosen private key. The adversary is, however, allowed to extract private keys of choice. In this respect, our model is similar to an identity based system. We call this *adaptive security under known key and chosen message attack*.

The reader should note that this is a weaker model than *adaptive security under chosen key and chosen message attack* used in the aggregate signatures of [5]. We feel, however, that our notation is more suitable in modeling the requirements of chain signatures (which are slightly different from aggregate signatures).

Game 1

1. *Setup*: The challenger \mathcal{C} sets a security parameter τ and gives it to the adversary. The adversary \mathcal{A} selects the parameter n . The challenger generates n key-pairs $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\} \stackrel{R}{\leftarrow} \text{KeyGen}$ and gives the set $Y = \{y_1, y_2, \dots, y_n\}$ of n public keys to the adversary. Denote by L the set of all non-empty sequences with elements from Y .


2. *Queries*: Working adaptively, the adversary \mathcal{A} issues at most q_s (chain) sign queries and q_e (private-key) extract queries as follows:
 - (a) *ChainSign*: A (chain) sign query i ($1 \leq i \leq q_s$) consists of a pair $(m_{s(i)}, L_{s(i)}) \in \Sigma^* \times L$. The challenger responds with a valid **chain signature** $(\sigma_{s(i)}, L_{s(i)})$ on $m_{s(i)}$ computed using the **ChainSign** algorithm.
 - (b) *Extract*: An extract query j ($1 \leq j \leq q_e$) consists of a public key $y_{e(j)}$. The challenger responds with the private key $x_{e(j)}$.
3. *Output*: Finally \mathcal{A} outputs a message-**chain-signature** pair $(m_A, (\sigma_A, L_A))$.
4. *Result*: \mathcal{A} wins the game if all the following conditions hold:
 - (a)
 - i. $L_A \in L$ and the **ChainVerify** algorithm accepts $(m_A, (\sigma_A, L_A))$ as valid.
 - ii. No sign query has been previously made on the pair (m_A, L_A) .
 - iii. At least one private key corresponding to L_A has not been extracted.
 - (b) For any sign query i , if $(m_{s(i)} = m_A) \wedge (\{L_A, L_{s(i)}\} \text{ overlap})$, then there is at least one key in $(L_{s(i)} \cup L_A) \setminus (L_{s(i)} \odot L_A)$ which has not been extracted.

We use the random oracle model [6], where a hash function is implemented using a random oracle. If the adversary needs to compute a hash value it queries the random oracle, which is also simulated by the challenger. The requirement of a fair game is that the responses of the challenger (to hash queries) are indistinguishable from the responses of a random oracle.


Definition 3.1. *We say that the chain signature scheme is $(n, \tau, t, q_s, q_e, q_h, \epsilon)$ -secure under an adaptive known-key and chosen-message attack if, for some parameters τ and n , there is no adversary \mathcal{A} that runs for at most time t ; makes at most q_s sign queries; makes at most q_e extract queries; makes at most q_h hash queries; and wins Game 1 with probability at least ϵ . Otherwise, we say that \mathcal{A} $(n, t, q_s, q_e, q_h, \epsilon)$ -breaks the chain signature scheme under an adaptive known-key and chosen-message attack.*

3.3 Analysis Of Game 1

Let us analyze the Result Section of Game 1. Clearly, Part (a) rules out the cases of a trivial win. The intuition of CS is captured in Part (b). To see this, consider the illustration of the **ChainSign** algorithm at some stage i .



Since we know that **ChainSign** can be “reversed” from σ_i (using the private key x_i) to obtain σ_{i-1} , we can also consider the following figure as a valid usage of this algorithm.



Therefore, if L_A and $L_{s(i)}$ overlap for some i and $m_A = m_{s(i)}$ (we call such queries *non-trivial queries*), then we know that **ChainSign** can be “reversed” from $\sigma_{s(i)}$ and then “forwarded” to obtain σ_A using only a subset of private keys for L_A .

The security requirement of CS is that this is the only other way to generate σ_A without knowing all the private keys of L_A . The condition of Part (b) of the Result Section implies that if the adversary does not know at least one private key needed for this “reverse-forward” operation then this is a valid win for the adversary.

3.4 Differences With Other Signature Schemes

In this section, we briefly demonstrate how chain signatures are different from other multi-user signature schemes such as sequential aggregate signatures [7, 8], multisignatures [9], aggregate signatures [5], and structured multisignatures [10]. We distinguish between two types of forgers for chain signatures.

Type 1 Forger This adversary either does not make any non-trivial queries or, if it makes one or more non-trivial queries then for each non-trivial query j , we have that $L_A \not\prec L_{s(j)}$. We call such a forger an *Ordinary Forgery*.

Type 2 Forger This adversary makes one or more non-trivial queries j such that $L_A \prec L_{s(j)}$. We call such a forger an *Extraction Forgery*.

All the above mentioned signature schemes only consider type 1 forgers, while chain signatures also consider type 2. To see this, consider the following instance of Game 1 with $n = 7$ and $Y = \{y_i | 1 \leq i \leq n\}$. The adversary outputs a valid message-chain-signature pair $(m_A, (\sigma_A, L_A))$ after making five extract queries on the keys $\{y_1, y_2, y_3, y_5, y_6\}$ and three non-trivial sign queries i (i.e, $m_A = m_{s(i)}$) such that $L_{s(i)} \odot L_A \neq \emptyset$. The sequences are (keys of extract queries have a gray box):

$$L_{s(1)} = \langle \boxed{y_1}, \boxed{y_2}, \boxed{y_3}, y_4, \boxed{y_5}, \boxed{y_6}, y_7 \rangle$$

$$L_{s(2)} = \langle \boxed{y_1}, \boxed{y_2}, \boxed{y_3}, y_4, \boxed{y_6}, y_7 \rangle$$

$$L_{s(3)} = \langle \boxed{y_1}, \boxed{y_2}, \boxed{y_3} \rangle$$

$$L_A = \langle \boxed{y_1}, \boxed{y_2}, \boxed{y_3}, y_4, \boxed{y_5} \rangle$$

Since for all the sequences $L_{s(i)}$ ($1 \leq i \leq 3$), at least one private key needed for the reverse-forward operation (described earlier) has not been extracted, the above configuration represents a win for the adversary of Game 1. The same configuration, however, represents a loss for the adversary of a suitably adapted game (adaptive chosen-key and chosen-message attack) in all the above mentioned schemes ([7, 8, 9, 5, 10]) when we keep y_4 as the challenge public key.

4 Chain Signatures Using Bilinear Maps

In this section we give a concrete example of chain signatures using bilinear maps. In our discussion, we will only use an abstract definition of such maps. The interested reader is referred to [11] for information on actual implementation of such maps.

4.1 Bilinear Maps

Let G_1 and G_2 be two cyclic multiplicative groups both of prime order q such that computing discrete logarithms in G_1 and G_2 is intractable. A bilinear pairing is a map $\hat{e} : G_1 \times G_1 \mapsto G_2$ that satisfies the following properties [11, 12]:

1. *Bilinearity*: $\hat{e}(a^x, b^y) = \hat{e}(a, b)^{xy} \forall a, b \in G_1$ and $x, y \in \mathbb{Z}_q$.
2. *Non-degeneracy*: If g is a generator of G_1 then $\hat{e}(g, g)$ is a generator of G_2 .
3. *Computability*: The map \hat{e} is efficiently computable.

The above properties also imply:

$$\begin{aligned}\hat{e}(ab, x) &= \hat{e}(a, x) \cdot \hat{e}(b, x) \quad \forall a, b, x \in G_1 \\ \hat{e}(a, xy) &= \hat{e}(a, x) \cdot \hat{e}(a, y) \quad \forall a, x, y \in G_1\end{aligned}$$

In a practical implementation, G_1 is a subgroup of the (additive²) group of points on the elliptic curve and G_2 is the multiplicative subgroup of a finite field. The map \hat{e} is derived either from the modified Weil pairing [11, 12] or the Tate pairing [13]. Without going into the details of generating suitable curves (since the same parameters of [5] will suffice³), we will assume that $q \geq 2^{171}$ so that the fastest algorithm for computing discrete logarithms in G_1 (Pollard's rho method [14, p.128]) takes $\geq 2^{85}$ iterations [11].

4.1.1 BDH Parameter Generator

Using the idea of [11], we define a Bilinear Diffie-Hellman (BDH) parameter generator **BDH** as a randomized algorithm that takes as input $\tau \in \mathbb{N}$ and outputs (\hat{e}, q, G_1, G_2) , where G_1, G_2 are the descriptions of two cyclic multiplicative groups, each of prime order q such that $q \approx 2^\tau$, and $\hat{e} : G_1 \times G_1 \mapsto G_2$ is a bilinear map as defined above.

4.1.2 The Computational Diffie-Hellman (CDH) Problem

The security of chain signatures depends on the hardness of the following problem.

- **Computational Diffie-Hellman (CDH_(g,G₁)) Problem:** Let $g \stackrel{R}{\leftarrow} G_1$ be a generator of G_1 . Given $g^x, g^y \in G_1$ for unknown $x, y \in \mathbb{Z}$, output $g^{xy} \in G_1$.

The Computational Diffie-Hellman Assumption (CDHA) states that the CDH_(g,G₁) problem is intractable for any PPT adversary. This is formally stated below.

CDHA. Let \mathcal{A} be an algorithm, and $\nu : \mathbb{N} \mapsto [0, 1]$ a function. We associate with any $\tau \in \mathbb{N}$ the following experiment.

Experiment $\mathbf{Exp}_{\mathcal{A}}^{cdh}(\tau)$

$(\hat{e}, q, G_1, G_2) \leftarrow \mathbf{BDH}(\tau); g \stackrel{R}{\leftarrow} G_1 \setminus \{1\}; \alpha, \beta \stackrel{R}{\leftarrow} \mathbb{Z}_q; a \leftarrow g^\alpha, b \leftarrow g^\beta$

$c \leftarrow \mathcal{A}(\hat{e}, q, G_1, G_1, g, a, b)$

If $c = g^{\alpha\beta}$ return 1 else return 0

We let

$$\mathbf{Adv}_{\mathcal{A}}^{cdh}(\tau) = \Pr [\mathbf{Exp}_{\mathcal{A}}^{cdh}(\tau) = 1]$$

denote the advantage of \mathcal{A} on input τ , the probability over the random choices of the inputs to \mathcal{A} and the coins of \mathcal{A} if any. We say that \mathcal{A} has success bound ν for τ if $\mathbf{Adv}_{\mathcal{A}}^{cdh}(\tau) \leq \nu(\tau)$

Assumption CDHA: For all poly-time algorithms \mathcal{A} there exists a negligible function ν and an integer $\tau' \in \mathbb{N}$, such that for all $\tau > \tau'$, \mathcal{A} has success bound ν .

Definition 4.1. We say that algorithm $\mathcal{A}(\epsilon, t, \tau)$ -breaks the BDH parameter generator **BDH** if \mathcal{A} runs for time at most t , and $\mathbf{Adv}_{\mathcal{A}}^{cdh}(\tau)$ is at least ϵ .

²Although it is conventional to use the additive notation for the group G_1 , it is more convenient for us to use the multiplicative one.

³Boneh et al.'s aggregate signature scheme of [5] uses an *asymmetric* pairing $G_0 \times G_1 \mapsto G_2$ such that an efficiently computable isomorphism $\psi : G_1 \mapsto G_0$ exists. Their construction can be directly adapted to our's simply by setting $G_0 = G_1$.

4.2 Chain Signature Protocol

In this scenario, n ordered distinct users $\langle 1, 2, \dots, n \rangle$, $m \in \Sigma^*$ is the contract.

System Parameters A Trusted Authority (TA) sets a parameter τ and generates $(\hat{e}, G_1, G_2, q) \xleftarrow{R} \text{BDH}(\tau)$.

Here, G_1, G_2 are descriptions of two groups each of prime order $q \approx \tau$ bits and $\hat{e} : G_1 \times G_1 \mapsto G_2$ is a bilinear mapping as defined above. The TA selects a one-way hash function $\mathcal{H} : \Sigma^* \mapsto G_1$ and picks a random generator g of G_1 . The system parameters are $\langle e, q, G_1, G_2, \mathcal{H}, g \rangle$.

KeyGen Each participant i generates $x_i \xleftarrow{R} \mathbb{Z}_q^*$ as the private key. The corresponding public key is $y_i = g^{x_i} \in G_1$.

ChainSign Let $L_i = \langle y_1, y_2, \dots, y_i \rangle$ and $h_i = \mathcal{H}(m, L_i)$ for $i \geq 1$. Define $\sigma_0 = 1 \in G_1$ and define recursively

$$\sigma_i = \sigma_{i-1} h_i^{x_i} \in G_1 \text{ for } i \geq 1$$

The chain signature of user i on m is (σ_i, L_i) .

ChainVerify We accept the signature (σ_i, L_i) on m as valid if the following equality holds:

$$\hat{e}(\sigma_i, g) \stackrel{?}{=} \prod_{j=1}^i \hat{e}(h_j, y_j)$$

The correctness of the verification process follows from the property of bilinear maps:

$$LHS = \hat{e}(\sigma_i, g) = \hat{e}\left(\prod_{j=1}^i h_j^{x_j}, g\right) = \prod_{j=1}^i \hat{e}(h_j, g^{x_j}) = RHS$$

4.3 Security Of The Construction

We use Definition 3.1 of Section 3.2 to prove adaptive security of the chain signature scheme. We will reuse the proof of security of the signature scheme of Boneh, Lynn and Shacham [11] (hereafter called BLS). The scheme is defined as follows.

Preliminary Setup For BLS Signatures A Trusted Authority (TA) sets a parameter τ and generates

$(\hat{e}, G_1, G_2, q) \xleftarrow{R} \text{BDH}(\tau)$. Here, G_1, G_2 are descriptions of two groups each of prime order $q \approx \tau$ bits and $\hat{e} : G_1 \times G_1 \mapsto G_2$ is a bilinear mapping as defined above. The TA selects a one-way hash function $\mathcal{H} : \Sigma^* \mapsto G_1$ and picks a random generator g of G_1 . The system parameters are $\langle e, q, G_1, G_2, \mathcal{H}, g \rangle$.

1. **KeyGen** Generate $x \xleftarrow{R} \mathbb{Z}_q^*$ as the private key. The public key is $y = g^x \in G_1$.
2. **Sign** To sign message m under public key y , compute $h = \mathcal{H}(m)$ and $\sigma = h^x \in G_1$. The algorithm outputs σ as a valid signature on message m under public key y .
3. **Verify** Accept the signature σ under the public key y as valid if the following holds:

$$\hat{e}(\sigma, g) \stackrel{?}{=} \hat{e}(\mathcal{H}(m), y)$$

Security of BLS signatures is defined using Game 2.

Game 2

1. *Setup*: The challenger sets some security parameter τ' and generates a key pair $(x, y) \xleftarrow{R} \text{KeyGen}$. It gives y , the challenge public key, along with τ' to the adversary.

2. *Queries:* Working adaptively, \mathcal{A} issues at most q'_s sign queries and q'_h hash queries:
 - (a) *Sign queries:* For each sign query i on distinct messages m_i for $1 \leq i \leq q'_s$, the challenger responds with a valid BLS signature $\sigma_i = \mathcal{H}(m_i)^x \in G_1$.
 - (b) *Hash queries:* For each hash query j on distinct messages m_h for $1 \leq j \leq q'_h$, the challenger responds with $\mathcal{H}(m_j)$.
3. *Output:* Finally \mathcal{A} outputs a BLS signature σ_A .
4. *Result:* \mathcal{A} wins if σ_A is a valid signature and no sign query was issued on m_A .

Definition 4.2. We say that the BLS scheme is $(\tau', t', q'_s, q'_h, \epsilon')$ -secure against existential forgery under an adaptive **chosen message** attack if for some parameter τ' , there is no adversary \mathcal{A} that runs for at most time t' ; makes at most q'_s sign queries and q'_h hash queries; and wins Game 2 with probability at least ϵ' . Otherwise, if such an adversary \mathcal{A} exists, then we say that \mathcal{A} $(\tau', t', q'_s, q'_h, \epsilon')$ -breaks the BLS scheme.

To prove the security of chain signatures, we will use the following result from [11].

Theorem 4.3. ([11, Theorem 3.2]) If there exists an algorithm \mathcal{A} that $(\tau', t', q'_s, q'_h, \epsilon')$ -breaks the BLS scheme under Definition 4.2, then there exists another algorithm \mathcal{B} that $(t'', \epsilon'', \tau'')$ -breaks the BDH parameter generator **BDH** under Definition 4.1, where;

$$\tau'' = \tau'; \quad t'' \leq t' - c'_{G_1}(q'_h + 2q'_s); \quad \epsilon'' \geq \frac{\epsilon'}{e(q'_s + 1)}$$

Here, c'_{G_1} is a constant that depends on G_1 and e is the base of natural logarithms.

Our security follows from the following Theorem.

Theorem 4.4. If there exists an algorithm \mathcal{A} that $(n, \tau, t, q_s, q_e, q_h, e)$ -breaks the chain signature scheme under Definition 3.1, then there exists another algorithm \mathcal{B} that $(\tau', t', q'_s, q'_h, \epsilon')$ -breaks the BLS scheme under Definition 4.2, where;

$$\tau = \tau'; \quad t' \leq t + c_{G_1} [n(q_h + 1) + q_s(n + 1)]; \quad q'_s \leq q_s; \quad q'_h \leq n(q_h + q_s); \quad \epsilon' \geq \frac{\epsilon(q_e - 1)}{eq_e^2}$$

Here, c_{G_1} is a constant that depends on G_1 and e is the base of natural logarithms.

Proof. Let there exist an adversary \mathcal{A} that $(n, \tau, t, q_s, q_e, q_h, e)$ -breaks the chain signature scheme under an adaptive known key and chosen message attack. Using \mathcal{A} , we construct another algorithm \mathcal{B} that $(\tau', t', q'_s, q'_h, \epsilon')$ -breaks the BLS signature scheme under an adaptive chosen key and chosen message attack.

Algorithm \mathcal{B} simulates the adversary of Game 2 and is given a challenge public key $(g, y) = (g, g^x) \in G_1^2$ (for unknown x), along with a security parameter τ' by challenger \mathcal{C} . Its goal is to forge a valid BLS signature under y using adversary \mathcal{A} , that can win Game 1. Algorithm \mathcal{B} simulates the challenger of Game 1 to adversary \mathcal{A} as follows.

Setup. Algorithm \mathcal{B} maintains a table of n entries called the K-List. Each entry i in the table is a tuple of the form $(a_i, r_i, y_i) \in \{0, 1\} \times \mathbb{Z}_q^* \times G_1$ and is created as follows. First \mathcal{B} generates $r_i \xleftarrow{R} \mathbb{Z}_q$ and a bit $a_i \xleftarrow{R} \{0, 1\}$ using a biased coin such that $\Pr[a_i = 1] = 1/q_e$. It then computes $y_i = y^{a_i} g^{r_i} \in G_1$ and gives the set $Y = \{y_1, y_2, \dots, y_n\}$ as the challenge public keys, along with τ to \mathcal{A} .

Queries. To handle the queries of \mathcal{A} , algorithm \mathcal{B} works as follows.

Extract: For an extract query on some key y_j , algorithm \mathcal{B} scans the K-List to find the tuple (a_j, r_j, y_j) . If $a_j = 1$, Algorithm \mathcal{B} reports **Failure** and terminates. Otherwise, it responds with r_j as the private key for y_j .

Hash: To respond to hash queries, \mathcal{B} maintains another table called the H-List (having up to $n(q_h + q_s)$ entries). Each entry i in the list is of the form,

$$(m_i, L_i, b_i, u_i, h_i, \gamma_i) \in \Sigma^* \times L \times \{0, 1\} \times \mathbb{Z}_q^* \times G_1 \times G_1,$$

and can be interpreted using the following Table.

Entry i	$b_i = 0$	$b_i = 1$
$u_i = 0$	$h_i = \mathcal{H}(m_i, L_i)$ $\gamma_i = \text{signature on } (m_i, L_i)$	$h_i = \mathcal{H}(m_i, L_i)$ $\gamma_i = \text{signature on } (m_i, L_i)/h_j^x$ for some $(m_j, L_j, b_j, u_j, h_j, \gamma_j) \in \text{H-List}$, $m_i = m_j \wedge L_j \prec L_i \wedge b_j = 1 \wedge u_j > 0$
$u_i > 0$	$h_i = g^{u_i}/h_j$ for some $(m_j, L_j, b_j, u_j, h_j, \gamma_j) \in \text{H-List}$, $m_i = m_j \wedge L_j \prec L_i \wedge b_j = 1 \wedge u_j > 0$ $\gamma_i = \text{signature on } (m_i, L_i)$	$h_i = \mathcal{H}(m_i, L_i)$ $\gamma_i = \text{signature on } (m_i, L_i)/h_i^x$
$u_i > 1$		$u_i = h_i^x$

For a hash query j on m_j^* , algorithm \mathcal{B} first parses m_j^* as (m_j, L_j) and scans the H-List for the the unique entry $(m_j, L_j, b_j, u_j, h_j, \gamma_j)$. If such an entry exists, \mathcal{B} returns $h_j = \mathcal{H}(m_j^*)$ as its response to the hash query. Otherwise, \mathcal{B} adds the entry $(m_j, L_j, b_j, u_j, h_j, \gamma_j)$ to the H-List as follows.

- First it parses L_j as $\langle y_{j(1)}, y_{j(2)}, \dots, y_{j(|L_j|)} \rangle$ and scans the K-List to find the entry (a_l, y_l, r_l) such that $y_l = y_{j(|L_j|)}$.
- If $|L_j| > 1$ then \mathcal{B} constructs the sequence $L'_j = \langle y_{j(1)}, y_{j(2)}, \dots, y_{j(|L_j|-1)} \rangle$ and simulates a *Hash* query to itself on the value (m_j, L'_j) .
- Let $(m_j, L'_j, b'_j, u'_j, h'_j, \gamma'_j)$ be the entry in the H-List corresponding to (m_j, L'_j) whenever $|L_j| > 1$.

Algorithm \mathcal{B} then uses the following table to compute the response. It adds $(m_j, h_j, b_j, u_j, h_j, \gamma_j)$ to the H-List and h_j as its response to the hash query.⁴

H-List		$ L_j > 1$	$ L_j = 1$
$a_l = 0$		$h_j \leftarrow \mathcal{H}(m_j, L_j)$ $b_j \leftarrow b'_j; u_j \leftarrow 0$ $\gamma_j \leftarrow \gamma'_j \cdot h_j^{r_l}$	$h_j \leftarrow \mathcal{H}(m_j, L_j)$ $b_j \leftarrow 0; u_j \leftarrow 0$ $\gamma_j \leftarrow h_j^{r_l}$
$a_l = 1$	$b'_j = 0$	$h_j \leftarrow \mathcal{H}(m_j, L_j)$ $b_j \leftarrow 1; u_j \leftarrow 1$ $\gamma_j = \gamma'_j \cdot h_j^{r_l}$	$h_j \leftarrow \mathcal{H}(m_j, L_j)$ $b_j \leftarrow 1; u_j \leftarrow 1$ $\gamma_j \leftarrow h_j^{r_l}$
	$b'_j = 1$	$b_j \leftarrow 0; u_j \xleftarrow{R} \mathbb{Z}_q^*$ $h_j \leftarrow g^{u_j}/h_k$, where, $(m_k, L_k, b_k, u_k, h_k, \gamma_k) \in \text{H-List}$ such that, $m_k = m_j \wedge L_k \prec L_j \wedge b_k = 1 \wedge u_k > 0 \wedge$ L_k is the largest sequence $\gamma_j \leftarrow \gamma'_j \cdot y^{u_j} \cdot h_j^{r_l}$	

ChainSign: For each (chain) sign query i ($1 \leq i \leq q_s$) on $(m_{s(i)}, L_{s(i)})$, algorithm \mathcal{B} scans the H-List to find the unique entry $(m_j, L_j, b_j, u_j, h_j, \gamma_j)$ such that $(m_j, L_j) = (m_{s(i)}, L_{s(i)})$. If such an entry does not exist, \mathcal{B} adds it by simulating a hash query on the message-sequence $(m_{s(i)}, L_{s(i)})$.

⁴**Notes:** (1) Algorithm \mathcal{B} will simulate the hash function \mathcal{H} by making hash queries to \mathcal{C} . (2) If $|L_j| = 0$, algorithm \mathcal{B} simply replies with $h_j = \mathcal{H}(m_j)$ without storing the value in the H-List. (3) Algorithm \mathcal{B} 's replies to \mathcal{A} 's hash queries are truly random, and therefore, the hash function simulation provided by \mathcal{B} is perfect.

1. If $b_j = 1$, algorithm \mathcal{B} finds the entry $(m_k, L_k, b_k, u_k, h_k, \gamma_k) \in \text{H-List}$ such that $(m_k = m_{s(i)}) \wedge (L_k \prec L_{s(i)}) \wedge (b_k = 1) \wedge (u_k > 0) \wedge (L_k \text{ is the **largest** sequence})$.⁵
 1. If $u_k = 1$, algorithm \mathcal{B} sets $u_k \leftarrow \text{Sign}(m_k, L_k) = \mathcal{H}(m_k, L_k)^x \in G_1$ by making a sign query to \mathcal{C} .
 2. \mathcal{B} sets $\sigma_{s(i)} \leftarrow u_k \cdot \gamma_j$ and returns $\sigma_{s(i)}$ as its response to the sign query.
2. If $b_j = 0$, algorithm \mathcal{B} sets $\sigma_{s(i)} \leftarrow \gamma_j$ and returns $\sigma_{s(i)}$ as a valid response to the chain sign query on $(m_{s(i)}, L_{s(i)})$.

Output. Finally \mathcal{A} outputs a message-chain-signature pair $(m_A, (\sigma_A, L_A))$.

Result. If $(m_A, (\sigma_A, L_A))$ is not a winning configuration of Game 1, Algorithm \mathcal{B} reports Failure and terminates.

We know that $(m_A, (\sigma_A, L_A))$ is a winning configuration of Game 1. Algorithm \mathcal{B} scans the H-List to find the unique entry $(m_k, L_k, b_k, u_k, h_k, \gamma_k)$. If such an entry does not exist in the list, it adds it by simulating a hash query on (m_A, L_A) .

1. If $b_A = 0$, algorithm \mathcal{B} reports Failure and terminates.
2. We know that $b_A = 1$. \mathcal{B} finds the entry $(m_k, L_k, b_k, u_k, h_k, \gamma_k) \in \text{H-List}$ such that $(m_k = m_A) \wedge (L_k \prec L_A) \wedge (b_k = 1) \wedge (u_k > 0) \wedge (L_k \text{ is the **largest** sequence})$.⁶
If $u_k > 1$, \mathcal{B} reports Failure and terminates.
3. We know that

$$b_A = 1 \wedge u_k = 1 \tag{1}$$

By definition, $\sigma_A/\gamma_A = h_k^x = \mathcal{H}(m_k, L_k)^x$. Therefore, σ_A/γ_A is a valid BLS signature under public key y on the message (m_k, L_k) .

For any key $y_l \in Y$, define a_l to be the first element of the entry (a_l, r_l, y_l) in the K-List corresponding to public key y_l . Observe that $b_A = 1$ implies that there are an odd number of keys $y_l \in L_A$ such that $a_l = 1$, while $u_k = 1$ implies that there are an even number of keys $y_l \in L_k$ such that $a_l = 1$. Since $u_k = 1$, we can conclude that, \mathcal{B} has **not** made a sign query to \mathcal{C} on the message (m_k, L_k) .⁷

Algorithm \mathcal{B} returns $(\sigma_A/\gamma_A, (m_k, L_k))$ to \mathcal{C} , thereby winning Game 2.

To find the probability of \mathcal{B} winning Game 2, we need to analyze the following events:

- \mathcal{E}_1 : \mathcal{B} does not abort as a result of \mathcal{A} 's extraction queries.
- \mathcal{E}_2 : \mathcal{A} wins Game 1.
- \mathcal{E}_3 : \mathcal{B} does not output Failure in the Result phase.

Clearly, $\epsilon' = \Pr[\mathcal{E}_3 \wedge \mathcal{E}_2 \wedge \mathcal{E}_1] = \Pr[\mathcal{E}_3 | \mathcal{E}_2 \wedge \mathcal{E}_1] \cdot \Pr[\mathcal{E}_2 | \mathcal{E}_1] \cdot \Pr[\mathcal{E}_1]$. The following three claims give the bound on ϵ' .

Claim 1. $\Pr[\mathcal{E}_3 | \mathcal{E}_2 \wedge \mathcal{E}_1] \geq \frac{1}{q_e} \left(1 - \frac{1}{q_e}\right)$

Proof. Assume that $q_e > 1$. Using Equation 1, define the events,

⁵It is possible that $j = k$.

⁶It is possible that $L_k = L_A$.

⁷To visualize this, consider the example given in Section 3.4 with $n = 7$. Assume that $b_5 = b_7 = 1$, while $b_i = 0 \forall y_i$ ($i \neq 5, 7$). In other words, keys y_5 and y_7 are “fake”, while the others are “real”. Since \mathcal{A} never obtained a chain signature on (m_A, L'_A) such that $(L'_A \prec L_A) \wedge (L'_A \text{ contains an odd number of fake keys})$, therefore, \mathcal{B} never made a sign query to \mathcal{C} on the message $(m_A, (y_1, y_2, y_3, y_4))$.

$$\begin{aligned}\mathcal{E}_4 &: b_A = 1 \\ \mathcal{E}_5 &: u_k = 1\end{aligned}$$

$$\therefore \Pr[\mathcal{E}_3|\mathcal{E}_2 \wedge \mathcal{E}_1] = \Pr[\mathcal{E}_4 \wedge \mathcal{E}_5|\mathcal{E}_2 \wedge \mathcal{E}_1] = \Pr[\mathcal{E}_4|\mathcal{E}_5 \wedge \mathcal{E}_2 \wedge \mathcal{E}_1] \cdot \Pr[\mathcal{E}_5|\mathcal{E}_2 \wedge \mathcal{E}_1].$$

$$\text{Clearly, } \mathcal{E}_4 \Rightarrow \sum_{y_i \in L_A} a_i \equiv 1 \pmod{2}, \quad \text{and} \quad \mathcal{E}_5 \Rightarrow \sum_{y_i \in L_k} a_i \equiv 0 \pmod{2}$$

Now consider the Result Section of Game 1 (reproduced below):

- (a) i. $L_A \in L$ and the **ChainVerify** algorithm accepts $(m_A, (\sigma_A, L_A))$ as valid.
 - ii. No sign query has been previously made on the pair (m_A, L_A) .
 - iii. At least one private key corresponding to L_A has not been extracted.
- (b) For any sign query i , if $(m_{s(i)} = m_A) \wedge (\{L_A, L_{s(i)}\} \text{ overlap})$, then there is at least one key in $(L_{s(i)} \cup L_A) \setminus (L_{s(i)} \odot L_A)$ which has not been extracted.

Since, we know that $(m_A, (\sigma_A, L_A))$ is a winning configuration of Game 1, therefore, for all i ($1 \leq i \leq q_s$) and all $L_{s(i)}$ such that $(m_A = m_{s(i)}) \wedge (L_A \odot L_{s(i)} \neq \emptyset)$, the following two variables are **independent** of \mathcal{A} 's view:

$$\sum_{y_i \in L_A} a_i \quad \text{and} \quad \sum_{y_i \in L_{s(i)}} a_i$$

$$\therefore \Pr[\mathcal{E}_4|\mathcal{E}_5 \wedge \mathcal{E}_2 \wedge \mathcal{E}_1] = \Pr[\mathcal{E}_4|\mathcal{E}_2 \wedge \mathcal{E}_1] = \alpha \text{ (say)}$$

$$\text{Also, } \Pr[\mathcal{E}_5|\mathcal{E}_2 \wedge \mathcal{E}_1] = 1 - \Pr[\mathcal{E}_4|\mathcal{E}_2 \wedge \mathcal{E}_1] = 1 - \alpha$$

$$\therefore \Pr[\mathcal{E}_3|\mathcal{E}_2 \wedge \mathcal{E}_1] = \alpha \cdot (1 - \alpha).$$

It can be shown (using standard techniques from calculus) that $1 - 1/q_e \geq \alpha \geq 1/q_e$.

$$\therefore \frac{1}{4} \geq \alpha(1 - \alpha) \geq \frac{1}{q_e} \left(1 - \frac{1}{q_e}\right)$$

from which we get the required bound. \square

Claim 2. $\Pr[\mathcal{E}_2|\mathcal{E}_1] \geq \epsilon$

Proof. If \mathcal{B} did not abort as a result of \mathcal{A} 's extract queries, then the simulation provided by \mathcal{B} is indistinguishable from a real game. Consequently, the probability of \mathcal{A} winning the simulated Game 2 is the same as the probability of \mathcal{A} winning Game 2. Thus, $\Pr[\mathcal{E}_2|\mathcal{E}_1] \geq \epsilon$. \square

Claim 3. $\Pr[\mathcal{E}_1] \geq 1/e$.

Proof. The probability that \mathcal{B} aborts on the i^{th} extract query on y_i depends on the value a_i of the entry (a_i, r_i, y_i) in the H-List. Consider the event,

$$\mathcal{E}_i^e = \mathcal{B} \text{ does not abort on the } i^{\text{th}} \text{ extract query}$$

$$\therefore \text{ If } j > 1, \text{ then, } \Pr[\mathcal{E}_j^e|\mathcal{E}_{j-1}^e] = \Pr[a_* = 0] \quad (\text{For some } a_* \text{ in the K-List.})$$

Since \mathcal{B} did not abort as a result of \mathcal{A} 's $j - 1^{\text{th}}$ extract query, the value a_* must be independent of \mathcal{A} 's view until now, and so,

$$\Pr[\mathcal{E}_j^e|\mathcal{E}_{j-1}^e] = \Pr[\mathcal{E}_1^e] = \Pr[a_1 = 0] = 1 - \frac{1}{q_e}$$

Therefore,

$$\Pr[\mathcal{E}_1] = \left(1 - \frac{1}{q_e}\right)^{q_e} \geq \frac{1}{e},$$

applying the induction hypothesis on q_e extract queries. \square

Combining the results of Claims 1, 2 and 3 above, we get the required bounds on ϵ' .

For each entry in the H-List, \mathcal{B} makes at most one hash query to \mathcal{C} . Also, \mathcal{A} can make up to q_h hash queries on arbitrary message-sequence pairs (m_*, L_*) , and each sequence L_* may contain up to n keys (and therefore up to n sub-sequences). Consequently, each hash query by \mathcal{A} can cause at most n entries to be added to the H-List. Additionally, adversary \mathcal{A} may make up to q_s distinct message-sequence pairs (m_*, L_*) without making any hash queries on them. These sign queries may cause up to nq_s more entries to be added to the H-List. Thus, for a total of q_h hash queries and q_s sign queries, the number of entries in the H-List (and the number of hash queries made by algorithm \mathcal{B} to challenger \mathcal{C}) is upper-bounded by $n(q_h + q_s)$. Consequently, $q'_h \leq n(q_h + q_s)$.

For each chain sign query by adversary \mathcal{A} , algorithm \mathcal{B} makes at most one sign query to challenger \mathcal{C} . Therefore, $q'_s \leq q_s$.

It only remains to bound the running time t' of \mathcal{B} . This is the running time of \mathcal{A} , plus the time required to add up to n entries in the K-List and up to $n(q_h + q_s)$ entries in the H-List. Each signature query involves up to one multiplication in G_1 . Assuming that the lists are efficiently indexed, the time for searching the H-List and K-List can be ignored.

Adding each entry in the H-List requires 1 exponentiation and up to 1 multiplication in G_1 . Adding each entry in the K-List requires at most 1 exponentiation. Therefore, for a maximum of $n(q_h + q_s)$ entries in the H-List, a maximum of n entries in the K-List, and a maximum of q_s signature queries, we have $t' \leq t + c_{G_1}(n(q_h + 1) + q_s(n + 1))$, where c_{G_1} is the time for 1 exponentiation and 1 multiplication in G_1 . □

4.4 Efficiency

The benchmarks of [15] indicate that each pairing operation using these parameters takes ≈ 8.6 ms and each elliptic curve point exponentiation takes ≈ 1.5 ms. These results were obtained on a desktop PC with an AMD Athlon 2100+ 1.8 GHz, 1 GB RAM and an IBM 7200 RPM, 40 GB, Ultra ATA/100 hard drive [15]. Using these values and neglecting the faster operations, we obtain the following performance estimates of the above protocol (assuming n users in the chain):

1. **Signing:** one exponentiation in G_1 , one multiplication in G_1 , and one computation of \mathcal{H} (total < 2 ms).
2. **Verification:** n pairing computations and multiplications in G_2 , and n computations of \mathcal{H} (giving < 1 second for $n = 100$).

5 Applications of Chain Signatures

Considering that chain signature enable us to correctly validate the path of any received message using very short signatures and provides non-repudiation, we can consider several applications: mobile agent authentication [16, 2], group e-commerce (e-commerce transactions where multiple entities are involved such that direct interaction is not possible between many of them), electronic work-flow enforcement (ensuring the order in which participants should be involved), secure routing, authenticated mail relaying, IP tracing, Grid computing and Mobile IP. Here, we will discuss one such application.

5.1 Secure Routing

We present a novel method for secure and efficient routing. The most common interior and exterior routing protocol is the Border Gateway Protocol (BGP), which is essentially a Distance Vector Routing (DVR) algorithm. In this protocol, routers repeatedly advertise 'better' routes to their immediate neighbors. On receiving an update, a router checks its routing table to decide if this advertised route is better than its existing routes. If so, the router updates its table and advertises the new updated route (from *its own tables*) to all its immediate neighbors except the one from which this update was received. Also,

an update is discarded once it reaches a preset maximum ‘cost’. The important thing to note here is that a routing table is updated every time a better route advertisement is received. Thus, an adversary could send a forged route advertisement and corrupt legitimate routing tables, which would lead to drastic consequences. A key security issue here, therefore, is to ensure the authenticity of the received advertisement before accepting it as valid.

Additionally, every time an advertisement is received, not only should it be ensured that the original initiator of the update was legitimate but also that the path through which this update propagated corresponds to the actual path indicated by the routing tables of the routers that accepted this update.

We propose the use of chain signatures that allow this to be done very efficiently. Denote by $L_i = \langle y_1, y_2, \dots, y_n \rangle$ some ordered sequence of (public keys of) routers that would be affected by a given routing update. Note that there will be many such distinct sequences for the same update and y_1 would be the first name (i.e., key) all these sequences. Let $\text{SIGN}_i, \text{VERIFY}_i$ be the sign and verify functions of user i (having public key y_i) under another existentially unforgeable signature scheme, such as BLS.

Denote by U_i the update message of i that is sent to $i + 1$ describing this particular update. First, the initiator, 1 generates a message $M \in \{0, 1\}^l$ describing this update (i.e., the source and a time-stamp). Denote by $(\sigma_i, (M, L_i))$, the **chain signature** of i on the message M using the sequence L_i and let $S_i = \text{SIGN}_i(M, U_i)$. Each update advertisement U_i sent by i to all its neighbors must include $\langle M, S_i, \sigma_i, L_i \rangle$. A neighboring router $i + 1$ accepts this advertisement as valid if the following checks pass:

1. $\text{VERIFY}_i((M, U_i), S_i) = \text{True}$.
2. σ_i is a valid chain signature on (M, L_i) .
3. The destination and time-stamp defined in M match with the route in question.
4. All routers specified in L_i are legitimate and trusted
5. Routes to each of the routers specified in L_i exist in the current routing table
6. The metrics of the existing routes (almost) match⁸ with the order of hosts specified in L_i .

We now briefly analyze this method:

- (b) *Storage*: To be able to validate the signatures, each host must be able to store/obtain public keys of all routers in question, which may not be feasible. This problem is easily solved using Identity Based Chain Signatures (IBCS) (briefly discussed in the conclusion) and Identity Based Signatures (IBS) where the IP address of a host acts as the public key.
- (a) *Overhead*: The signature includes the tuple $\langle M, S, \sigma_i, L_i \rangle$ in addition to the message U_i . The message U_i is of the form $\langle \text{DestinationAddress}, \text{NextHop}, \text{Metric} \rangle$. If we are using the BLS signature scheme [11] and each host i is represented using 32 bit IP addresses, the signature payload is $32i + 513$ bits at the i^{th} step. Assuming that a typical router advertisement propagates to up to 20 levels, the signature size is between 69 – 144 bytes. Also, verification requires between 20 – 200ms. Note that this overhead is incurred for each unique destination in the router advertisement.

In the above description, we assumed that each advertisement U_i contains only one route and is transmitted instantaneously. In the real world, each advertisement contains multiple routes and is sent periodically. Fortunately, both the chain signature and individual signature schemes used above allow for *signature aggregation* and *aggregate verification* where a large number of (chained or individual) signatures can be verified at once [5].

The use of aggregate signatures in BGP advertisement verification has already been discussed in [5]. We additionally suggest using chain signatures for increased security provided that the extra computation and transmission overhead can be afforded. (Since the actual size of the signature increases as $\mathcal{O}(m)$ where m is the size of the advertised message.)

⁸We say almost rather than an exact match to allow for dynamic route changes while an update is in progress

- (c) *Security*: The use of the time-stamp avoids any replay attacks. The use of M ensures that some sort of correlation is maintained between an advertised route and the actual destination (due to the time stamp, existing routes, path of the route propagation specified by the chain signature, etc). If a corrupted router sends a forged route advertisement, it is traceable.
- (d) *Link State Routing*: This system offers more efficiency/security with Link State Routing (LSR) because here, unlike DVR, the messages (or Link State Advertisements) are passed on unmodified.

6 Summary

In this paper, we introduced the notion of Chain Signatures as an extension of Boneh et al.'s short signatures [11]. Although chain signature arise naturally from the aggregate signatures of [5] due to the inherent properties of bilinear maps, the security requirements of chain signatures is significantly different as demonstrated in Sections 3.2 and 3.4. We note that chain signatures without using bilinear maps were independently proposed in [2, 3] in which the authors used hypothetical primitives called *Strong Associative One-Way Functions* (SAOWFs⁹).

The protocol presented here uses a standard certificate-based PKI. However, it is possible to construct Identity Based Chained Signatures (IBCS) because of the observation that the Identity Based Signature (IBS) schemes of [17, 18] support signature aggregation with the property that once aggregated, individual signatures cannot be extracted.

Considering that chained signatures enable us to correctly validate the path of any received message and provide non-repudiation, we can consider several applications: mobile agent authentication [3, 2], electronic auctions, relaying, token based authentication. As a practical demonstration of applications, we presented a novel method for secure routing.

The main feature of chain signatures that distinguishes them from other multi-user signature schemes is that chain signatures provide *delete-protection* (See Section 2.1). The chain signature scheme presented here, however, does not provide *strong delete-protection*. Signatures that also provide strong delete protection are called *Strong Chain Signatures* (SCS). However, whether practical SCS schemes exist or not, is still an open question at this stage.

References

- [1] Masahiro Mambo, Keisuke Usuda, and Eiji Okamoto. Proxy signatures for delegating signing operation. In *CCS '96: Proceedings of the 3rd ACM conference on Computer and communications security*, pages 48–57, New York, NY, USA, 1996. ACM Press.
- [2] Amitabh Saxena and Ben Soh. Authenticating mobile agent platforms using signature chaining without trusted third parties. In *Proceedings of The 2005 IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE-05)*, pages 282–285, Hong kong, 2005. IEEE computer press.
- [3] Amitabh Saxena and Ben Soh. A novel method for authenticating mobile agents with one-way signature chaining. In *Proceedings of The 7th International Symposium on Autonomous Decentralized Systems (ISADS 05)*, pages 187–193, China, 2005. IEEE Computer Press.
- [4] Silvio Micali and Ronald L. Rivest. Transitive signature schemes. In *CT-RSA*, pages 236–243, 2002.
- [5] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In Eli Biham, editor, *EUROCRYPT*, volume 2656 of *Lecture Notes in Computer Science*, pages 416–432. Springer, 2003.
- [6] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.

⁹SAOWFs exhibit properties similar to multilinear maps.

- [7] H. Shacham. Sequential aggregate signatures from trapdoor homomorphic permutations, 2003.
- [8] Steve Lu, Rafail Ostrovsky, Amit Sahai, Hovav Shacham, and Brent Waters. Sequential aggregate signatures and multisignatures without random oracles. Cryptology ePrint Archive, Report 2006/096, 2006.
- [9] Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In *PKC '03: Proceedings of the 6th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46, London, UK, 2003. Springer-Verlag.
- [10] Mike Burmester, Yvo Desmedt, Hiroshi Doi, Masahiro Mambo, Eiji Okamoto, Mitsuru Tada, and Yuko Yoshifuji. A structured elgamal-type multisignature scheme. In *PKC '00: Proceedings of the Third International Workshop on Practice and Theory in Public Key Cryptography*, pages 466–483, London, UK, 2000. Springer-Verlag.
- [11] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In *ASIACRYPT '01: Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security*, pages 514–532, London, UK, 2001. Springer-Verlag.
- [12] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.
- [13] Paulo S. L. M. Barreto, Hae Yong Kim, Ben Lynn, and Michael Scott. Efficient algorithms for pairing-based cryptosystems. In *CRYPTO '02: Proceedings of the 22nd Annual International Cryptology Conference on Advances in Cryptology*, pages 354–368, London, UK, 2002. Springer-Verlag.
- [14] Alfred J. Menezes, Scott A. Vanstone, and Paul C. Van Oorschot. *Handbook of Applied Cryptography*. CRC Press, Inc., Boca Raton, FL, USA, 1996.
- [15] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. Improved proxy re-encryption schemes with applications to secure distributed storage. Cryptology ePrint Archive, Report 2005/028, 2005.
- [16] Amitabh Saxena and Ben Soh. A mobile agent authentication protocol using signature chaining with bilinear pairings. Cryptology ePrint Archive, Report 2005/272, 2005.
- [17] Jae Choon Cha and Jung Hee Cheon. An identity-based signature from gap Diffie-Hellman groups. In Yvo Desmedt, editor, *Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 18–30. Springer, 2003.
- [18] B. Libert and J. Quisquater. The exact security of an identity based signature and its applications. Technical Report 2004/102, Cryptology ePrint Archive, 2004.