

Aggregate and Verifiably Encrypted Signatures from Bilinear Maps

Dan Boneh
dabo@cs.stanford.edu

Craig Gentry
cgentry@docomolabs-usa.com

Ben Lynn
blynn@cs.stanford.edu

Hovav Shacham
hovav@cs.stanford.edu

Abstract

An aggregate signature scheme is a digital signature that supports aggregation: Given n signatures on n distinct messages from n distinct users, it is possible to aggregate all these signatures into a single short signature. This single signature (and the n original messages) will convince the verifier that the n users did indeed sign the n original messages (i.e., user i signed message M_i for $i = 1, \dots, n$). In this paper we introduce the concept of an aggregate signature scheme, present security models for such signatures, and give several applications for aggregate signatures. We construct an efficient aggregate signature from a recent short signature scheme based on bilinear maps due to Boneh, Lynn, and Shacham. Aggregate signatures are useful for reducing the size of certificate chains (by aggregating all signatures in the chain) and for reducing message size in secure routing protocols such as SBGP. We also show that aggregate signatures give rise to verifiably encrypted signatures. Such signatures enable the verifier to test that a given ciphertext C is the encryption of a signature on a given message M . Verifiably encrypted signatures are used in contract-signing protocols. Finally, we show that similar ideas can be used to extend the short signature scheme to give simple ring signatures.

1 Introduction

Many real-world applications involve signatures on many different messages generated by many different users. For example, in a Public Key Infrastructure (PKI) of depth n , each user is given a chain of n certificates. The chain contains n signatures by n Certificate Authorities (CAs) on n distinct certificates. Similarly, in the Secure BGP protocol (SBGP) [14] each router receives a list of n signatures attesting to a certain path of length n in the network. A router signs its own segment in the path and forwards the resulting list of $n + 1$ signatures to the next router. As a result, the number of signatures in routing messages is linear in the length of the path. Both applications would benefit from a method for compressing the list of signatures on distinct messages issued by distinct parties. Specifically, X.509 certificate chains could be shortened by compressing the n signatures in the chain into a single signature.

An aggregate signature scheme enables us to achieve precisely this type of compression. Suppose each of n users has a public-private key pair (PK_i, SK_i) . User u_i signs message M_i to obtain a signature σ_i . Then there is a public aggregation algorithm that takes as input all of $\sigma_1, \dots, \sigma_n$ and outputs a short compressed signature σ . Anyone can aggregate the n signatures. Moreover, the aggregation can be performed incrementally. That is, signatures σ_1, σ_2 can be aggregated into σ_{12} which can then be further aggregated with σ_3 to obtain σ_{123} . When aggregating signatures in a certificate chain, each CA can incrementally aggregate its own signature into the chain. There is also an aggregate verification algorithm that takes $PK_1, \dots, PK_n, M_1, \dots, M_n$, and σ and decides

whether the aggregate signature is valid. Intuitively, the security requirement is that the aggregate signature σ is declared valid only if the aggregator who created σ was given all of $\sigma_1, \dots, \sigma_n$. Precise security definitions are given in Section 3.2. Thus, an aggregate signature provides non-repudiation at once on many different messages by many users.

We construct an aggregate signature scheme based on a recent short signature due to Boneh, Lynn, and Shacham (BLS) [6]. This signature scheme works in any group where the Decision Diffie-Hellman problem (DDH) is easy, but the Computational Diffie-Hellman problem (CDH) is hard. We refer to such groups as gap groups [6, 22]. Recently there have been a number of constructions using such gap groups [6, 15, 7, 4]. Surprisingly, gap groups are insufficient for constructing efficient aggregate signatures. Instead, our construction uses a pair of groups G_1, G_T for which there is a bilinear map $e : G_1 \times G_1 \rightarrow G_T$ and CDH is hard in G_1 . Joux and Nguyen [13] showed that the map e can be used to solve DDH, and so G_1 is a gap group. However, the extra structure provided by the bilinear map enables us to construct an efficient aggregate signature scheme. We do not know how to build efficient aggregate signatures from general gap groups. Thus, our construction is an example where the bilinear map provides extra functionality beyond a simple algorithm for solving DDH. Bilinear maps were previously used for three-way Diffie-Hellman [13], Identity-Based Encryption (IBE) [5], and Hierarchical IBE [12, 10].

Aggregate signatures are related to multisignatures [17, 21, 20, 4]. In multisignatures, a set of users all sign the *same message* and the result is a single signature. Recently, Micali *et al.* [17] defined a security model for multisignatures and gave some constructions and applications. Multisignatures are insufficient for the applications we have in mind, such as certificate chains and SBGP. For these applications we must be able to aggregate signatures on distinct messages. We note that recently Boldyreva [4] showed that general gap groups are sufficient for constructing multisignatures from BLS signatures. As noted above, to obtain aggregate signatures, one needs the extra structure provided by bilinear maps.

As a further application for aggregate signatures we show in Section 4 that certain aggregate signature schemes give rise to simple verifiably encrypted signatures. These signatures enable user Alice to give Bob a signature on a message M encrypted using a third party's public key and Bob to verify offline that the encrypted signature is valid. Verifiably encrypted signatures are used in optimistic contract signing protocols [1, 2] to enable fair exchange. Previous constructions [1, 23] require zero knowledge proofs and interaction to verify the encrypted signature. The verifiably encrypted signatures in Section 4 can be verified without any interaction.

As a third application of these ideas we construct in Section 5 a simple ring signature [24] using bilinear maps. As above, the construction using a bilinear map is simpler and more efficient than constructions that only make use of gap groups.

2 Signature schemes based on co-gap Diffie-Hellman

We first review a few concepts related to bilinear maps and Gap Diffie-Hellman signatures [6]. Throughout the paper we use the following notation:

1. G_1 and G_2 are two (multiplicative) cyclic groups of prime order p ;
2. g_1 is a generator of G_1 and g_2 is a generator of G_2 ;
3. ψ is a computable isomorphism from G_1 to G_2 , with $\psi(g_1) = g_2$; and
4. e is a computable bilinear map $e : G_1 \times G_2 \rightarrow G_T$ as described below.

The isomorphism ψ is mostly needed for the proofs of security. To keep the discussion general, we simply assume that ψ exists and is efficiently computable. When G_1, G_2 are subgroups of the group

of points of an elliptic curve E/\mathbb{F}_q , the trace map on the curve can be used as this isomorphism (we assume $G_1 \subseteq E/\mathbb{F}_{q^r}$ and $G_2 \subseteq E/\mathbb{F}_q$).

Throughout the paper, we consider bilinear maps $e : G_1 \times G_2 \rightarrow G_T$ where all groups G_1, G_2, G_T are multiplicative and of prime order p . One could set $G_1 = G_2$. However, we allow for the more general case where $G_1 \neq G_2$ so that our constructions can make use of certain families of non-supersingular elliptic curves defined by Miyaji *et al.* [18]. These curves give rise to very short signatures [6]. This will lead in turn to short aggregate signatures, ring signatures, etc. To handle the case $G_1 \neq G_2$ we define the co-CDH and co-DDH problems [6]. When $G_1 = G_2$, these problems reduce to the standard CDH and DDH problems. Hence, for the remainder of the paper, although we handle arbitrary G_1, G_2 , the reader may assume $G_1 = G_2$ for simplicity.

With this setup we obtain natural generalizations of the CDH and DDH problems:

Computational co-Diffie-Hellman. Given $g_1, g_1^a \in G_1$ and $g_2, g_2^b \in G_2$ compute $g_2^{ab} \in G_2$.

Decision co-Diffie-Hellman. Given $g_1, g_1^a \in G_1$ and $g_2, g_2^b, g_2^c \in G_2$ decide whether $ab = c$. When $G_1 = G_2$ and $g_1 = g_2$, these problems reduce to the standard CDH and DDH. Next we define co-GDH gap groups to be group pairs G_1 and G_2 on which co-DDH is easy but co-CDH is hard.

Definition 2.1. The groups G_1 and G_2 are, together, decision groups for co-Diffie-Hellman if the group action on both G_1 and G_2 and the map ψ from G_1 to G_2 can be computed in one time unit, and Decision co-Diffie-Hellman on G_1 and G_2 can be solved in one time unit.

Definition 2.2. The advantage of an algorithm \mathcal{A} in solving the Computational co-Diffie-Hellman problem in groups G_1 and G_2 is

$$\text{Adv co-CDH}_{\mathcal{A}} \stackrel{\text{def}}{=} \Pr \left[\mathcal{A}(g_1, g_1^a, g_2, g_2^b) = g_2^{ab} : a, b \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_p \right].$$

The probability is taken over the choice of a and b and \mathcal{A} 's coin tosses. An algorithm \mathcal{A} (t, ϵ) -breaks Computational co-Diffie-Hellman on G_1 and G_2 if \mathcal{A} runs in time at most t , and $\text{Adv co-CDH}_{\mathcal{A}}$ is at least ϵ . Groups G_1 and G_2 are, together, (t, ϵ) -co-GDH groups if they are decision groups for co-Diffie-Hellman and no algorithm (t, ϵ) -breaks Computational co-Diffie-Hellman on them.

2.1 Bilinear maps

Let G_1 and G_2 be two groups as above, with an additional group G_T such that $|G_1| = |G_2| = |G_T|$. A bilinear map is a map $e : G_1 \times G_2 \rightarrow G_T$ with the following properties:

1. Bilinear: for all $u \in G_1, v \in G_2$ and $a, b \in \mathbb{Z}$, $e(u^a, v^b) = e(u, v)^{ab}$.
2. Non-degenerate: $e(g_1, g_2) \neq 1$.

These properties imply two more: for any $u \in G_1, v_1, v_2 \in G_2$, $e(u, v_1 v_2) = e(u, v_1) \cdot e(u, v_2)$; for any $u, v \in G_1$, $e(u, \psi(v)) = e(v, \psi(u))$.

Definition 2.3. Two groups G_1 and G_2 are, together, bilinear groups if the group action on either can be computed in one time unit, the map ψ from G_1 to G_2 can be computed in one time unit, a bilinear map $e : G_1 \times G_2 \rightarrow G_T$ exists, and e is computable in one time unit.

Definition 2.4. Two groups G_1 and G_2 are, together, (t, ϵ) -bilinear groups for co-Diffie-Hellman if they are bilinear groups and no algorithm (t, ϵ) -breaks Computational co-Diffie-Hellman on them.

Joux and Nguyen [13] showed that an efficiently-computable bilinear map e provides an algorithm for solving the decision co-Diffie-Hellman problem. For a tuple $(g_1, g_1^a, g_2^b, g_2^c)$ we have

$$ab = c \pmod{p} \iff e(g_1, g_2^c) = e(g_1^a, g_2^b).$$

Consequently, if two groups G_1 and G_2 are together, (t, ϵ) -bilinear groups for co-Diffie-Hellman, then they are also (t, ϵ) -co-GDH groups. The converse is probably not true.

2.2 The co-GDH signature scheme

We review the signature scheme of [6], which can be based on any gap group. It comprises three algorithms, *KeyGen*, *Sign*, and *Verify*, and uses a full-domain hash function $h : \{0, 1\}^* \rightarrow G_2$, viewed as a random oracle [3].

Key Generation. Pick random $x \xleftarrow{R} \mathbb{Z}_p$, and compute $v \leftarrow g_1^x$. The public key is $v \in G_1$. The secret key is $x \in \mathbb{Z}_p$.

Signing. Given a secret key x and a message $M \in \{0, 1\}^*$, Compute $h \leftarrow h(M)$, where $h \in G_2$, and $\sigma \leftarrow h^x$. The signature is $\sigma \in G_2$.

Verification. Given a public key v , a message M , and a signature σ , compute $h \leftarrow h(M)$ and verify that (g_1, v, h, σ) is a valid co-Diffie-Hellman tuple.

A co-GDH signature is a single element of G_2 . On certain elliptic curves these signatures are very short: they are half the size of DSA signatures with similar security. Theorem 1 of [6] proves the existential unforgeability of the scheme under a chosen message attack [11] in the random oracle model assuming G_1 and G_2 are co-gap groups for Diffie-Hellman.

3 Aggregate signatures

We define aggregate signatures and describe an aggregate signature scheme based on co-GDH signatures. Unlike the co-GDH scheme, aggregate signatures require the existence of a bilinear map. We define security models and provide proofs of security for aggregate signatures.

Consider a set \mathbb{U} of users. Each user $u \in \mathbb{U}$ has a signing keypair (PK_u, SK_u) . We wish to aggregate the signatures of some subset $U \subseteq \mathbb{U}$. Each user $u \in U$ produces a signature σ_u on a message M_u of her choice. These signatures are then combined into a single aggregate σ by an aggregating party. The aggregating party, who can be different from and untrusted by the users in U , has access to the users' public keys, to the messages, and to the signatures on them, but not to any private keys. The result of this aggregation is an aggregate signature σ whose length is the same as that of any of the individual signatures. This aggregate has the property that a verifier given σ along with the identities of the parties involved and their respective messages is convinced that each user signed his respective message.

3.1 Bilinear aggregate signatures

We describe a bilinear aggregate signature scheme. Individual signatures in the aggregate signature scheme are created and verified precisely as are signatures in the co-GDH scheme (Section 2.2). Aggregate verification makes use of a bilinear map on G_1 and G_2 .

The aggregate signature scheme allows the creation of signatures on arbitrary distinct messages $M_i \in \{0, 1\}^*$. An individual signature σ_i is an element of G_2 . The base groups G_1 and G_2 , their respective generators g_1 and g_2 , the computable isomorphism ψ from G_1 to G_2 , and the bilinear map $e : G_1 \times G_2 \rightarrow G_T$, with target group G_T , are system parameters.

The scheme comprises five algorithms: *KeyGen*, *Sign*, *Verify*, *Aggregate*, and *AggregateVerify*. The first three are as in ordinary signature schemes; the last two provide the aggregation capability. The scheme employs a full-domain hash function $h : \{0, 1\}^* \rightarrow G_2$, viewed as a random oracle.

Key Generation. For a particular user, pick random $x \xleftarrow{R} \mathbb{Z}_p$, and compute $v \leftarrow g_1^x$. The user's public key is $v \in G_1$. The user's secret key is $x \in \mathbb{Z}_p$.

Signing. For a particular user, given the secret key x and a message $M \in \{0, 1\}^*$, compute $h \leftarrow h(M)$, where $h \in G_2$, and $\sigma \leftarrow h^x$. The signature is $\sigma \in G_2$.

Verification. Given user's public key v , a message M , and a signature σ , compute $h \leftarrow h(M)$; accept if $e(g_1, \sigma) = e(v, h)$ holds.

Aggregation. For the aggregating subset of users $U \subseteq \mathbb{U}$, assign to each user an index i , ranging from 1 to $k = |U|$. Each user $u_i \in U$ provides a signature $\sigma_i \in G_2$ on a message $M_i \in \{0, 1\}^*$ of his choice. The messages M_i must all be distinct. Compute $\sigma \leftarrow \prod_{i=1}^k \sigma_i$. The aggregate signature is $\sigma \in G_2$.

Aggregate Verification. We are given an aggregate signature $\sigma \in G_2$ for an aggregating subset of users U , indexed as before, and are given the original messages $M_i \in \{0, 1\}^*$ and public keys $v_i \in G_1$ for all users $u_i \in U$. To verify the aggregate signature σ ,

1. ensure that the messages M_i are all distinct, and reject otherwise; and
2. compute $h_i \leftarrow h(M_i)$ for $1 \leq i \leq k = |U|$, and accept if $e(g_1, \sigma) = \prod_{i=1}^k e(v_i, h_i)$ holds.

A bilinear aggregate signature, like a co-GDH signature, is a single element of G_2 . Note that aggregation can be done incrementally.

The intuition behind bilinear aggregate signatures is as follows. Each user u_i has a secret key $x_i \in \mathbb{Z}_p$ and a public key $v_i = g_1^{x_i}$. User u_i 's signature, if correctly formed, is $\sigma_i = h_i^{x_i}$, where h_i is the hash of the user's chosen message, M_i . The aggregate signature σ is thus $\sigma = \prod_i \sigma_i = \prod_i h_i^{x_i}$. Using the properties of the bilinear map, the left-hand side of the verification equation expands:

$$e(g_1, \sigma) = e(g_1, \prod_i h_i^{x_i}) = \prod_i e(g_1, h_i)^{x_i} = \prod_i e(g_1^{x_i}, h_i) = \prod_i e(v_i, h_i),$$

which is the right-hand side, as required. It remains to prove security of the scheme.

3.2 Aggregate signature security

Informally, the security of aggregate signature schemes is equivalent to the nonexistence of an adversary capable, within the confines of a certain game, of existentially forging an aggregate signature. Existential forgery here means that the adversary attempts to forge an aggregate signature, on messages of his choice, by some set of users.

We formalize this intuition as the aggregate chosen-key security model. In this model, the adversary \mathcal{A} is given a single public key. His goal is the existential forgery of an aggregate signature. We give the adversary power to choose all public keys except the challenge public key. The adversary is also given access to a signing oracle on the challenge key. His advantage, $\text{Adv AggSig}_{\mathcal{A}}$, is defined to be his probability of success in the following game.

Setup. The aggregate forger \mathcal{A} is provided with a public key PK_1 , generated at random.

Queries. Proceeding adaptively, \mathcal{A} requests signatures with PK_1 on messages of his choice.

Response. Finally, \mathcal{A} outputs $k - 1$ additional public keys PK_2, \dots, PK_k . Here k is at most N , a game parameter. These keys, along with the initial key PK_1 , will be included in \mathcal{A} 's forged aggregate. \mathcal{A} also outputs messages M_1, \dots, M_k ; and, finally, an aggregate signature σ by the k users, each on his corresponding message.

The forger wins if the aggregate signature σ is a valid aggregate on messages M_1, \dots, M_k under keys PK_1, \dots, PK_k , and σ is nontrivial, i.e., \mathcal{A} did not request a signature on M_1 under PK_1 . The probability is over the coin tosses of the key-generation algorithm and of \mathcal{A} .

Definition 3.1. An aggregate forger \mathcal{A} $(t, q_H, q_S, N, \epsilon)$ -breaks an N -user aggregate signature scheme in the aggregate chosen-key model if: \mathcal{A} runs in time at most t ; \mathcal{A} makes at most q_H queries

to the hash function and at most q_S queries to the signing oracle; $\text{AdvAggSig}_{\mathcal{A}}$ is at least ϵ ; and the forged aggregate signature is by at most N users. An aggregate signature scheme is $(t, q_H, q_S, N, \epsilon)$ -secure against existential forgery in the aggregate chosen-key model if no forger $(t, q_H, q_S, N, \epsilon)$ -breaks it.

A potential attack on aggregate signatures. The adversary’s ability in the chosen-key model to generate keys suggests the following attack, previously considered in the context of multisignatures [17, 4]. Alice publishes her public key v_A . Bob generates a private key x'_B and a public key $v'_B = g_1^{x'_B}$, but publishes as his public key $v_B = v'_B/v_A$, a value whose discrete log he does not know. Then $h(M)^{x'_B}$ verifies as an aggregate signature on M by both Alice and Bob. Note that in this forgery Alice and Bob both sign the same message M .

One countermeasure is to require the adversary to prove knowledge of the discrete logarithms (to base g_1) of his published public keys. For example, Boldyreva, in her multisignature scheme [4], requires, in effect, that the adversary disclose the corresponding private keys x_2, \dots, x_k . Micali *et al.* [17] discuss a series of more sophisticated approaches based on zero-knowledge proofs, again with the effect that the adversary is constrained in his key selection. These defenses apply equally well to our aggregate signature scheme. For aggregate signatures, though, there is a simpler defense.

A simple defense for aggregate signatures. In the context of aggregate signature we can defend against the attack above by simply requiring that an aggregate signature is valid only if it is an aggregation of signatures on *distinct* messages. This restriction, codified in Step 1 of *Aggregate Verify*, suffices to prove the security of the bilinear aggregate signature scheme in the chosen-key model. There is no need for zero-knowledge proofs or the disclosure of private keys.

The requirement that all messages in an aggregate be distinct is naturally satisfied for the applications to certificate chains and SBGP we have in mind. Even in more general environments it is easy to ensure that all messages are distinct: The signer simply prepends her public key to every message she signs prior to the application of the hash function h . The implicit prefix need not be transmitted with the signature, so signature and message length is unaffected.

The next theorem shows that this simple constraint is sufficient for proving security in the chosen-key model. The proof is given in Appendix A.

Theorem 3.2. *Let G_1 and G_2 together be (t', ϵ') -bilinear groups for co-Diffie-Hellman, of order p , with respective generators g_1 and g_2 , with an isomorphism ψ computable from G_1 to G_2 , and with a bilinear map $e : G_1 \times G_2 \rightarrow G_T$. Then the bilinear aggregate signature scheme on G_1 and G_2 is $(t, q_H, q_S, N, \epsilon)$ -secure against existential forgery in the aggregate chosen-key model, where $t \leq t' - 2c_c(\lg p)(q_H + q_S + N + 2)$, $\epsilon \geq 2e \cdot (q_S + N - 1)\epsilon'$, and c_c is a small constant (less than 2). Here e is the base of natural logarithms.*

Aggregate verification time. Let σ be an aggregate of the n signature $\sigma_1, \dots, \sigma_n$. The time to verify the aggregate signature σ is linear in n . In the special case when all n signatures are issued by the same public key v , aggregate verification time is lower. One need only verify that $e(g_2, \sigma) = e(\prod_{i=1}^n h(M_i), v)$, where M_1, \dots, M_n are the signed messages.

4 Verifiably encrypted signatures

Next, we show an application of aggregate signatures to verifiably encrypted signatures. Verifiably encrypted signatures are used in applications such as online contract signing [1, 2]. Suppose Alice

wants to show Bob that she has signed a message, but does not want Bob to possess her signature of that message. (Alice will give her signature to Bob only when a certain event has occurred, e.g., Bob has given Alice his signature on the same message.) Alice can achieve this by encrypting her signature using the public key of a trusted third party, and sending this to Bob along with a proof that she has given him a valid encryption of her signature. Bob can see Alice has signed the message, but cannot deduce any information about her signature. Later in the protocol, if Alice is unable or unwilling to reveal her signature, Bob can ask the third party to recover Alice’s signature.

We show that a variant of the bilinear aggregate signature scheme allows the creation of very efficient verifiably encrypted signatures.

4.1 Verifiably encrypted signature security

A verifiably encrypted signature scheme comprises seven algorithms. Three, *KeyGen*, *Sign*, and *Verify*, are analogous to those in ordinary signature schemes. The others, *AdjKeyGen*, *VESigCreate*, *VESigVerify*, and *Adjudicate*, provide the verifiably encrypted signature capability. The algorithms are described below. We refer to the trusted third party as the adjudicator.

Key Generation, Signing, Verification. As in standard signature schemes.

Adjudicator Key. Generate a public-private key pair (APK, ASK) for the adjudicator.

VESig Creation. Given a secret key SK , a message M , and an adjudicator’s public key APK , compute (probabilistically) a verifiably encrypted signature ω on M .

VESig Verification. Given a public key PK , a message M , an adjudicator’s public key APK , and a verifiably encrypted signature ω , verify that ω is a valid verifiably encrypted signature on M under key PK .

Adjudication. Given an adjudicator’s keypair (APK, ASK) , a public key PK , and a verifiably encrypted signature ω on some message M , extract and output σ , an ordinary signature on M under PK .

Besides the ordinary notions of signature security in the signature component, we require three security properties of verifiably encrypted signatures: validity, unforgeability, and opacity.

Validity requires that verifiably encrypted signatures verify, and that adjudicated verifiably encrypted signatures verify as ordinary signatures, i.e., that $VESigVerify(M, VESigCreate(M))$ and $Verify(M, Adjudicate(VESigCreate(M)))$ hold for all M and for all properly-generated keypairs and adjudicator keypairs. (The keys provided to the algorithms are here omitted for brevity.)

Unforgeability requires that it be difficult to forge a valid verifiably encrypted signature. The advantage in existentially forging a verifiably encrypted signature of an algorithm \mathcal{F} , given access to a verifiably-encrypted-signature creation oracle S and an adjudication oracle A , along with a hash oracle, is

$$\text{Adv VSigF}_{\mathcal{F}} \stackrel{\text{def}}{=} \Pr \left[\begin{array}{l} (PK, SK) \stackrel{\mathbb{R}}{\leftarrow} \text{KeyGen}, \\ (APK, ASK) \stackrel{\mathbb{R}}{\leftarrow} \text{AdjKeyGen}, \\ (M, \omega) \stackrel{\mathbb{R}}{\leftarrow} \mathcal{F}^{S,A}(PK, APK) \end{array} \mid VESigVerify(PK, APK, M, \omega) = \text{valid} \right].$$

The probability is taken over the coin tosses of the key-generation algorithms, of the oracles, and of the forger. The forger is additionally constrained in that its forgery on M must be nontrivial: It must not previously have queried either oracle at M . Note that an ordinary signing oracle is not provided; it can be simulated by a call to S followed by a call to A .

Definition 4.1. A verifiably encrypted signature forger \mathcal{F} $(t, q_H, q_S, q_A, \epsilon)$ -forges a verifiably encrypted signature if \mathcal{F} runs in time at most t , makes at most q_H queries to the hash function, at most q_S queries to the verifiably encrypted signature creation oracle S , at most q_A queries to the adjudication oracle, and $\text{Adv VSigF}_{\mathcal{F}}$ is at least ϵ . A verifiably encrypted signature scheme is $(t, q_H, q_S, q_A, \epsilon)$ -secure against existential forgery if no forger $(t, q_H, q_S, q_A, \epsilon)$ -breaks it.

Opacity requires that it be difficult, given a verifiably encrypted signature, to extract an ordinary signature on the same message. The advantage in extracting a verifiably encrypted signature of an algorithm \mathcal{E} , given access to a verifiably-encrypted-signature creation oracle S and an adjudication oracle A , along with a hash oracle, is

$$\text{Adv VSigE}_{\mathcal{E}} \stackrel{\text{def}}{=} \Pr \left[\begin{array}{l} (PK, SK) \stackrel{\text{R}}{\leftarrow} \text{KeyGen}, \\ \text{Verify}(PK, M, \sigma) = \text{valid} : (APK, ASK) \stackrel{\text{R}}{\leftarrow} \text{AdjKeyGen}, \\ (M, \sigma) \stackrel{\text{R}}{\leftarrow} \mathcal{E}^{S,A}(PK, APK) \end{array} \right].$$

The probability is taken over the coin tosses of the key-generation algorithms, of the oracles, and of the forger. The extraction must be nontrivial: the adversary must not have queried the adjudication oracle A at M . (It is allowed, however, to query S at M .) Verifiably encrypted signature extraction is thus no more difficult than forgery in the underlying signature scheme.

Definition 4.2. An algorithm \mathcal{E} $(t, q_H, q_S, q_A, \epsilon)$ -extracts a verifiably encrypted signature if \mathcal{E} runs in time at most t , makes at most q_H queries to the hash function, at most q_S queries to the verifiably-encrypted-signature creation oracle S , at most q_A queries to the adjudication oracle, and $\text{Adv VSigE}_{\mathcal{E}}$ is at least ϵ . A verifiably encrypted signature scheme is $(t, q_H, q_S, q_A, \epsilon)$ -secure against extraction if no algorithm $(t, q_H, q_S, q_A, \epsilon)$ -extracts it.

4.2 Aggregate extraction

Our verifiably encrypted signature scheme depends on the assumption that given an aggregate signature of k signatures it is difficult to extract the individual signatures.

Consider the bilinear aggregate signature scheme on G_1 and G_2 . We posit that it is difficult to recover the individual signatures σ_i given their aggregate σ , the public keys, and the message hashes. In fact, we posit that it is difficult to recover an aggregate σ' of any proper subset of the signatures. This we term the k -element aggregate extraction problem.

We formalize this assumption as follows. Let G_1 and G_2 be, together, bilinear groups for co-Diffie-Hellman of order p , with respective generators g_1 and g_2 , a computable isomorphism $\psi : G_1 \rightarrow G_2$ such that $g_2 = \psi(g_1)$, and a computable bilinear map $e : G_1 \times G_2 \rightarrow G_T$.

Consider a k -user aggregate in this setting. Each user has a private key $x_i \in \mathbb{Z}_p$ and a public key $v_i = g_1^{x_i} \in G_1$. Each user selects a distinct message $M_i \in \{0, 1\}^*$ whose hash is $h_i \in G_2$ and creates a signature $\sigma_i = h_i^{x_i} \in G_2$. Finally, the signatures are aggregated, yielding $\sigma = \prod_i \sigma_i \in G_2$.

Let I be the set $\{1, \dots, k\}$. Each public key v_i can be expressed as $g_1^{x_i}$, each hash h_i as $g_2^{y_i}$, each signature σ_i as $g_2^{x_i y_i}$, and the aggregate signature σ as g_2^z , where $z = \sum_{i \in I} x_i y_i$. The advantage of an algorithm \mathcal{E} in extracting a subaggregate from a k -element aggregate is

$$\text{Adv } k\text{-Extr}_{\mathcal{E}} \stackrel{\text{def}}{=} \Pr \left[\left((I' \subsetneq I) \wedge (\sigma' = g_2^{\sum_{i \in I'} x_i y_i}) : \begin{array}{l} x_1, \dots, x_k, y_1, \dots, y_k \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_p, \sigma \leftarrow g_2^{\sum_{i \in I} x_i y_i}, \\ (\sigma', I') \stackrel{\text{R}}{\leftarrow} \mathcal{E}(g_1^{x_1}, \dots, g_1^{x_k}, g_2^{y_1}, \dots, g_2^{y_k}, \sigma) \end{array} \right) \right].$$

The probability is taken over the choices of all x_i and y_i , and the coin tosses of \mathcal{E} .

Definition 4.3. An algorithm \mathcal{E} (t, k, ϵ) -extracts a subaggregate from an k -element bilinear aggregate signature if \mathcal{E} runs in time at most t and $\text{Adv } k\text{-Extr}_{\mathcal{E}}$ is at least ϵ . An instantiation of the bilinear aggregate signature scheme is (t, k, ϵ) -secure against aggregate extraction if no algorithm (t, k, ϵ) -extracts it.

We will be particularly concerned with the case $k = 2$. In this case, the aggregate extraction problem reduces to this one: given $g_1^a, g_1^b, g_2^u, g_2^v$, and g_2^{au+bv} , calculate g_2^{au} . (If the extractor outputs g_2^{bv} instead, we may recover g_2^{au} as g_2^{au+bv}/g_2^{bv} .)

4.3 Verifiably encrypted signatures via aggregation

We motivate our construction for verifiably encrypted signatures by considering aggregate signatures as a launching point. An aggregate signature scheme can give rise to a verifiably encrypted signature scheme if it is difficult to extract individual signatures from an aggregate, but easy to forge existentially under the adjudicator's key. Consider the following:

1. Alice wishes to create a verifiably encrypted signature, which Bob will verify; Carol is the adjudicator. Alice and Carol's keys are both generated under the underlying signature scheme's key-generation algorithm.
2. Alice creates a signature σ on M under her public key. She forges a signature σ' on some random message M' under Carol's public key. She then combines σ and σ' , obtaining an aggregate ω . The verifiably encrypted signature is the pair (ω, M') .
3. Bob validates Alice's verifiably encrypted signature (ω, M') on M by checking that ω is a valid aggregate signature by Alice on M and by Carol on M' .
4. Carol adjudicates, given a verifiably encrypted signature (ω, M') on M by Alice, by computing a signature σ' on M' under her key, and removing σ' from the aggregate; what remains is Alice's ordinary signature σ .

In the bilinear aggregate signature scheme, it is difficult to extract individual signatures, under the aggregate extraction assumption. Moreover, existential forgery is easy when the random oracle hash function is set aside: Given a public key $v \in G_1$ and $r \in \mathbb{Z}_p$, $\psi(v)^r$ is a valid signature on a message whose hash is $\psi(g_1)^r = g_2^r$. Below, we formalize and prove the security of the verifiably encrypted signature scheme created in this way.

4.4 The bilinear verifiably encrypted signature scheme

The bilinear verifiably encrypted signature scheme is built on the bilinear aggregate signature scheme of the previous section. It shares the key-generation algorithm with the underlying aggregate scheme. Moreover, the adjudicator's public and private information is simply an aggregate-signature keypair. The scheme comprises the seven algorithms described below:

Key generation. *KeyGen* and *AdjKeyGen* are the same as *KeyGen* in the co-GDH scheme.

Signing, Verification. *Sign* and *Verify* are the same as in the co-GDH scheme.

VESig Creation. Given a secret key $x \in \mathbb{Z}_p$ a message $M \in \{0, 1\}^*$, and an adjudicator's public key $v' \in G_1$, Compute $h \leftarrow h(M)$, where $h \in G_2$, and $\sigma \leftarrow h^x$. Select r at random from \mathbb{Z}_p and set $\mu \leftarrow \psi(g_1)^r$ and $\sigma' \leftarrow \psi(v')^r$. Aggregate σ and σ' as $\omega \leftarrow \sigma\sigma' \in G_2$. The verifiably encrypted signature is the pair (ω, μ) .

VESig Verification. Given a public key v , a message M , an adjudicator's public key v' , and a verifiably encrypted signature (ω, μ) , set $h \leftarrow h(M)$; accept if $e(g_1, \omega) = e(v, h) \cdot e(v', \mu)$ holds.

Adjudication. Given an adjudicator’s public key v' and corresponding private key $x' \in \mathbb{Z}_p$, a public key v , and a verifiably encrypted signature (ω, μ) on some message M , ensure that the verifiably encrypted signature is valid; then output $\sigma = \omega/\mu^{x'}$.

(If the adjudicator does not first validate a purported verifiably encrypted signature, a malicious user can trick him into signing arbitrary messages under his adjudication key.)

It is easy to see that validity holds. A verifiably encrypted signature correctly validates under $VESigVerify$, which is simply the aggregate signature verification algorithm. Moreover, for any valid verifiably encrypted signature, $e(g_1, \omega/\mu^{x'}) = e(g_1, \omega) \cdot e(g_1, \mu)^{-x'} = e(v, h) \cdot e(v', \mu) \cdot e(v', \mu)^{-1} = e(v, h)$, so the output of $Adjudicate$ is a valid signature on message M under the key v .

The next two theorems prove the unforgeability and opacity of the scheme. The proofs are given in Appendix A.

Theorem 4.4. *Let G_1 and G_2 be cyclic groups of prime order p , with respective generators g_1 and g_2 , with a computable bilinear map $e : G_1 \times G_2 \rightarrow G_T$. Suppose that the co-GDH signature scheme is $(t', q'_H, q'_S, \epsilon')$ -secure against existential forgery on G_1 and G_2 . Then the bilinear verifiably encrypted signature scheme is $(t, q_H, q_S, q_A, \epsilon)$ -secure against existential forgery on G_1 and G_2 , where $t \leq t' - 4c_{\mathcal{F}}(\lg p)(q'_S + q_A + 1)$, $q_H \leq q'_H$, $q_S \leq q'_S$, q_A is at most q'_S , and $c_{\mathcal{F}}$ is a small constant (at most 2).*

Theorem 4.5. *Let G_1 and G_2 be cyclic groups of prime order p , with respective generators g_1 and g_2 , with a computable bilinear map $e : G_1 \times G_2 \rightarrow G_T$. Suppose that the bilinear aggregate signature scheme on G_1 and G_2 is $(t', 2, \epsilon')$ -secure against aggregate extraction. Then the bilinear verifiably encrypted signature scheme is $(t, q_H, q_S, q_A, \epsilon)$ -secure against extraction on G_1 and G_2 , where $t \leq t' - 2c_{\mathcal{A}}(\lg p)(q_H + 3q_S + q_A + 2)$, $\epsilon \geq 2e \cdot q_A \epsilon'$, and $c_{\mathcal{A}}$ is a small constant (at most 2). Here e is the base of natural logarithms.*

4.5 Observations on verifiably encrypted signatures

We note some extensions of the verifiably encrypted signature scheme discussed above. Some of these rely for security on the k -element aggregate extraction assumption with $k > 2$.

- A message-signature pair in the co-GDH signature scheme is of the same form as an identity-private-key pair in the Boneh-Franklin Identity-Based Encryption Scheme [5]. Thus the verifiably encrypted signature scheme can be modified to yield a verifiably encrypted encryption scheme for IBE private keys. Verifiably encrypted private keys have many applications [23].
- Anyone can convert an ordinary unencrypted signature to a verifiably encrypted signature. The same applies to unencrypted aggregate signatures.
- An adjudicator’s private key can be shared amongst n parties using k -of- n threshold cryptography [9, 8], so that k parties are needed to adjudicate a verifiably encrypted signature.

5 Ring signatures

Rivest, Shamir and Tauman define ring signature schemes and construct some using RSA and Rabin cryptosystems [24]. Naor defines the closely-related notion of deniable ring authentication and proposes such a scheme that relies only on the existence of a strong encryption function [19]. We shall see that co-GDH signatures give natural ring signatures.

5.1 Ring signatures

Consider a set U of users. Each user $u \in U$ has a signing keypair (PK_u, SK_u) . A ring signature on U is a signature that is constructed using all the public keys of the users in U , and a single private key of any user in U . A ring signature has the property that a verifier is convinced that the signature was produced using one of the private keys of U , but is not able to determine which one. This property is called *signer-ambiguity* [24]. Applications for ring signatures include authenticated (yet repudiable) communication and leaking secrets [24].

Zhang and Kim [25] devised a bilinear ring signature in an identity-based setting. Our scheme differs from theirs, as our goal is to extend co-GDH signatures to obtain efficient ring signatures; the system parameters and key generation algorithm in our system are identical to those of the co-GDH scheme.

5.2 Bilinear ring signatures

The ring signature scheme comprises three algorithms: *KeyGen*, *RingSign*, and *RingVerify*. Recall g_1, g_2 are generators of groups G_1, G_2 respectively, and $e : G_1 \times G_2 \rightarrow G_T$ is a bilinear map, and a computable isomorphism $\psi : G_1 \rightarrow G_2$ exists, with $\psi(g_1) = g_2$. Again we use a full-domain hash function $h : \{0, 1\}^* \rightarrow G_2$. The security analysis views h as a random oracle.

Key Generation. For a particular user, pick random $x \xleftarrow{R} \mathbb{Z}_p$, and compute $v \leftarrow g_1^x$. The user's public key is $v \in G_1$. The user's secret key is $x \in \mathbb{Z}_p$.

Ring Signing. Given public keys $v_1, \dots, v_n \in G_1$, a message $M \in \{0, 1\}^*$, and a private key x corresponding to one of the public keys v_s for some s , choose random $a_i \xleftarrow{R} \mathbb{Z}_p$ for all $i \neq s$. Compute $h \leftarrow h(M) \in G_2$ and set

$$\sigma_s \leftarrow \left(h / \psi \left(\prod_{i \neq s} v_i^{a_i} \right) \right)^{1/x}.$$

For all $i \neq s$ let $\sigma_i \leftarrow g_2^{a_i}$. Output the ring signature $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle \in G_2^n$.

Ring Verification. Given public keys $v_1, \dots, v_n \in G_1$, a message $M \in \{0, 1\}^*$, and a ring signature σ , compute $h \leftarrow h(M)$ and verify that $e(g, h) = \prod_{i=1}^n e(v_i, \sigma_i)$.

Using the bilinearity and nondegeneracy of the pairing e , it is easy to show that a signature produced by the *RingSign* algorithm will verify under the *RingVerify* algorithm.

5.3 Security

There are two aspects a security analysis for ring signatures must consider. Firstly, signer ambiguity must be ensured. We first show that the identity of the signer is unconditionally protected.

Theorem 5.1. *For any algorithm \mathcal{A} , any set of users U , and a random $u \in U$, the probability $\Pr[\mathcal{A}(\sigma) = u]$ is at most $1/|U|$, where σ is any ring signature on U generated with private key SK_u .*

Proof. The theorem follows from a simple probability argument: for any $h \in G_2$, and any $1 \leq s \leq n$, the distribution $\{g_2^{a_1}, \dots, g_2^{a_n} : a_i \xleftarrow{R} \mathbb{F}_q \text{ for } i \neq s, a_s \text{ chosen such that } \prod_{i=1}^n g_2^{a_i} = h\}$ is identical to the distribution $\{g_2^{a_1}, \dots, g_2^{a_n} : \prod_{i=1}^n g_2^{a_i} = h\}$, since the value of any one of the a_i 's is uniquely determined by the values of the other a_i 's. \square

Secondly, we need to examine the scheme’s resistance to forgery. We adopt the security model of Rivest, Shamir and Tauman [24]. Consider the following game played between an adversary and a challenger. The adversary is given the public keys v_1, \dots, v_n of a set of users U , and is given oracle access to h and a ring-signing oracle. The adversary may work adaptively. The goal of the adversary is to output a valid ring signature on U of a message M subject to the condition that M has never been presented to the ring-signing oracle. An adversary A ’s advantage Adv RingSig_A in existentially forging a bilinear ring signature is the probability, taken over the coin tosses of the key-generation algorithm and of the forger, that A succeeds in creating a valid ring signature in the above game. The proof of the following security theorem is given in Appendix A.

Theorem 5.2. *Suppose \mathcal{F} is a (t', ϵ') -algorithm that can produce a forgery of a ring signature on a set of users of size n . Then there exists an (t, ϵ) -algorithm that can solve the co-CDH problem where $t \leq 2t' + c(\lg p)nq_H$ for some constant c and $\epsilon \geq (\epsilon'/e(1 + q_S))^2$, where \mathcal{F} issues at most q_S ring-signature queries and at most q_H hash queries.*

5.4 Observations on ring signatures

Any ring signature scheme restricts to an ordinary signature scheme when $n = 1$. Our scheme restricts to a short signature scheme similar to the co-GDH scheme [6]. In this modified co-GDH scheme, σ equals $h^{1/x}$ rather than h^x , and one verifies that $e(g_1, h) = e(v, \sigma)$ rather than that $e(g_1, \sigma) = e(v, h)$.

However, bilinear ring signatures have interesting properties that do not appear to be shared by ring signatures in general. For any set of users U with $u \in U$, *anyone* can convert a modified co-GDH signature by u into a ring signature by U . Specifically, to convert a modified co-GDH signature σ_1 on M for public key v_1 into a ring signature $\sigma = \langle \sigma'_1, \dots, \sigma'_n \rangle$ on M for public keys v_1, \dots, v_n , we choose $r_i \xleftarrow{R} \mathbb{Z}_p$ for $2 \leq i \leq n$, and set $\sigma'_1 \leftarrow \sigma_1 \prod_{i=2}^n \psi(v_i^{r_i})$ and $\sigma'_i \leftarrow \psi(v_1^{-r_i})$ for $2 \leq i \leq n$. More generally, anyone can further anonymize a ring signature by adding users to U .

6 Conclusions

We introduced the concept of aggregate signatures and constructed an efficient aggregate signature scheme based on bilinear maps. Key generation, aggregation, and verification require no interaction. We proved security of the system in a model that gives the adversary his choice of public keys and messages to forge. For security, we introduced the additional constraint that an aggregate signature is valid only if it is an aggregation of signatures on *distinct* messages. This constraint is satisfied naturally for the applications we have in mind.

We gave several applications for aggregate signatures. For example, they can be used to reduce the size of certificate chains and reduce communication bandwidth in protocols such as SBGP. We also showed that our specific aggregate signature scheme gives verifiably encrypted signatures and verifiably encrypted private keys.

Previous signature constructions using bilinear maps [6, 15, 7, 4] only required a gap Diffie-Hellman group (i.e., DDH easy, but CDH hard). The signature constructions in this paper require the extra structure provided by the bilinear map. These constructions are an example where a bilinear map provides more power than a generic gap Diffie-Hellman group.

Acknowledgments

The authors thank Leonid Reyzin for helpful discussions about this work.

References

- [1] N. Asokan, V. Shoup, and M. Waidner. Optimistic Fair Exchange of Digital Signatures. *IEEE J. Selected Areas in Comm.*, 18(4):593–610, April 2000.
- [2] F. Bao, R. Deng, and W. Mao. Efficient and Practical Fair Exchange Protocols with Offline TTP. In *Proceedings of IEEE Symposium on security and privacy*, pages 77–85, 1998.
- [3] M. Bellare and P. Rogaway. The Exact Security of Digital Signatures: How to Sign with RSA and Rabin. In U. Maurer, editor, *Proceedings of Eurocrypt '96*, volume 1070 of *LNCS*, pages 399–416. Springer-Verlag, 1996.
- [4] A. Boldyreva. Efficient Threshold Signature, Multisignature and Blind Signature Schemes Based on the Gap-Diffie-Hellman-Group Signature Scheme. Cryptology ePrint Archive, Report 2002/118, Aug 2002. <http://eprint.iacr.org/>.
- [5] D. Boneh and M. Franklin. Identity-Based Encryption from the Weil Pairing. In J. Kilian, editor, *Proceedings of Crypto 2001*, volume 2139 of *LNCS*, pages 213–29. Springer-Verlag, 2001.
- [6] D. Boneh, B. Lynn, and H. Shacham. Short Signatures from the Weil Pairing. In *Proceedings of Asiacrypt 2001*, volume 2248 of *LNCS*, pages 514–32. Springer-Verlag, 2001.
- [7] Y. Dodis. Efficient Construction of (Distributed) Verifiable Random Functions. Cryptology ePrint Archive, Report 2002/133, Aug 2002. <http://eprint.iacr.org/2002/133/>.
- [8] P. Gemmel. An Introduction to Threshold Cryptography. *RSA CryptoBytes*, 2(3):7–12, 1997.
- [9] R. Gennaro, T. Rabin, S. Jarecki, and H. Krawczyk. Robust and Efficient Sharing of RSA Functions. *J. of Cryptology*, 13(2):273–300, 2000.
- [10] C. Gentry and A. Silverberg. Hierarchical ID-Based Cryptography. In Y. Zheng, editor, *Proceedings of Asiacrypt 2002*, LNCS. Springer-Verlag, 2002. To appear.
- [11] S. Goldwasser, S. Micali, and R. Rivest. A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM J. on Computing*, 17(2):281–308, 1988.
- [12] J. Horwitz and B. Lynn. Toward Hierarchical Identity-Based Encryption. In L. Knudsen, editor, *Proc. of Eurocrypt 2002*, volume 2332 of *LNCS*, pages 466–81. Springer-Verlag, 2002.
- [13] A. Joux. A One Round Protocol for Tripartite Diffie-Hellman. In W. Bosma, editor, *Proceedings of ANTS IV*, volume 1838 of *LNCS*, pages 385–94. Springer-Verlag, 2000.
- [14] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE J. Selected Areas in Comm.*, 18(4):582–92, April 2000.
- [15] A. Lysyanskaya. Unique Signatures and Verifiable Random Functions from the DH-DDH Separation. In M. Yung, editor, *Proceedings of Crypto 2002*, volume 2442 of *LNCS*, pages 597–612. Springer-Verlag, 2002.
- [16] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1997.
- [17] S. Micali, K. Ohta, and L. Reyzin. Accountable-Subgroup Multisignatures: Extended Abstract. In *Proceedings of CCS 2001*, pages 245–54. ACM Press, 2001.
- [18] A. Miyaji, M. Nakabayashi, and S. Takano. New Explicit Conditions of Elliptic Curve Traces for FR-Reduction. *IEICE Trans. Fundamentals*, E84-A(5):1234–43, May 2001.
- [19] M. Naor. Deniable Ring Authentication. In M. Yung, editor, *Proceedings of Crypto 2002*, volume 2442 of *LNCS*, pages 481–98. Springer-Verlag, 2002.

- [20] K. Ohta and T. Okamoto. Multisignature Schemes Secure Against Active Insider Attacks. *IEICE Trans. Fundamentals*, E82-A(1):21–31, 1999.
- [21] T. Okamoto. A Digital Multisignature Scheme Using Bijective Public-Key Cryptosystems. *ACM Trans. Computer Systems*, 6(4):432–441, 1998.
- [22] T. Okamoto and D. Pointcheval. The Gap Problems: A New Class of Problems for the Security of Cryptographic Primitives. In K. Kim, editor, *Proceedings of PKC 2001*, volume 1992 of *LNCS*, pages 104–118. Springer-Verlag, 2001.
- [23] G. Poupard and J. Stern. Fair Encryption of RSA Keys. In B. Preneel, editor, *Proceedings of Eurocrypt 2000*, volume 1807 of *LNCS*, pages 172–89. Springer-Verlag, 2000.
- [24] R. Rivest, A. Shamir, and Y. Tauman. How to Leak a Secret. In C. Boyd, editor, *Proceedings of Asiacrypt 2001*, volume 2248 of *LNCS*, pages 552–65. Springer-Verlag, 2001.
- [25] F. Zhang and K. Kim. ID-Based Blind Signature and Ring Signature from Pairings. In Y. Zheng, editor, *Proceedings of Asiacrypt 2002*, LNCS. Springer-Verlag, 2002. To appear.

A Proofs

Proof of Theorem 3.2. Given an aggregate forger \mathcal{A} for the bilinear groups for co-Diffie-Hellman G_1 and G_2 , we build an algorithm \mathcal{C} that uses \mathcal{A} to break co-CDH on G_1 and G_2 . \mathcal{C} is given a challenge (g_1, g_1^a, g_2^b) . It uses this challenge to construct a public key v_1 that it provides to \mathcal{A} . It then allows \mathcal{A} to run. At times, \mathcal{A} makes queries to two oracles, one for message hashes and one for message signatures by the private key corresponding to v_1 . These oracles are puppets of \mathcal{C} , which it manipulates in constructive ways. Finally, if all goes well, the aggregate forgery which \mathcal{A} outputs is transformed by \mathcal{C} into an answer to the co-CDH challenge.

We assume that \mathcal{A} is well-behaved in the sense that it always requests the hash of a message M before it requests a signature for M , and that it always requests a hash of each message M_i that it outputs in its aggregate forgery. It is trivial to modify any forger algorithm \mathcal{A} to have this property.

\mathcal{C} needs to engage in a certain amount of bookkeeping. In particular, it must maintain a list of the messages on which \mathcal{A} requests hashes or signatures. Each message M , as it arrives from \mathcal{A} , is assigned an index i ; i is obviously bounded above by q_H . The message is stored in $M^{(i)}$, its hash in $h^{(i)}$, and its signature (if available) in $\sigma^{(i)}$. \mathcal{C} also maintains a bit array $s^{(i)}$, to be described later.

The behavior of \mathcal{C} will depend on a probability constant ζ , which will be optimized later to yield the best possible reduction. We define B_ζ to be the probability distribution over $\{0, 1\}$ where 1 is drawn with probability ζ , and 0 with probability $1 - \zeta$.

\mathcal{C} is given a challenge (g_1, g_1^a, g_2^b) . In setup, it constructs $v_1 \leftarrow g_1^a$. Then, for each i , $1 \leq i \leq q_H$, \mathcal{C} does the following. It picks a random bit $s^{(i)} \stackrel{R}{\leftarrow} B_\zeta$, and a random number $r^{(i)} \stackrel{R}{\leftarrow} \mathbb{Z}_p$. If $s^{(i)}$ equals 1, \mathcal{C} sets $h^{(i)} \leftarrow g_2^b \cdot \psi(g_1)^{r^{(i)}}$, and $\sigma^{(i)} \leftarrow \star$, a placeholder value; if $s^{(i)}$ equals 0, \mathcal{C} sets $h^{(i)} \leftarrow \psi(g_1)^{r^{(i)}}$, and $\sigma^{(i)} \leftarrow \psi(g_1^a)^{r^{(i)}}$. Note that, when $s^{(i)}$ is 0, $(g_1, v_1, h^{(i)}, \sigma^{(i)})$ is a valid co-Diffie-Hellman tuple, so $\sigma^{(i)}$ is a signature on any message whose hash is $h^{(i)}$.

Now \mathcal{C} runs \mathcal{A} with public key v_1 as input. When \mathcal{A} requests a hash on a message $M^{(i)}$, \mathcal{C} responds with $h^{(i)}$. When \mathcal{A} requests a signature on a message $M^{(i)}$, \mathcal{C} declares failure and halts if $s^{(i)} = 1$, and otherwise responds with $\sigma^{(i)}$.

Finally, \mathcal{A} halts. It either concedes failure or returns a value k (where $k \leq N$), $k - 1$ public keys v_2, \dots, v_k , k messages M_1, \dots, M_k , and a forged aggregate signature σ . The messages M_i must all be distinct, and \mathcal{A} must not have requested a signature on M_1 .

Since \mathcal{A} is well-behaved, it must have asked for the hash of each of the messages M_i in the aggregate. Let μ be the mapping from aggregate-message index to \mathcal{C} 's internal-bookkeeping message

index. Then $M_i = M^{(\mu(i))}$ for all i , $1 \leq i \leq k$. For each i , $1 \leq i \leq k$, \mathcal{C} sets $s_i \leftarrow s^{(\mu(i))}$, $r_i \leftarrow r^{(\mu(i))}$ and $h_i \leftarrow h^{(\mu(i))}$.

\mathcal{C} now proceeds only if $s_1 = 1$ and, for $2 \leq i \leq k$, $s_i = 0$; otherwise it declares failure and halts. Since $s_1 = 1$, it follows that $h_1 = g_2^b \cdot \psi(g_1)^{r_1}$; for $i > 1$, since $s_i = 0$, that $h_i = \psi(g_1)^{r_i}$. The aggregate signature σ must satisfy the aggregate verification equation, $e(g_1, \sigma) = \prod_{i=1}^k e(v_i, h_i)$. For each $i > 1$, \mathcal{C} sets $\sigma_i \leftarrow \psi(v_i)^{r_i}$. Then, for $i > 1$,

$$e(g_1, \sigma_i) = e(g_1, \psi(v_i)^{r_i}) = e(g_1, \psi(v_i))^{r_i} = e(v_i, \psi(g_1))^{r_i} = e(v_i, \psi(g_1)^{r_i}) = e(v_i, h_i),$$

So σ_i is a valid signature on M_i (whose hash is h_i) by the key whose public component is v_i . Now \mathcal{C} constructs a value σ_1 : $\sigma_1 \leftarrow \sigma \cdot (\prod_{i=2}^k \sigma_i)^{-1}$. Then

$$e(g_1, \sigma_1) = e(g_1, \sigma) \cdot \prod_{i=2}^k e(g_1, \sigma_i)^{-1} = e(g_1, \sigma) \cdot \prod_{i=2}^k e(v_i, h_i)^{-1} = \prod_{i=1}^k e(v_i, h_i) \cdot \prod_{i=2}^k e(v_i, h_i)^{-1} = e(v_1, h_1).$$

Thus σ_1 is a valid co-GDH signature by key v_1 on a message whose hash is h_1 . Equivalently, $(g_1, v_1, h_1, \sigma_1) = (g_1, g_1^a, g_2^b \cdot \psi(g_1)^{r_1}, \sigma_1)$ is a valid co-Diffie-Hellman tuple; but then so is $(g_1, g_1^a, g_2^b \cdot \psi(g_1)^{r_1} / \psi(g_1)^{r_1}, \sigma_1 / \psi(g_1^a)^{r_1}) = (g_1, g_1^a, g_2^b, \sigma_1 / \psi(g_1^a)^{r_1})$; therefore $\sigma_1 / \psi(g_1^a)^{r_1}$ is the solution to the posed computational co-Diffie-Hellman challenge; \mathcal{C} calculates this value, outputs it, and halts.

It remains only to analyze the success probability and running time of \mathcal{F} , and to optimize the success probability by choosing a suitable value for ζ .

\mathcal{C} succeeds whenever \mathcal{A} does, so long as (1) \mathcal{A} never requests a signature on a message $M^{(i)}$ for which $s^{(i)} = 1$, and (2) \mathcal{A} aggregates messages M_i such that $s^{(\mu(1))} = 1$ but $s^{(\mu(i))} = 0$ for $i > 1$.

Multiplication by a group element (such as g_2^b) induces a permutation, so each $h^{(i)}$ is uniformly distributed regardless of the values $s^{(i)}$. Thus each signature request on some message $M^{(i)}$ is an independent trial that causes \mathcal{C} to halt with probability ζ . Since \mathcal{A} makes at most q_S signature requests, the probability of \mathcal{C} 's successfully answering all signing queries is at least $(1 - \zeta)^{q_S}$.

Disqualification at the signature-query phase is independent of \mathcal{A} 's eventual success. For comparison, consider a game in which \mathcal{C} computes each $h^{(i)}$ as $\psi(g_1)^{r^{(i)}}$ regardless of $s^{(i)}$, answers any signature queries as asked, allows \mathcal{A} to complete its run, and then fails if it ever asked for a signature on a message $M^{(i)}$ with $s^{(i)} = 1$. In this game, the independence of disqualifying trials and eventual success is obvious; the immediate halt in unlucky runs is a shortcut and does not affect the outcome distribution.

\mathcal{A} cannot previously have asked for a signature on M_1 , so it can have no information about the value of s_1 ; in the forged aggregate, $s_1 = 1$ occurs with probability ζ . For each $i > 1$, \mathcal{A} either asked for a signature on M_i , in which case $s_i = 0$ with probability 1, or it didn't, and $s_i = 0$ with probability $1 - \zeta$. Regardless, the probability that $s_i = 0$ for all i , $2 \leq i \leq k$, is at least $(1 - \zeta)^{k-1} \geq (1 - \zeta)^{N-1}$.

Thus, if \mathcal{A} succeeds with probability ϵ , \mathcal{C} succeeds with probability at least $\zeta \cdot (1 - \zeta)^{q_S + N - 1} \cdot \epsilon$. This probability is maximized at $\zeta = 1/(q_S + N)$, where it has the value

$$\frac{1}{q_S + N} \cdot \left(1 - \frac{1}{q_S + N}\right)^{q_S + N - 1} \cdot \epsilon = \frac{1}{q_S + N - 1} \cdot \left(1 - \frac{1}{q_S + N}\right)^{q_S + N} \cdot \epsilon.$$

Now \mathcal{C} 's success probability ϵ' is at least as great as this.

\mathcal{C} 's running time includes the running time of \mathcal{A} . The overhead imposed by \mathcal{C} is dominated by the need to evaluate group exponentiation for each signature and hash request from \mathcal{A} , and the two inversions and k exponentiations required to generate the co-CDH solution from \mathcal{A} 's forged

aggregate. each component of \mathcal{A} 's aggregate. Any one such operation may be completed by using at most $2\lg p$ group actions, and thus at most $2\lg p$ time units, on G_2 (see [16]). \mathcal{C} may need to perform as many as $q_H + q_S + N + 2$ such operations, so its overall running time is $t' \leq t + 2c_C(\lg p)(q_H + q_S + N + 2)$, Where c_C is a small constant that accounts for the remainder of \mathcal{C} 's administrative overhead; in practice, c_C should be at most 2.

To summarize: if there exists a forger algorithm \mathcal{A} that $(t, q_H, q_S, N, \epsilon)$ -breaks the bilinear aggregate signature scheme on G_1 and G_2 , then there exists an algorithm \mathcal{C} that (t', ϵ') -breaks co-CDH on G_1 and G_2 , where

$$t' \leq t + 2c_C(\lg p)(q_H + q_S + N + 2) \quad \text{and} \quad \epsilon' \geq \frac{1}{q_S + N - 1} \cdot \left(1 - \frac{1}{q_S + N}\right)^{q_S + N} \cdot \epsilon.$$

Conversely, if G_1 and G_2 are, together, (t', ϵ') -co-GDH groups, then there can exist no algorithm \mathcal{A} that (t, q_H, q_S, ϵ) -breaks the co-GDH Signature Scheme on G_1 and G_2 , where

$$t \leq t' - 2c_C(\lg p)(q_H + q_S + N + 2) \quad \text{and} \quad \epsilon \geq (q_S + N - 1)\epsilon' / \left(1 - \frac{1}{q_S + N}\right)^{q_S + N}.$$

For all positive $q_S + N$, the radicand in the latter equation is greater than $1/(2e)$ where e is the base of natural logarithms, so the equation may be rewritten as $\epsilon \geq (q_S + N - 1)\epsilon' \cdot 2e$. \square

Proof of Theorem 4.4. Given a verifiably-encrypted-signature forger algorithm \mathcal{V} , we construct a forger algorithm \mathcal{F} for the underlying co-GDH signature scheme. The co-GDH forger \mathcal{F} is given a public key v , and has access to a signing oracle for v and a hash oracle. It runs \mathcal{V} , answering its oracle calls, and uses \mathcal{V} 's verifiably encrypted signature forgery to forge a co-GDH signature.

We assume that \mathcal{V} is well-behaved in the sense that it always requests the hash of a message M before it requests a verifiably encrypted signature or an adjudication involving M , and that it never requests adjudication on a message M on which it had not previously asked for a verifiably encrypted signature. It is trivial to modify any forger algorithm \mathcal{V} to have the first property. The second property is reasonable since the input to the adjudication oracle in this case would be a nontrivial verifiably encrypted signature forgery; \mathcal{V} can be modified simply to output it and halt.

\mathcal{F} generates a key, (x', v') $\stackrel{R}{\leftarrow}$ $KeyGen$, which serves as the adjudicator's key. Now \mathcal{F} runs \mathcal{V} , providing as input the public keys v and v' . When \mathcal{V} makes an oracle query, \mathcal{F} answers, as follows.

Hash. \mathcal{V} requests a hash on some string M . \mathcal{F} makes a query on M to its own hash oracle, receiving some value $h \in G_2$, with which it responds to \mathcal{V} 's query.

VerSig Creation. \mathcal{V} requests a signature on some string M . (It will have already queried the hash oracle at M .) \mathcal{F} queries its signing oracle (for v) at M , obtaining $\sigma \in G_2$. It then selects r at random from \mathbb{Z}_p , and returns to \mathcal{V} the pair $(\sigma \cdot \psi(v')^r, \psi(g_1)^r)$.

Adjudication. \mathcal{V} requests adjudication for a verifiably encrypted signature (ω, μ) on a message M . \mathcal{F} checks that the verifiably encrypted signature is valid, then returns $\omega/\mu^{x'}$.

Finally, \mathcal{V} halts, either declaring failure, in which case \mathcal{F} , too, declares failure and halts, or providing a nontrivial verifiably encrypted signature (ω^*, μ^*) on a message M^* . \mathcal{F} sets $\sigma^* \leftarrow \omega^*/(\mu^*)^{x'}$ which, by the validity property, is a valid signature on M^* under key v .

That the forgery is nontrivial means that \mathcal{V} did not query the verifiably encrypted signature oracle at M^* , from which it follows that \mathcal{F} did not query its signing oracle at M^* . Thus (M^*, σ^*) is a nontrivial co-GDH forgery; \mathcal{F} outputs it and halts.

It remains only to analyze the success probability and running time of \mathcal{F} . \mathcal{F} succeeds whenever \mathcal{V} does, that is, with probability at least ϵ . \mathcal{F} 's running time includes \mathcal{V} 's, which is at most t .

The overhead imposed by \mathcal{F} is dominated by the need to evaluate two group exponentiations for each verifiably encrypted signature request, an exponentiation and an inversion for each adjudication request, and an exponentiation and an inversion in computing σ^* . Any one such operation may be performed using at most $2 \lg p$ group actions, and thus at most $2 \lg p$ time units, on G_2 . Thus \mathcal{F} 's overall running time is at most $t + 4c_{\mathcal{F}}(\lg p)(q_S + q_A + 1)$, Where $c_{\mathcal{F}}$ is a small constant that accounts for the remainder of \mathcal{F} 's administrative overhead; in practice, $c_{\mathcal{F}}$ should be at most 2.

\mathcal{F} queries its hash oracle whenever \mathcal{V} queries its hash oracle, and its signing oracle whenever \mathcal{V} queries its verifiably encrypted signature oracle.

Combining all this, we see that if \mathcal{V} $(t, q_H, q_S, q_A, \epsilon)$ -forges a bilinear verifiably encrypted signature on G_1 and G_2 , then \mathcal{F} $(t + 4c_{\mathcal{F}}(\lg p)(q_S + q_A + 1), q_H, q_S, \epsilon)$ -breaks the co-GDH signature scheme on G_1 and G_2 . Conversely, if the co-GDH signature scheme is $(t', q'_H, q'_S, \epsilon')$ -secure, then the bilinear verifiably encrypted signature scheme is $(t' - 4c_{\mathcal{F}}(\lg p)(q'_S + q_A + 1), q'_H, q'_S, q_A, \epsilon')$ -secure against existential forgery, where $c_{\mathcal{F}}$ is a small constant and q_A is at most q'_S . \square

Proof of Theorem 4.5. Given a verifiably-encrypted-signature extractor algorithm \mathcal{V} , we construct an aggregate extractor algorithm \mathcal{A} . The co-GDH forger \mathcal{A} is given values g_1^a and g_1^b in G_1 , g_2^u , g_2^v , and g_2^{au+bv} in G_2 . It runs \mathcal{V} , answering its oracle calls, and uses \mathcal{V} 's verifiably encrypted signature extraction to calculate g_2^{au} , the answer to its own extraction challenge.

We assume that \mathcal{V} is well-behaved in the sense that it always requests the hash of a message M before it requests a verifiably encrypted signature or an adjudication involving M . It is trivial to modify any extractor algorithm \mathcal{V} to have the this property.

\mathcal{A} needs to engage in a certain amount of bookkeeping. In particular, it must maintain a list of the messages on which \mathcal{V} requests hashes or signatures. Each message M , as it arrives from \mathcal{V} , is assigned an index i ; i is obviously bounded above by q_H . The message is stored in M_i and its hash in h_i . \mathcal{A} also maintains an array s_i of bits and an array x_i of elements of \mathbb{Z}_p , to be described later.

The behavior of \mathcal{A} will depend on a probability constant ζ , which will be optimized later to yield the best possible reduction. We define B_ζ to be the probability distribution over $\{0, 1\}$ where 1 is drawn with probability ζ , and 0 with probability $1 - \zeta$.

In setup, \mathcal{A} sets $v \leftarrow g_1^a$, the signer's public key, and $v' \leftarrow g_1^b$, the adjudicator's public key. Then, for each i , $1 \leq i \leq q_H$, \mathcal{A} does the following. It picks a random bit $s_i \stackrel{R}{\leftarrow} B_\zeta$, and a random number $x_i \stackrel{R}{\leftarrow} \mathbb{Z}_p$. If s_i equals 1, \mathcal{A} sets $h_i \leftarrow g_2^u \cdot g_2^{x_i}$; if s_i equals 0, \mathcal{A} sets $h_i \leftarrow g_2^{x_i}$.

Now \mathcal{A} runs \mathcal{V} , providing as input the public keys v and v' . When \mathcal{V} makes an oracle query, \mathcal{A} makes answer, as follows.

Hash. \mathcal{V} requests a hash on some string M_i . \mathcal{A} outputs h_i .

VerSig Creation. \mathcal{V} requests a signature on some string M_i . \mathcal{A} selects y at random from \mathbb{Z}_p . If s_i equals 1, \mathcal{A} computes and returns $(\omega, \mu) = (\psi(g_1^a)^{x_i} \cdot g_2^{au+bv} \cdot \psi(g_1^b)^y, g_2^v \cdot g_2^y)$. If s_i equals 0, \mathcal{A} computes and returns $(\omega, \mu) = (\psi(g_1^a)^{x_i} \cdot \psi(g_1^b)^y, g_2^y)$. It is easy to verify that (ϵ, μ) is in either case a correct verifiably encrypted signature on the message with hash h_i .

Adjudication. \mathcal{V} requests adjudication for a verifiably encrypted signature (ω, μ) on a message M .

\mathcal{A} first checks that the verifiably encrypted signature is valid. Then, if s_i equals 1, \mathcal{A} declares failure and halts. Otherwise, it computes and returns $\psi(g_1^a)^{x_i}$.

Finally, \mathcal{V} halts. It either concedes failure or returns a nontrivial extracted signature σ^* on some message $M^* = M_{i^*}$ for some i^* . For the extraction to be nontrivial, \mathcal{V} must not have asked for adjudication on a verifiably encrypted signature of M_{i^*} .

\mathcal{A} now proceeds only if $s_{i^*} = 1$; otherwise it declares failure and halts. Since $s_{i^*} = 1$, it follows that $h_{i^*} = g_2^u \cdot g_2^{x_{i^*}}$. The extracted signature σ must satisfy the co-GDH verification equation,

$e(g_1, \sigma^*) = e(v, h_{i^*})$. \mathcal{A} sets $\sigma \leftarrow \sigma^*/\psi(v)^{x_{i^*}}$. Then

$$\begin{aligned} e(g_1, \sigma) &= e(g_1, \sigma^*) \cdot e(g_1, \psi(v))^{-x_{i^*}} = e(v, h_{i^*}) \cdot e(v, \psi(g_1))^{-x_{i^*}} \\ &= e(v, g_2^u) \cdot e(v, g_2)^{x_{i^*}} \cdot e(v, g_2)^{-x_{i^*}} = e(g_1^a, g_2^u). \end{aligned}$$

Where in the last equality we substitute $v = g_1^a$. Thus $(g_1, g_1^a, g_2^u, \sigma)$ is a valid co-Diffie-Hellman tuple, so σ equals g_2^{au} , the answer to the aggregate extraction problem; \mathcal{A} outputs it and halts.

It remains only to analyze the success probability and running time of \mathcal{A} , and to optimize the success probability by choosing a suitable value for ζ .

\mathcal{A} succeeds whenever \mathcal{V} does, provided \mathcal{V} never requests adjudication for a verifiably encrypted signature on a message M_i for which $s_i = 1$ and \mathcal{V} extracts a signature on a message M_{i^*} such that $s_{i^*} = 1$.

Multiplication by g_2^u induces a permutation, so each h_i is uniformly distributed independently of s_i . Each adjudication request on some message M_i is an independent trial that causes \mathcal{A} to halt with probability ζ . Since \mathcal{V} makes at most q_A adjudication requests, the probability of \mathcal{A} 's successfully answering all adjudication queries is at least $(1 - \zeta)^{q_A}$. Disqualification at this phase is independent of \mathcal{A} 's eventual success, by an argument analogous to that deployed in Theorem 3.2; it amounts to a shortcutting of unlucky runs.

\mathcal{V} cannot previously have asked for a signature on M_{i^*} , so it can have no information about the value of s_{i^*} ; in the extracted signature, $s_{i^*} = 1$ occurs with probability ζ .

Thus, if \mathcal{V} succeeds with probability ϵ , \mathcal{A} succeeds with probability at least $\zeta \cdot (1 - \zeta)^{q_A} \cdot \epsilon$. This probability is maximized at $\zeta = 1/(q_A + 1)$, where it has the value

$$\frac{1}{q_A + 1} \cdot \left(1 - \frac{1}{q_A + 1}\right)^{q_A} \cdot \epsilon = \frac{1}{q_A} \cdot \left(1 - \frac{1}{q_A + 1}\right)^{q_A + 1} \cdot \epsilon.$$

\mathcal{A} 's success probability ϵ' is at least as great as this.

\mathcal{A} 's running time includes \mathcal{V} 's running time. The overhead imposed by \mathcal{A} is dominated by group exponentiations, of which one is needed for each hash, three for each verifiably encrypted signature creation, and one for each adjudication, along with the exponentiation and inversion required to compute σ . Any one such operation may be completed using at most $2 \lg p$ group actions, and thus at most $2 \lg p$ time units, on G_2 . \mathcal{A} 's overall running time is $t' \leq t + 2c_{\mathcal{A}}(\lg p)(q_H + 3q_S + q_A + 2)$, where $c_{\mathcal{A}}$ is a small constant, in practice at most 2, that accounts for \mathcal{A} 's administrative overhead.

To summarize: if there exists an encrypted-signature extractor algorithm \mathcal{V} that $(t, q_H, q_S, q_A, \epsilon)$ -extracts the bilinear verifiably encrypted signature scheme on G_1 and G_2 , then there exists an algorithm \mathcal{A} that $(t', 2, \epsilon')$ -extracts a subaggregate on G_1 and G_2 , where

$$t' \leq t + 2c_{\mathcal{A}}(\lg p)(q_H + 3q_S + q_A + 2) \quad \text{and} \quad \epsilon' \geq \frac{1}{q_A} \cdot \left(1 - \frac{1}{q_A + 1}\right)^{q_A + 1} \cdot \epsilon.$$

Conversely, if G_1 and G_2 are, together, $(t', 2, \epsilon')$ -secure against aggregate extraction, then there can exist no algorithm \mathcal{V} that $(t, q_H, q_S, q_A, \epsilon)$ -extracts the bilinear verifiably encrypted signature scheme on G_1 and G_2 , where

$$t \leq t' - 2c_{\mathcal{A}}(\lg p)(q_H + 3q_S + q_A + 2) \quad \text{and} \quad \epsilon \geq (q_A)\epsilon' / \left(1 - \frac{1}{q_A + 1}\right)^{q_A + 1}.$$

For all $q_A > 1$, the radicand in the latter equation is greater than $1/(2e)$ (where e again is the base of natural logarithms), so the equation may be rewritten as $\epsilon \geq q_A \epsilon' \cdot 2e$. \square

Proof of Theorem 5.2. The co-CDH problem can be solved by first solving two random instances of the following problem: given g_2^{ab}, g_1^a (and g_1, g_2), compute g_2^b . We shall construct an algorithm \mathcal{A} that solves this problem.

Initially \mathcal{A} picks x_2, \dots, x_n at random and sets $x_1 = 1$. It sets $v_i = (g_1^a)^{x_i}$. Algorithm \mathcal{F} is given the public keys v_1, \dots, v_n . Without loss of generality we may assume \mathcal{F} submits distinct queries (as previous replies can be cached); that for every ring-signing query on a message M , \mathcal{F} has previously issued a hash query for M ; and that \mathcal{F} issues a hash query on the message on which it attempts to forge a signature some time before giving its final output.

On a hash query, \mathcal{A} flips a coin that shows 0 with probability p and 1 otherwise (p shall be determined later). Then \mathcal{A} picks a random $r \xleftarrow{\mathbb{R}} \mathbb{F}_q$, and if the coin shows 0, \mathcal{A} returns $(g_2^{ab})^r$, otherwise it returns $\psi(g_1^a)^r$.

Suppose \mathcal{F} issues a ring sign query for a message M . By assumption, \mathcal{A} has previously issued a hash query for M . If the coin \mathcal{A} flipped for this h -query showed 0, then \mathcal{A} fails and exits. Otherwise \mathcal{A} had returned $h(M) = \psi(g_1^a)^r$ for some r . In this case \mathcal{A} chooses random $a_2, \dots, a_n \xleftarrow{\mathbb{R}} \mathbb{F}_q$, computes $a_1 = r - (a_2 x_2 + \dots + a_n x_n)$, and returns the signature $\sigma = \langle g_2^{a_1}, \dots, g_2^{a_n} \rangle$.

Eventually \mathcal{F} outputs a forgery $\langle \sigma_1, \dots, \sigma_n \rangle$ for a message M . Again by assumption, \mathcal{F} has previously issued a h -query for M . If the coin flipped by \mathcal{A} for this query did not show 0 then \mathcal{A} fails. Otherwise $h(M) = g_2^{abr}$ for some r chosen by \mathcal{A} , and \mathcal{A} outputs the r th root of $\sigma_1 \sigma_2^{x_2} \dots \sigma_n^{x_n}$.

\mathcal{F} cannot distinguish between \mathcal{A} 's simulation and real life. Also, \mathcal{A} will not fail with probability $p^{q_S}(1-p)$ which is maximized when $p = q_S/(q_S + 1)$, giving a bound of $1/e(1 + q_S)$. If it does not fail and \mathcal{F} successfully forges a ring signature then \mathcal{A} is successful and outputs g_2^b . \square