

A preliminary version of this paper appears in *Proceedings of the 14th ACM Conference on Computer and Communications Security, CCS 2007*, pp. 276–285, S. De Capitani di Vimercati and P. Syverson eds., ACM Press, 2007. This is the full version.

Ordered Multisignatures and Identity-Based Sequential Aggregate Signatures, with Applications to Secure Routing

ALEXANDRA BOLDYREVA* CRAIG GENTRY† ADAM O’NEILL‡
DAE HYUN YUM§

Abstract

We construct two new multiparty digital signature schemes that allow multiple signers to sequentially produce a compact, fixed-length signature. First, we introduce a new primitive that we call *ordered multisignatures* (OMS), which allows signers to attest to a common message as well as the order in which they signed. Our OMS construction substantially improves computational efficiency and scalability over any existing scheme with suitable functionality. Second, we design a new identity-based sequential aggregate signature scheme, where signers can attest to different messages and signature verification does not require knowledge of traditional public keys. The latter property permits savings on bandwidth and storage as compared to public-key solutions. In contrast to the only prior scheme to provide this functionality, ours offers improved security that does not rely on synchronized clocks or a trusted first signer. We provide formal security definitions and support the proposed schemes with security proofs under appropriate computational assumptions. We focus on potential applications of our schemes to secure network routing, but we believe they will find many other applications as well.

Keywords: Digital signatures, identity-based signatures, multisignatures, aggregate signatures, pairings, network security.

*School of Computer Science, College of Computing, Georgia Institute of Technology, 266 Ferst Drive, Atlanta, GA 30332, USA. E-mail: aboldyre@cc.gatech.edu. URL: <http://www.cc.gatech.edu/~aboldyre>. Supported in part by NSF CAREER award 0545659.

†Dept. of Computer Science, Stanford University, USA. E-Mail: cgentry@cs.stanford.edu.

‡School of Computer Science, College of Computing, Georgia Institute of Technology, 266 Ferst Drive, Atlanta, GA 30332, USA. E-mail: amoneill@cc.gatech.edu. URL: <http://www.cc.gatech.edu/~amoneill>. Supported in part by the grant of the first author.

§Pohang University of Science and Technology, Korea. E-Mail: dhyum@postech.ac.kr.

Contents

1	Introduction	3
1.1	Overview	3
1.2	Ordered Multisignatures	3
1.3	Identity-based Sequential Aggregate Signatures	5
1.4	Versions of this Paper and Corrections	6
2	Preliminaries	6
3	Ordered Multisignatures	8
3.1	OMS Schemes and Their Security	8
3.2	Our OMS Construction and Analysis	10
4	Identity-Based Sequential Aggregate Signatures	13
4.1	IBSAS Schemes and Their Security	13
4.2	Our IBSAS Construction	14
5	On the Hardness of IBSAS-CDH	17
6	Conclusions and Open Problems	18
7	Acknowledgments	19
A	Proof of Theorem 3.6	22
B	Proof of Theorem 4.3	25
C	Proof of Theorem 5.1	28
C.1	A Useful Lemma	28
C.2	Main Argument	28

1 Introduction

1.1 Overview

The current Internet design largely lacks the principles of AAA: Authentication, Authorization, and Accountability. It is understood that incorporation of these principles would make tackling security and reliability problems more tractable.

A large body of recent research focuses on identifying weak points in the current design and proposing fixes to the deployed infrastructure. For example, the Secure Border Gateway Protocol (S-BGP) initiative (and its variants) [26, 1, 28, 48, 46, 15, 17, 16], whose primary goal is to patch authenticity of route announcements in BGP, a path-vector protocol used in Internet routing, is currently under consideration for standardization by the IETF. While it is accepted that new security measures are necessary, many remain skeptical about the prospects of widespread adoption and deployment in the near future. The main technical reason is that secure networking adds additional overhead to in-use protocols. We view our role as cryptographers in this regard as designing suitable provably-secure mechanisms to address some of the identified weaknesses, which maintain as best as possible the design goals of the original protocols, especially in terms of router processing time and storage, bandwidth overhead, scalability, etc.

In line with this view, this work introduces two new “multiparty” digital signature schemes for efficiently enhancing authenticity in several network routing applications. Our schemes offer important performance and security improvements as compared to previous candidate solutions. We show that they provably provide security according to the corresponding security definitions (that are of independent interest) and under appropriate computational assumptions.

We clarify that we do *not* attempt to rigorously analyze all possible threats and assumptions about adversarial abilities in the network routing applications we discuss. Indeed, much work remains to be done in this regard, and the specific security requirements in these applications remains an issue of contention [38]. What we suggest is that our schemes appear to be useful towards future resolution of some of the security concerns that have been raised. We next discuss our schemes and their applications in more detail.

1.2 Ordered Multisignatures

DEFINITION AND MOTIVATION. We introduce a new primitive that we call *ordered multisignatures* (OMS). Recall that a multisignature scheme [6, 32] is a public-key primitive that allows multiple signers who want to sign some message to produce a single compact (constant-size) signature convincing a verifier that each signer signed the message. However, some network routing applications that we discuss below require verifying the *path* (i.e. the ordered list of routers) in which a packet travels to reach its destination, where routers have incentive to lie. Although by using multisignatures the routers could each sign (a fixed part of) the packet while keeping total packet overhead due to signatures fixed to a constant, this would be insufficient from a security standpoint because it does not allow to verify the order in which they signed. one way to ensure signing order in a (non-interactive) multisignature scheme would be to have each signer use a separate public-private key-pair for signing messages in each position on a path. But the resulting scheme would be impractical due to large combined key size. Accordingly, the OMS primitive we introduce produces a compact (constant-size) multisignature, uses constant-size keys, is “sequential” in that signers sign one after another and no further interaction among the signers is required, and ensures authenticity of both the signing order and that of the message.

Note that, as for (plain) multisignatures, construction of ordered multisignatures are only interesting if they are more efficient than aggregate signature schemes [10, 33, 32, 3], which allow

multiple signers to sign *different* messages while keeping total signature size constant. So they can of course immediately provide the needed functionality if the signers sign, in addition to the packet, their position on the path. But aggregate signatures are usually much less efficient than multisignatures.

FURTHER CONTRIBUTIONS. After introducing and defining the new primitive, we propose a formal security model for OMS. It adapts the notion of security for multisignatures first presented in [6] to also ensure authenticity of the signing order. Intuitively, a secure OMS scheme, in addition to being secure as a “plain” (un-ordered) multisignature scheme, must enforce an additional unforgeability with respect to the ordering of the signers. We then provide a construction that we prove secure in our model, under a standard computational assumption on the groups equipped with the “bilinear maps” (aka. pairings) we use, in the random oracle (RO) model of [5]. As compared to known aggregate schemes, our construction offers substantial computational savings. Namely, the work per required on both signing and verification is essentially *constant* in the number of signatures currently in the OMS. Section 3.2 gives detailed efficiency comparisons.

APPLICATIONS. We sketch some potential applications of our OMS scheme in more detail. One problem that appears suitable, raised in [22], is “data plane” security in S-BGP. This means allowing autonomous systems (ASes, i.e. networks under control of a single entity such as Georgia Tech or AT&T) to verify that data packets they send and receive/forward actually travel via previously-authenticated AS paths. (Authentication of AS paths is handled in the “control plane” by route attestation, explained in the following subsection.) To do so, a data packet should be signed, in order, by egress routers of ASes that forward it, allowing ingress routers to accept and forward only packets that followed an authenticated path, and the originating AS to later verify that the packet actually took an authenticated path to reach its destination (an OMS attesting to which could be piggybacked onto traffic on the reverse path).

Another setting where OMS could help arises in the recent in-band network troubleshooting system Orchid [37, 36]. In order to quickly and accurately diagnose faults (e.g. packet drops, re-orderings, duplications) along a flow from a sender to a receiver, Orchid has routers along the flow “mark” a fixed-size header in the data packets being sent. The first packet triggers a probe, which is sent to find out which routers are on the path. Later, a certain pattern of marks in the data packets by a router can implicate packet re-ordering or duplication by the *previous* router on the path, according to the data collected by the probe. When deployed across multiple networks (i.e. ASes), a router may wish to “frame” a router in another network by making it appear that the latter is directly upstream from it. Thus the probe should be signed, in order, by all the routers on the path, before fault data collected by the receiver can be considered authentic.

We suggest that the computational savings our scheme provides over existing solutions is desirable in the above-mentioned applications because it (1) distributes processing time more equitably amongst routers, (2) offers a sizable gain in total processing time, and (3) scales much better in the number of routers or ASes in the network.

RELATED WORK. Verifiability of signing order in multisignatures has been considered before, specifically in [20, 21, 13, 35, 45, 14], where they are usually called “structured” or “order-specified.” However, this line of work is in the *interactive* setting, meaning the schemes they consider require multiple rounds of communication between co-signers (sometimes even in key generation, requiring a separate interactive key-generation protocol for each subgroup of signers), which is impractical in applications we consider. Other differences between these works and ours include that they mainly treat more complicated structures on the group of signers than just linear ordering, and they do not give concrete applications of their schemes. In the interactive setting, the literature on “plain” multisignatures is extensive; see [6] for a comprehensive account.

One-way signature chains [40] are designed for a different setting than ours, in which a signer attests to which specific signers came before her (cf. Remark 3.4), and moreover their construction does not provide any efficiency gain over existing aggregate signature schemes.

1.3 Identity-based Sequential Aggregate Signatures

MOTIVATION AND PREVIOUS WORK. It has been pointed out in numerous works and further explored and tested in detail by [48] that aggregate signatures [10, 33, 32, 3], which allow multiple signers to sign different messages while keeping total signature size constant, can be used to address route announcement authenticity in S-BGP while significantly reducing associated bandwidth overhead and memory space for signatures. According to the proposal, each AS forwarding an update message should add its signature on the label of the *next* AS on the advertised route, so that route authenticity can be verified upon receipt of the aggregate. (This is a simplification that omits some details and optimizations such as “signature amortization;” see [48].) This technique is to prevent an unauthorized AS from extending the path and means that that signers need to sign genuinely *different* messages (as opposed to applications of OMS).

However, any public-key-infrastructure-based cryptographic proposal for networking applications requires all parties to know the authentic public keys of all other parties involved. This means that, in routing-based applications, these protocols incur the setup and storage overhead of distributing the public keys and corresponding certificates of all users out-of-band, and participating routers storing them indefinitely. Otherwise, public keys (which cannot be aggregated) and certificates of the signer in each signature would always have to be sent along with the latter for verification, defeating the purpose of using constant-size multiparty signatures to minimize bandwidth in the first place. As noted in previous works [24, 4, 47], identity-based cryptography [12], in which an arbitrary string (e.g. an IP address) acts as a user’s public key (the corresponding private key for which can be obtained by authenticating oneself to a trusted private key generator or PKG) and verifying a signature requires knowledge only of a sender’s identity in addition to a “master” public key of the PKG, can offer a superior alternative for such applications (subject to various trade-offs). This is because most of the information needed for verifying an aggregate signature is then already contained in the description of “who signed what.” It is a compelling setting in which to design and deploy aggregate schemes.

Yet the only (non-interactive) identity-based aggregate signature scheme to date is that of [24], which has the restriction that signers in a given aggregate must agree on a “common nonce” never used by any of them before; indeed, if a signer ever re-uses such a nonce in two different signatures, it then becomes simple to forge a signature by that signer on any message of one’s choice. From a functionality perspective, then, in order for the scheme to remain non-interactive, one possibility would be to simply trust the first signer in an aggregate (when signing is done in a “sequential” fashion) to pick a fresh random nonce each time. But there is no reason for this trust. Alternatively, one could rely on synchronized clocks of the signers and instantiate the nonce as a time-stamp; however, an honest computer’s perceived clock-time could be altered by a simple virus or after a power failure, which would lead to new potential attacks in practice. Therefore, the above restriction seems rather imprudent from the standpoint of security.

CONTRIBUTIONS. After defining the primitive, we design a security model for identity-based sequential aggregate signatures (IBSAS) which adapts the security model of [33] to the identity-based setting. (“Sequential” means that, as for OMS, signatures are aggregated one-by-one as the aggregate-so-far moves along the path, as is natural in the routing-based applications we consider.) Then we provide the first construction of an IBSAS scheme that does *not* place any such “common nonce” restriction on the signers. In other words, we show that in the sequential setting such a

restriction is not necessary. We prove our construction secure in the RO model under an interactive “CDH-type” computational assumption that arises from our scheme. To help justify the new assumption, we prove it holds in the generic bilinear group model of [8]. This proof constitutes a “heuristic” security argument showing the assumption holds unless adversarial algorithms exploit specific properties of the underlying algebraic group (i.e. special properties beyond its basic structure), which has become a common way of building confidence in new cryptographic assumptions about the “bilinear” groups we use (see e.g. [8, 9]). (On the other hand, we do not claim our assumption is particularly natural, and thus our IBSAS scheme should probably be viewed as secure in the generic bilinear group model only. It remains an important open problem to design such a scheme secure under a more standard computational assumption.)

APPLICATIONS IN MORE DETAIL. As we mentioned, our scheme seems to fit in nicely with route attestation in S-BGP, especially because storage overhead of the protocol has been cited as a major concern [16, 48]. Identities here would roughly consist of an organization name, AS number, and IP address range, which all together are vastly smaller than traditional public keys and certificates. Each identity would be bound to a secret key by a PKG, e.g. ICANN (Internet Corporation for Assigned Names and Numbers). Actually, in practice, the PKG would be in a hierarchy rooted at the latter (cf. [30]), whereby it can delegate generation of user secret keys to descendant PKGs, signatures under which can be verified with the public keys of the PKGs along the path to the root. (We provide an extension of our IBSAS scheme based on [25] that allows this to be done efficiently.) Note that the overhead associated with obtaining and storing these keys – which is equivalent to that of obtaining and storing public keys of a hierarchical certificate authority (CA) – is typically much smaller than that of obtaining traditional public keys and certificates of the signers themselves, which the identity-based setting eliminates.

FURTHER RELATED WORK. Herranz and Galindo et al. [27, 23] obtain results about identity-based signature schemes permitting aggregation of signatures from the same signer only. Append-only signatures [31] is an interesting public-key primitive suggested for use in S-BGP route attestation, but no construction yielding less than $\omega(\sqrt{n})$ -size signatures for n signers is currently known. We clarify that [24] appears to be the only previous (non-interactive) identity-based aggregate signature scheme in the literature; another recent scheme of [18] is interactive. Interactive (i.e. multi-round) identity-based multisignatures are also studied in [4].

1.4 Versions of this Paper and Corrections

Regrettably, the IBSAS scheme proposed in the prior version of this paper has been withdrawn. As shown by Hwang et al. [29], the assumption upon which its security was based is false and the scheme is insecure. However, in this full version of the paper we propose a corrected scheme, new computational assumption to base its security on, and a revised proof that the latter holds in the generic bilinear group model. The scheme has basically the same functionality and (even better) efficiency as the prior (insecure) one, thereby essentially recovering our claimed results.

2 Preliminaries

NOTATION AND CONVENTIONS. Let \mathbb{Z} denote the set of integers, \mathbb{N} the positive integers, and \mathbb{Z}_n the integers modulo a number $n \geq 2$. If \mathbb{G} is a prime-order group then \mathbb{G}^* is its set of generators, i.e. $\mathbb{G} = \mathbb{G} \setminus 1_{\mathbb{G}_T}$, where $1_{\mathbb{G}_T}$ denotes the identity in \mathbb{G} . We denote by $\{0, 1\}^*$ the set of all (binary) strings of finite length and by ε the empty string. If X is a string then $|X|$ is its length in bits. If X, Y are strings then $X||Y$ denotes an encoding from which X and Y are uniquely recoverable.

If S is a finite set then $s \stackrel{\$}{\leftarrow} S$ means that s is selected uniformly at random from S . For any $l \in \mathbb{N}$, we often write $s_1, s_2, \dots, s_l \stackrel{\$}{\leftarrow} S$ as shorthand for $s_1 \stackrel{\$}{\leftarrow} S; s_2 \stackrel{\$}{\leftarrow} S; \dots; s_l \stackrel{\$}{\leftarrow} S$. The notation $b \stackrel{\delta}{\leftarrow} \{0, 1\}$ means that a bit b is assigned value 1 with some probability $0 \leq \delta \leq 1$ and 0 otherwise. If A is a randomized algorithm then $x \stackrel{\$}{\leftarrow} A(y, z, \dots)$ means that x is assigned the output of running A on inputs y, z, \dots . If A is deterministic then we drop the dollar sign above the arrow. In either case, $A(y, z, \dots) \Rightarrow x$ denotes that A outputs x after being run on inputs y, z, \dots . All algorithms considered in this paper are efficient and possibly randomized unless indicated otherwise. An adversary is an algorithm with an overlying experiment. By convention, the running-time of an adversary includes that of its overlying experiment.

BILINEAR MAPS. Our schemes use bilinear maps (aka. pairings). Let \mathbb{G}, \mathbb{G}_T be groups of the same prime order p . Following a convention in the cryptographic community, we write both groups multiplicatively. A *pairing* is an efficiently computable map $\mathbf{e}: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ such that the following two conditions hold:

- Bilinearity: For all $u, v \in \mathbb{G}$ and $a, b \in \mathbb{Z}$, we have $\mathbf{e}(u^a, v^b) = \mathbf{e}(u, v)^{ab}$.
- Non-degeneracy: For any generator $g \in \mathbb{G}^*$, we have $\mathbf{e}(g, g) \neq 1_{\mathbb{G}_T}$, i.e. $\mathbf{e}(g, g)$ generates \mathbb{G}_T .

Observe that $\mathbf{e}(\cdot, \cdot)$ is symmetric since $\mathbf{e}(g^a, g^b) = \mathbf{e}(g, g)^{ab} = \mathbf{e}(g^b, g^a)$.

Definition 2.1 We call an algorithm \mathcal{G} that outputs (descriptions of) $p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}$ as above a *bilinear-group generation algorithm*, and \mathbb{G} a *bilinear group*. ■

In practice, \mathbb{G} is typically a subgroup of the group of rational points of an elliptic curve over a finite field. Using embedding degree (the degree of certain extension of the ground field) $k = 2$, for standard security levels (meaning discrete log computation in \mathbb{G} is believed to take at least 2^{80} basic operations), elements in \mathbb{G} can be represented using about 512 bits. It is also currently possible to reduce this length to 237 bits for the same security level by choosing $k = 6$, but there are fewer suitable curves known in this case [11]. It is possible that this bit-length will be further reduced in future research. Note that we purposely do not consider the “asymmetric” setting, as in $\mathbf{e}: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ on groups $\mathbb{G}_1 \neq \mathbb{G}_2$, because, although in this case elements in \mathbb{G}_1 could be represented using only 160 bits in this case, representation of elements in \mathbb{G}_2 would then require at least 1024 bits (due to the “MOV” attack [34]). Since our signatures would contain elements of both, their total length would actually be longer.

Although the bit-length of the representation of elements in \mathbb{G} is 512 bits with embedding degree $k = 2$, for computational efficiency the *order* of \mathbb{G} is usually be chosen to be about 2^{160} . In particular, this means that exponentiations in \mathbb{G} use exponents of only about 160 bits in length. Fast pairing implementation is a rapidly advancing area and the cost of computing a pairing is now down to about twice that of an exponentiation [39], which is quite practical. (However, the cost with embedding degree $k = 6$ is slightly more than that with embedding degree $k = 2$.)

Below, we formulate some computational problems in such groups that we use for our security proofs. Note that we omit formally defining what it means for these problems to be “hard” and therefore we do not explicitly make any assumptions about them as such, but such definitions and assumptions can be easily derived from our treatment in standard ways.

CDH PROBLEM. First we recall the well-known *computational Diffie-Hellman problem* (CDH) in bilinear groups.

Definition 2.2 Fix a bilinear group generator \mathcal{G} . We define the *CDH*-advantage of an algorithm A relative to \mathcal{G} as

$$\begin{aligned} \text{Adv}_{\mathcal{G}}^{\text{CDH}}(A) \\ \stackrel{\text{def}}{=} \Pr \left[C = g^{ab} : (p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}) \xleftarrow{\$} \mathcal{G} ; g \xleftarrow{\$} \mathbb{G}^* ; a, b \xleftarrow{\$} \mathbb{Z}_p ; C \xleftarrow{\$} A(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, g^a, g^b) \right]. \blacksquare \end{aligned}$$

IBSAS-CDH PROBLEM. For the security analysis of our IBSAS scheme we introduce the following computational problem that we call the *IBSAS-CDH problem*, which is a CDH-type problem that arises from our scheme.

Definition 2.3 Fix a bilinear group generator \mathcal{G} . For $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$ output by \mathcal{G} , we define for all $a_1, b_1, a_2, b_2 \in \mathbb{Z}_p$ and $g \in \mathbb{G}^*$ the associated oracle $\mathcal{O}_{g, a_1, b_1, a_2, b_2}^{\text{IBSAS-CDH}}(\cdot)$ taking as input $m \in \mathbb{Z}_p$ and defined as

$$\begin{aligned} \text{Oracle } \mathcal{O}_{g, a_1, b_1, a_2, b_2}^{\text{IBSAS-CDH}}(m) \\ r, x \xleftarrow{\$} \mathbb{Z}_p \\ \text{Return } (g^{rx} g^{a_1 b_1} g^{m a_2 b_2}, g^r, g^x) \end{aligned}$$

We then define the *IBSAS-CDH*-advantage of an algorithm A relative to \mathcal{G} as

$$\begin{aligned} \text{Adv}_{\mathcal{G}}^{\text{IBSAS-CDH}}(A) \stackrel{\text{def}}{=} \Pr \left[C = (m', g^{r'x'} g^{a_1 b_1} g^{m a_2 b_2}, g^{r'}, g^{x'}) : (p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}) \xleftarrow{\$} \mathcal{G} ; g \xleftarrow{\$} \mathbb{G}^* ; \right. \\ \left. a_1, b_1, a_2, b_2 \xleftarrow{\$} \mathbb{Z}_p ; C \xleftarrow{\$} A^{\mathcal{O}_{g, a_1, b_1, a_2, b_2}^{\text{IBSAS-CDH}}(\cdot)}(g, g^{a_1}, g^{b_1}, g^{a_2}, g^{b_2}) \right], \end{aligned}$$

where we require that $m' \in \mathbb{Z}_p^*$ was not queried by A to its oracle. \blacksquare

In Section 5, we show that the IBSAS-CDH problem is hard in the generic bilinear group model of [8]. This has become a standard way of building confidence in the hardness of computational problems in groups equipped with bilinear maps. (On the other hand, we do not claim that the IBSAS-CDH problem is a “natural” computational problem, and thus our construction should essentially only be viewed as secure in the generic bilinear group model. However, for the purposes of our security analysis, it is simpler to introduce the IBSAS-CDH problem.)

3 Ordered Multisignatures

Ordered multisignatures (OMS) are a natural extension of the notion of multisignatures [6] in which, intuitively, a (constant-size) ordered multisignature on a message attests not only to the fact that some specified group of signers signed it (as in a “plain” multisignature scheme), but also to the order in which they signed. Note that such (non-interactive) schemes are “sequential” by nature. As discussed in the Introduction, potential applications include BGP data-plane security [22] and security in in-band fault localization [37, 36]. One can easily construct an OMS scheme from any aggregate signature scheme; however, the benefit of our OMS construction is that it significantly improves computational efficiency and scalability over existing aggregate signature schemes.

3.1 OMS Schemes and Their Security

SYNTAX. We formally define the syntax of an OMS scheme.

Definition 3.1 We specify an OMS scheme $\text{OMS} = (\text{OPg}, \text{OKg}, \text{OSign}, \text{OVf})$ by four algorithms:

- A *parameter generation algorithm* OPg that returns some global information I for the scheme. This algorithm can be run by a trusted third-party or standards bodies.
- A *key generation algorithm* OKg run by a user that on the input global information I returns a public-private key-pair (pk, sk) .
- A *signing algorithm* OSign run by a user that on inputs its secret key sk , a message $m \in \{0, 1\}^*$, an OMS-so-far σ , and a list of $i - 1$ public keys $L = (pk_1, \dots, pk_{i-1})$. It returns a new OMS σ' , or \perp if the input is deemed invalid.
- A *verification algorithm* OVf that on inputs a list of public keys (pk_1, \dots, pk_n) , a message m , and an OMS σ returns a bit.

For consistency, we require that the probability that $\text{OVf}(L_n, m, \sigma_n) \Rightarrow 1$ is 1, for all $n \in \mathbb{N}$ and all $m \in \{0, 1\}^*$, where the probability is over the experiment

$$\begin{aligned}
& I \stackrel{\$}{\leftarrow} \text{OPg}; (pk_1, sk_1), \dots, (pk_n, sk_n) \stackrel{\$}{\leftarrow} \text{OKg}(I) \\
& \sigma_0, L_0 \leftarrow \varepsilon \\
& \text{For } i = 1, \dots, n \text{ do} \\
& \quad \sigma_i \stackrel{\$}{\leftarrow} \text{OSign}(sk_i, m, \sigma_{i-1}, L_{i-1}) \\
& \quad L_i \leftarrow (pk_1, \dots, pk_i).
\end{aligned}$$

We also require that $\text{OSign}(sk, m, \sigma, L) \Rightarrow \perp$ if $|L| > 1$ and $\text{OVf}(L, m, \sigma) \Rightarrow 0$ (see Remark 3.3). ■

SECURITY. We adapt the notion of security for multisignatures given in [6] to our context. Intuitively, a secure OMS scheme, in addition to being secure as a “plain” multisignature scheme, must enforce an additional unforgeability with respect to the ordering of the signers; it should not be possible to re-order the positions of honest signers in an OMS, even if all other signers are malicious. (Note that this also implies that ordered multisignatures cannot be “adversarially combined;” e.g. a forger who sees two ordered multisignatures on a message m by signers (A, B) and (separately) by (C, D) cannot produce a single ordered multisignature on m by signers (A, B, C, D) . Security of plain multisignatures does not prevent this.)

Similarly to the model of [6], we require users to prove knowledge of their secret keys during public-key registration with a CA. For simplicity, this is modeled by requiring an adversary to hand over secret keys of malicious signers. This is known as the registered- or certified-key model.

Definition 3.2 Let $\text{OMS} = (\text{OPg}, \text{OKg}, \text{OSign}, \text{OVf})$ be an OMS scheme. We consider the following *UF-OMS experiment* associated to OMS and a forger F with access to an oracle, which runs in three stages.

Setup: The experiment first runs OPg to obtain output I and then generates a challenge key-pair (pk, sk) by running OKg on input I .

Attack: F runs on inputs I, pk . F may query a key registration oracle with a key-pair (pk', sk') and coins c used for key generation, which records pk' as *registered* if $\text{OKg}(I; c) \Rightarrow (pk', sk')$. (This is a simplified model of a possibly more-complex key registration protocol with a CA that involves proofs of knowledge of secret keys.) F also has access to a signing oracle $\mathcal{O}_{\text{OSign}}(sk, \cdot, \cdot, \cdot)$, which on inputs m, σ, L returns \perp if not all public keys in L are registered and $\text{OSign}(sk, m, \sigma, L)$ otherwise.

Forgery: Eventually, F halts with outputs a list of public keys $L^* = (pk_1^*, \dots, pk_n^*)$, a message m^* , and a purported OMS signature σ^* . This output is considered to be a *forgery* if it holds that (1)

$\text{Ovf}(L^*, m^*, \sigma^*) = 1$, (2) $pk_{i^*}^* = pk$ for some $i^* \in \{1, \dots, n\}$, (3) all public keys in L^* except pk are registered, and (4) F did not query m^*, σ', L' to its signing oracle where $|L'| = i^* - 1$ for any $\sigma' \in \{0, 1\}^*$.

We define the *UF-OMS*-advantage $\text{Adv}_{\text{OMS}}^{\text{UF-OMS}}(F)$ of F against OMS as the probability that F outputs a forgery in the above experiment, taken over the coin flips of the parameter generation algorithm, the oracles, and any by F itself. We say that F *outputs lists* of length at most n_{\max} if all its lists of public keys used in calls to its signing oracle have length at most $n_{\max} - 1$ and that in its final output (i.e. L^* above) has length at most n_{\max} . ■

Remark 3.3 Our security model does not capture the natural requirement that an honest user should only sign at position i in an OMS if there are really currently $i - 1$ signers in it. (As is not the case in secure routing protocols, we do not assume that a signer knows *a priori* its signing position. Instead, she is to obtain this information from the data transmitted by the previous signer.) Otherwise, an adversary that modifies data in transit might simply tell the third signer on the path to sign at the tenth position, and the tenth to sign at the third, for example; the resulting OMS is not required to be invalid. The way we ensure this requirement is instead by the syntactic condition that the signing algorithm in the OMS definition above implicitly must verify validity of the signature-so-far relative to the other data in its input, in order to confirm the signing position.

Remark 3.4 Note that our security model guarantees authenticity of the message signed by an honest signer and her position in an OMS, but not of which specific signers signed before or will sign after her. For example, it would not correspond to a forgery in our model if an OMS σ on a message m valid for public keys (pk_1, pk_2, pk_3) is modified by a malicious signer to some σ' on m valid for (pk'_1, pk_2, pk'_3) , where pk'_1, pk'_3 belong to the latter. But this seems to be acceptable in the applications we consider:

- In in-band fault localization [37, 36], reports of packet loss or reordering by a particular router typically indicate a problem upstream, so a main security property we want is that an honest router should not appear to a receiver collecting fault statistics to be further upstream than it actually is — but this does not concern *who* is upstream from the router.
- In S-BGP data plane security [22], since the previously-authenticated AS paths that a data packet may travel are known, if such a packet (having been signed and verified by the previous nodes who have received it to be traveling on an authenticated path) is incorrectly routed to a malicious node, our security model still ensures the latter cannot modify the packet to then be accepted by an honest node.

However, it is beyond the scope of this paper to rigorously analyze the security requirements needed in these emerging applications (cf. [38]).

3.2 Our OMS Construction and Analysis

THE SCHEME. Our construction extends Boldyreva’s multisignature scheme [6] to suitably encode in an OMS the ordering of the signers in addition to the message they signed, by using a technique similar to that of [32]. Our scheme yields a constant-size OMS consisting of 2 group elements (about 1024 or 474 bits depending on implementation details; see Section 2) and is substantially more efficient than all existing aggregate signature alternatives. Unlike these alternatives, it requires essentially *constant work* (in the number of current signers in the OMS) by a user on both signing and verification (see below).

Construction 3.5 Let \mathcal{G} be a bilinear-group generation algorithm. To it we associate the following construction:

Global Parameters: The algorithm first runs \mathcal{G} to obtain output $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$ and chooses a random generator $g \in \mathbb{G}^*$ and cryptographic hash function $\mathbf{H}: \{0, 1\}^* \rightarrow \mathbb{G}$. (The analysis will model the latter as a random oracle (RO) [5], adjusting security definitions accordingly.) It returns $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, \mathbf{H})$ as the global information I for the scheme.

Key Generation: On input I , the algorithm chooses random $s, t, u \in \mathbb{Z}_p$ and returns $(S = g^s, T = g^t, U = g^u)$ as pk and (s, t, u) as sk .

Signing: On inputs $sk_i, m, \sigma, L = (pk_1, \dots, pk_{i-1})$, the algorithm first verifies that $\text{OVf}(L, m, \sigma) \Rightarrow 1$ (as defined below) and if not, outputs \perp . (This step is skipped for a first signer, i.e. if $i = 1$, for whom σ is defined as $(1_{\mathbb{G}}, 1_{\mathbb{G}})$.) Then it parses σ as (Q, R) and chooses random $r \in \mathbb{Z}_p$. It computes:

1. $R' \leftarrow R \cdot g^r$
 2. $X \leftarrow (R')^{t_i + iu_i}$
 3. $Y \leftarrow (\prod_{j=1}^{i-1} T_j(U_j)^j)^r$
 4. $Q' \leftarrow \mathbf{H}(m)^{s_i} \cdot Q \cdot X \cdot Y$
- Finally, it returns (Q', R') .

Verification: On inputs $(pk_1, \dots, pk_n), m, \sigma$, the algorithm first checks that all of pk_1, \dots, pk_n are distinct and outputs 0 if not.¹ Then it parses σ as (Q, R) and checks if

$$\mathbf{e}(Q, g) \stackrel{?}{=} \mathbf{e}(\mathbf{H}(m), \prod_{i=1}^n S_i) \cdot \mathbf{e}(\prod_{i=1}^n T_i(U_i)^i, R).$$

If so, it outputs 1. If not, it outputs 0. Consistency follows straightforwardly from the bilinearity condition of a pairing. **■**

Thus, an ordered multisignature in our scheme on a message m by n signers with public keys pk_1, \dots, pk_n , respectively, has the form

$$\left(\mathbf{H}(m)^{\sum_i s_i} (g^{\sum_i t_i + iu_i})^{\sum_i r_i}, g^{\sum_i r_i} \right),$$

where r_i is the randomness chosen by the i -th signer.

SECURITY. Intuitively, the following implies that our OMS scheme is secure (in the RO model) if the CDH is hard relative to its associated bilinear-group generator \mathcal{G} .

Theorem 3.6 Let \mathcal{G} be a bilinear-group generation algorithm and OMS be the associated OMS construction given by to Construction 3.5. Suppose there exists a forger F against OMS in the RO model that makes at most q_h queries to its hash oracle, at most q_s queries to its signing oracle, and outputs lists of length at most $n_{\max} \geq 1$. Then there is an algorithm B against the CDH relative to \mathcal{G} such that

$$\mathbf{Adv}_{\text{OMS}}^{\text{UF-OMS}}(F) \leq n_{\max} e(q_s + 1) \cdot \mathbf{Adv}_{\mathcal{G}}^{\text{CDH}}(B). \quad (1)$$

Furthermore, the running-time of B is at most that of A plus $\tau(\mathcal{G}) \cdot O((q_h + n_{\max}(q_s + 1)))$, where $\tau(\mathcal{G})$ is the maximum time for an exponentiation in the bilinear groups output by \mathcal{G} . **■**

Above and in Theorem 4.3, e denotes the base of the natural logarithm.

Proof: See Appendix A. **■**

¹This is needed for our security proof, but in all applications we consider repeating signers in the “signature path” is not needed anyway.

Interestingly, our security proof relies on specific properties of Boldyreva’s multisignature scheme [6], in the sense that if the recent standard model (random oracle devoid) multisignature scheme of Lu et al. [32] is “substituted” for the former in our OMS construction, our approach to proving security no longer seems to work. Our proof also leverages a technique of [7], originally developed for achieving “selectively-secure” identity-based encryption in the standard model.

RUNNING-TIME ANALYSIS. In our efficiency analysis, we assume that $|\mathbb{G}| = 2^{160}$, i.e. $|p| = 160$; see Section 2. Then, step 1 in the signing algorithm requires one 160-bit exponentiation. (By which we mean that the bit-length of the exponent here is about 160 bits.) In typical applications, steps 2, 3, and 4 can essentially be executed together in the time of one 3-term multi-exponentiation, which is faster than computing 1.5 individual exponentiations. This ignores the cost of computing (we re-name $i - 1$ as n here for consistency with the below) $\prod_{j=1}^n T_j(U_j)^j$ in step 3, so let us justify this. Computing $\prod_{j=1}^n (U_j)^j$ requires n $O(\log n)$ -bit exponentiations. So, even if n is a hundred, this is only about the cost of computing three 160-bit exponentiations. (In most applications, n will be much less.) Thus, signing time will remain dominated by the 3 pairing computations in the verification call – which can be reduced to the time of about 2.5 individual pairing computations using a technique of [42]) – and similarly verification requires essentially *constant work* in the number of signers in the OMS.

EFFICIENCY COMPARISON WITH [32]. As noted in the Introduction, one can construct an OMS scheme from any aggregate signature scheme, basically by having the i -th signer add its signature on $m \parallel i$ to the aggregate-so-far, where m is the common message. (We also enforce the requirement in Definition 3.1 that the signing algorithm of the derived OMS scheme verify validity of the signature-so-far in case this is not done by the signing algorithm of the aggregate scheme already, which only affects the comparison with [10] below anyway.) In terms of performance, the best alternative to our OMS scheme seems to be obtained from the “RO version” of the recent aggregate scheme of Lu et al. [32, Section 3.4], which, for basically the same amount of security and signature size,² requires an additional n 160-bit exponentiations on both signing and verification (where n is the number of signers currently in the aggregate).

Note that while an n -term multi-exponentiation could be used for these, it would require $(2^n - 2) + 2(160 - 1)$ multiplications (a 160-bit exponentiation by the square-and-multiply method requires 240 multiplications on average) and thus would only provide a speed-up for relatively small values of n . Moreover, it would incur an extra $512 \cdot (2^n - 1)$ bits of memory usage. We also stress that the bases for these n exponentiations vary from aggregate to aggregate. So, regardless of the computational technique employed, without requiring a prohibitively large amount of memory for pre-computation (and routing platforms are quite memory-constrained in the first place) the cost of computing the n exponentiations will still grow linearly in n . Therefore, the RO version of [32] scales more poorly and provides less equitable distribution of processing time amongst signers (e.g. different ASes) as n grows, giving it more limited applicability in real-world deployment scenarios as compared to our OMS scheme.

EFFICIENCY COMPARISON WITH [10, 33]. Aggregate signature length is just 160 bits and signing is, strictly speaking, more efficient in the aggregate scheme of [10] than in our OMS scheme, but verification is vastly slower, requiring a linear amount of *pairing computations* in the number of signatures in an aggregate, and verifying the aggregate-so-far is needed anyway upon signing in the derived OMS scheme (cf. Remark 3.3). (We comment that even if it was not, our OMS scheme may still be preferable to [10] due to slow verification time in the latter, which could be of particular

²Though [32] claims that if using “asymmetric” pairings, as in $\mathbf{e}: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, their aggregate signatures can have length 320 bits, this appears to be an oversight, since their signatures, like ours, would also contain an element of G_2 , whose representation, as we mentioned, would actually require much longer bit-length in this case.

concern with denial-of-service attacks on a verifier.) In most routing-based applications, however, a fixed 1024-bit packet size overhead for signatures is still within reason, and, as discussed in Section 2, some implementations may reduce this overhead to less than 500 bits, which is even more manageable.

Finally, we observe that the RSA-based aggregate scheme of [33] either requires proofs of knowledge of RSA keys on key-registration or having the signers' public exponents bigger than their 1024-bit moduli. As for RSA the former are much more expensive than those for discrete log and are not used in practice, this means that their scheme will similarly require a linear number of such costly 1024-bit exponentiations on both signing and verification.

4 Identity-Based Sequential Aggregate Signatures

It has been pointed out in numerous works and tested in [48] that aggregate signatures [10, 33, 32], which allow multiple signers to sign different messages while keeping total signature size constant, can be used to address route announcement authenticity in S-BGP while significantly reducing associated bandwidth overhead. According to the proposal, each AS forwarding an update message should add its signature on the address of the *next* AS on the route (the latter is to prevent an unauthorized AS from picking up the path and makes OMS insufficient here), so that route authenticity can be verified upon receipt of the aggregate.

However, as explained in the Introduction, using public-key schemes that necessitate a public-key infrastructure (PKI) dramatically increases set-up and storage or bandwidth overhead of secure networking protocols. But in the identity-based setting [12], where an arbitrary identifier can be used as a public key, most of the information needed to verify a signature is already included in the transmission anyway. We treat sequential aggregate signatures in this setting. IBSAS schemes appear well-suited for secure routing applications and in particular for route attestation in S-BGP, where storage overhead of the protocol is of particular concern [16]. Our construction improves upon the security of a previous solution in this setting by removing an undue restriction on the signers, making it more useful in practice.

4.1 IBSAS Schemes and Their Security

SYNTAX. We specify an *identity-based sequential aggregate signature* (IBSAS) scheme $AS = (\text{Setup}, \text{KeyDer}, \text{Sign}, \text{Vf})$ by four algorithms:

- A *setup algorithm* Setup initially run by the trusted private-key generator (PKG) to generate its master public key mpk and corresponding master secret key msk .
- A *private-key derivation* algorithm KeyDer run by the PKG on input msk, ID for any user's identity $ID \in \{0, 1\}^*$, to generate the private key sk_{ID} for user ID .
- A *signing algorithm* Sign run by a user ID on input its secret key sk_{ID} , a message $m \in \{0, 1\}^*$, an aggregate-so-far σ , a list $((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1}))$ of identity-message pairs to return a new aggregate signature σ' , or \perp to indicate that the input was deemed invalid.
- A *verification algorithm* Vf that takes as input the master public key mpk , a list $((ID_1, m_1), \dots, (ID_n, m_n))$ of identity-message pairs, and an IBSAS σ to return a bit.

For consistency, we require that $\text{Vf}(L_n, \sigma_n)$ always outputs 1, for all $n \in \mathbb{N}$ and all $\{(ID_i, m_i) \mid 1 \leq i \leq n, ID_i \in \{0, 1\}^*, m_i \in \{0, 1\}^*\}$ where the probability is over the experiment

$(mpk, msk) \xleftarrow{\$} \text{Setup}$
 For all $i = 1, \dots, n$ do
 $sk_{ID_i} \xleftarrow{\$} \text{KeyDer}(msk, ID_i)$
 $\sigma_0, L_0 \leftarrow \varepsilon$
 For $i = 1, \dots, n$ do
 $\sigma_i \xleftarrow{\$} \text{Sign}(sk_{ID_i}, m_i, L_{i-1}, \sigma_{i-1})$
 $L_i \leftarrow ((ID_1, m_1), \dots, (ID_i, m_i)). \blacksquare$

SECURITY. Our notion of security for IBSAS adapts of the notion of security for sequential aggregate signatures presented in [32] to the identity-based setting.

Definition 4.1 Let $AS = (\text{Setup}, \text{KeyDer}, \text{Sign}, \text{Vf})$ be an IBSAS scheme. We consider an experiment with a forger F with access to two oracles, that runs in three stages.

Setup: The experiment first generates a master key-pair via $(mpk, msk) \xleftarrow{\$} \text{Setup}$.

Attack: F runs on input mpk with access to a key-derivation oracle $\text{KeyDer}(msk, \cdot)$ and signing oracle $\mathcal{O}_{\text{Sign}}(\cdot, \cdot, \cdot, \cdot)$. The first operates according to the above definition of the private-key derivation algorithm for IBSAS. The second on input an identity ID , a message m , an aggregate-so-far σ , and a list of identity-message pairs $((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1}))$, sets $sk_{ID} \leftarrow \text{KeyDer}(msk, ID)$ and returns

$$\text{Sign}(sk_{ID}, m, ((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1})), \sigma) .$$

Forgery: Eventually, F halts with output a list of identity-message pairs $L^* = ((ID_1^*, m_1^*), \dots, (ID_n^*, m_n^*))$ and a purported aggregate signature σ^* . This output is considered to be a forgery if (1) all of ID_1^*, \dots, ID_n^* are distinct, (2) $\text{Vf}(L^*, \sigma^*) = 1$ and (3) there exists some $i^* \in \{1, \dots, n\}$ such that $ID_{i^*}^*$ was not queried by F to its key-derivation oracle and F did not query

$$(ID_{i^*}^*, m_{i^*}^*, ((ID_1, m_1), \dots, (ID_k^*, m_k^*)), \sigma')$$

to oracle $\mathcal{O}_{\text{Sign}}(\cdot, \cdot, \cdot, \cdot)$ for any $\sigma' \in \{0, 1\}^*$ and any $k \in \mathbb{N}$.

We denote by $\text{Adv}^{\text{UF-IBSAS}}(F)$ the probability that F outputs a forgery in the above experiment, taken over the coin flips of the setup algorithm, the oracles, and any by F itself. We say that F outputs lists of length at most n_{\max} if all its lists of identity-message pairs used in calls to its signing oracle have length at most $n_{\max} - 1$ and that in its final output (i.e. L^* above) has length at most n_{\max} . \blacksquare

COMPARISON TO PREVIOUS DEFINITIONS. Our definition of security for IBSAS, unlike that for public-key sequential aggregate signatures in [33] but similarly to that in [32], does not make the requirement that a signature appended to an aggregate cannot be re-used in another aggregate in which the signers and messages that come before it are different. However, one can easily convert a scheme meeting an “unordered” security definition into one meeting an “ordered” one by a simple modification; namely, have the i th signer ID_i sign not just m_i but $m_i \| ID_{i-1} \| m_{i-1} \| \dots \| ID_1 \| m_1$.

4.2 Our IBSAS Construction

THE SCHEME. We present our IBSAS construction. The scheme is quite efficient and yields constant-size aggregate signatures of 3 group elements. As we explain below, it improves function-

ality/security over the previous solution of [24] in that it lifts a “common nonce” restriction that can lead to some attacks on the scheme in practice.

Construction 4.2 Let \mathcal{G} be a bilinear-group generation algorithm. To it we associate the following construction:

Setup: The algorithm first runs \mathcal{G} to obtain output $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$. It then chooses random $\alpha_1, \alpha_2 \in \mathbb{Z}_p$, and cryptographic hash functions $H_1, H_2: \{0, 1\}^* \rightarrow \mathbb{G}$ and $H_3: \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$. It returns

$$(\mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, g^{\alpha_1}, g^{\alpha_2}, \{H_i\}_{i=1, \dots, 3})$$

as the *mpk* and (α_1, α_2) as the *msk*.

Key Derivation: On input *msk* and $ID \in \{0, 1\}^*$, the algorithm returns

$$(H_1(ID)^{\alpha_1}, H_2(ID)^{\alpha_2})$$

as sk_{ID} .

Signing: On input sk_{ID}, m, σ, L where $L = ((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1}))$ the algorithm first parses σ as $(\sigma_1, \sigma_2, \sigma_3)$ checks that $(\sigma_1, \sigma_2, \sigma_3)$ is a valid aggregate signature (according to the verification algorithm below) and returns \perp if not. (This step is skipped for a first signer, i.e. if $i = 1$, for whom σ is defined as $(1_{\mathbb{G}}, 1_{\mathbb{G}}, 1_{\mathbb{G}})$.) If so the algorithm computes:

1. $r, x \xleftarrow{\$} \mathbb{Z}_p$
2. $\sigma'_3 \leftarrow \sigma_3 \cdot g^x$
3. $\sigma'_2 \leftarrow \sigma_2 \cdot g^r$
4. $\sigma'_1 \leftarrow \sigma_1 \cdot (\sigma_3)^r \cdot (\sigma'_2)^x \cdot H_2(ID)^{\alpha_2} \cdot H_1(ID)^{\alpha_1 \cdot H_3(ID \| m)}$

It returns $(\sigma'_1, \sigma'_2, \sigma'_3)$.

Verification: On input *mpk*, $((ID_1, m_1), \dots, (ID_n, m_n)), \sigma$, the algorithm first checks that all of ID_1, \dots, ID_n are distinct and outputs 0 if not.³ It then parses σ as $(\sigma_1, \sigma_2, \sigma_3)$ and checks if

$$\mathbf{e}(\sigma_1, g) \stackrel{?}{=} \mathbf{e}(\sigma_2, \sigma_3) \cdot \mathbf{e}\left(\prod_{i=1}^n H_2(ID_i), g^{\alpha_2}\right) \cdot \mathbf{e}\left(\prod_{i=1}^n H_1(ID_i)^{H_3(ID_i \| m_i)}, g^{\alpha_1}\right)$$

If so, the algorithm returns 1; else it returns 0. ■

Thus, an aggregate signature in our scheme on messages m_1, \dots, m_n by signers ID_1, \dots, ID_n , respectively, has the form

$$\left(g^{xr} \cdot \prod_i H_2(ID_i)^{\alpha_2} \cdot H_1(ID_i)^{\alpha_1 \cdot H_3(ID_i \| m_i)}, g^r, g^x \right)$$

where $r = \sum_i r_i$ and $x = \sum_i x_i$.

SECURITY. The following establishes that our IBSAS scheme is secure (in the RO model) if the IBSAS-CDH problem is hard relative to its associated bilinear-group generator \mathcal{G} . We note that, alternatively, one could use the assumption in our analyses that the base scheme associated to our IBSAS construction is secure instead of the IBSAS-CDH assumption; however, the latter is slightly simpler and easier to work with.

³As we mentioned, the fact that we require all signers in an aggregate to be distinct is not a restriction in the applications we consider.

Theorem 4.3 Let \mathcal{G} be a bilinear-group generation algorithm and AS be the associated IBSAS scheme given by Construction 4.2. Suppose there exists a forger F against AS in the RO model that make at most q_{h_1} queries to its hash oracles, at most q_{k} queries to its key-derivation oracle, at most q_{s} queries to its signing oracle, and outputs lists of length at most n_{max} . Then there is an algorithm B for the IBSAS-CDH problem relative to \mathcal{G} such that

$$\mathbf{Adv}_{\text{AS}}^{\text{UF-IBSAS}}(F) \leq e(n_{\text{max}}(q_{\text{s}} + 1) + q_{\text{k}}) \cdot \mathbf{Adv}_{\mathcal{G}}^{\text{IBSAS-CDH}}(B) + q_{\text{h}}^2/l_{\text{min}}(\mathcal{G}),$$

where $l_{\text{min}}(\mathcal{G})$ is the minimum bit-length of the order p of a bilinear group output by \mathcal{G} . Furthermore, B makes at most q_{s} queries to its oracle and its running-time is at most that of A plus $\tau(\mathcal{G}) \cdot O(q_{\text{h}} + n_{\text{max}}(q_{\text{s}} + 1) + q_{\text{k}})$, where $\tau(\mathcal{G})$ is the maximum time for an exponentiation in a bilinear group output by \mathcal{G} . ■

Proof: See Appendix B. ■

We note that, in the random oracle model, hardness of the IBSAS-CDH problem as we have defined it is actually stronger than what we need to prove our IBSAS scheme secure. It suffices that given $g, g^{a_1}g^{b_1}, g^{a_2}, g^{b_2}$ and q -many instances of $(m, g^{rx}g^{a_1b_1}g^{ma_2b_2}, g^r, g^x)$ for *random* m (where q is an upper-bound on the number of signing queries the adversary against the IBSAS scheme may make), it is hard to produce $(m', g^{rx}g^{a_1b_1}g^{m'a_2b_2}, g^r, g^x)$ for (independently) *random* m' . Alternatively, if messages are small enough to be embedded injectively into \mathbb{Z}_p^* , then hash function H_3 (and indeed prefix ID_i of $ID_i\|m_i$) can be dropped from the scheme, which results in greater efficiency at the expense of relying on the stronger version of the assumption.

COMPARISON WITH [24] AND DESIGN RATIONALE. For comparison, we recall that the recent identity-based aggregate signature scheme of [24] produces an aggregate signature on messages m_1, \dots, m_n by signers ID_1, \dots, ID_n , respectively, of the form

$$\left(\text{H}_3(w)^{\sum_i^n r_i} \cdot \prod_i^n \text{H}_1(0\|ID_i)^\alpha \cdot \prod_i^n \text{H}_1(1\|ID_i)^{\alpha \text{H}_2(w\|ID_i\|m_i)}, g^{\sum_i^n r_i} \right);$$

here the private key of user ID_i consists of the pair $(\text{H}_1(0\|ID_i)^\alpha, \text{H}_1(1\|ID_i)^\alpha)$, $\text{H}_1, \text{H}_3: \{0, 1\}^* \rightarrow \mathbb{G}$ and $\text{H}_2: \{0, 1\}^* \rightarrow \mathbb{Z}_p$ are hash functions, and w is a nonce picked by the first signer. The element $\text{H}_3(w)$ provides a “common place” for the signers to “aggregate their randomness.” However, for security, the string w is then required to be a “common nonce” specific to each aggregate, meaning each of ID_1, \dots, ID_n above need to be sure that they have not used it in any other aggregate before. Indeed, if any signer repeats the same w in two different signatures, it becomes simple for an adversary to forge a signature by the signer on any message, via simple linear algebraic attacks in the exponent (cf. [24, Remark 4]). Unfortunately, this restriction still leaves their scheme vulnerable in practice in a different way: typical implementations of the scheme in the context of secure routing protocols would likely use a time-stamp as w and require signers to check that an aggregate-so-far has w sufficiently close to the signer’s current clock-time, but then the possibility of malicious altering of the latter, say by installing a simple virus that can get no information about the secret key, introduces potential real-world attacks.

Our IBSAS construction shows that in the sequential setting (e.g., for secure routing) $\text{H}_3(w)$ can be effectively formed as $g^{\sum_i x_i}$, where x_i is randomly chosen by the i -th signer, and thus it is not necessary to require the signers to use any “common nonce” in forming an aggregate signature. However, simply making the scheme to produce aggregate signatures of the form

$$\left(g^{x^r} \cdot \prod_i \text{H}_2(ID_i)^\alpha \cdot \text{H}_1(ID_i)^{\alpha \cdot \text{H}_3(ID_i\|m_i)}, g^r, g^x \right)$$

is *not* secure. To forge, it suffices for the adversary to find $(\sigma_1, \sigma_2, \sigma_3) \in \mathbb{G}$ such that $e(g, \sigma_1) \cdot e(\sigma_2, \sigma_3) = \mathbf{e}(g^\alpha, h)$, where h is some known element of \mathbb{G} . But this is easy: the adversary could just set $(\sigma_1, \sigma_2, \sigma_3) = (g^0, g^\alpha, h)$. Then there are straightforward ways to “randomize” this forgery so that it is not so obvious that it is a forgery. To fix this we use two different exponents α_1, α_2 in place of the two instances of α in the above signature. Then for the analogous attack the adversary would need to find $(\sigma_1, \sigma_2, \sigma_3) \in \mathbb{G}$ such that $\mathbf{e}(g, \sigma_1) \cdot \mathbf{e}(\sigma_2, \sigma_3) = \mathbf{e}(g^{\alpha_1}, h_1) \cdot \mathbf{e}(g^{\alpha_2}, h_2)$ for some known elements $h_1, h_2 \in \mathbb{G}$, which seems hard if the adversary does not know the discrete log to the base g of one of the latter.

FURTHER DISCUSSION. We caution that the IBSAS-CDH problem we introduce to prove our IBSAS scheme secure is quite strong. However, in Section 5, we prove its hardness in the generic bilinear group model of [8]. Intuitively, this result means that breaking it in practice is likely to nevertheless require a significant new advance in our current understanding of the appropriate elliptic curve groups in which our scheme can be implemented. We also point out that, given the “common nonce” restriction in the previous scheme of [24] (which, under this restriction, is shown to be secure based on CDH) and the potential attacks on that scheme in practice that can result, it is unclear in general why one should prefer to use the former, even if one strongly favors schemes that are “proven secure” based on more standard computational problems.

EXTENSION TO HIERARCHICAL PKGS. We sketch an extension of our IBSAS scheme to the setting where users’ private keys are generated via PKGs in a hierarchy. (That is, where there are multiple PKGs situated as nodes in a tree, and any PKG can delegate user secret key generation to descendant PKGs; a user obtains its private key from one of them determined by the particular application.) This is useful for S-BGP (cf. [30]). Although this functionality can be achieved in a generic way by using any public-key signature scheme to create certificates binding PKGs to their public keys appropriately, the advantage of our extension is that it eliminates the need for these certificates as well as the overhead of using an additional scheme.

The extension, based on the scheme of [25], goes as follows. Let P_i denote a PKG at a level i in the hierarchy. In the initial setup, the root PKG P_0 selects $\mathbb{G}, \mathbb{G}_T, g, \alpha_1, \alpha_2, H_1, H_2, H_3$ as in the basic scheme and additionally chooses another cryptographic hash function $H_4: \{0, 1\}^* \rightarrow \mathbb{G}$. The global parameters are $\mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, H_1, H_2, H_3, H_4$, and in addition each P_i chooses random $\{\alpha_{i,b}\} \in \mathbb{Z}_p$ for $b \in \{1, 2\}$ as its secret key and publishes $\{g^{\alpha_{i,b}}\}$ for $b \in \{1, 2\}$ as its public key. For a PKG P_ℓ below the root, P_ℓ ’s parent provides it with a secret value $S_\ell = \prod_{j=1}^{\ell} H_4(P_1 \| \dots \| P_j)^{\alpha_{j-1}}$, where (P_1, \dots, P_ℓ) is the path in the tree from the root PKG P_0 to P_ℓ . As the private key for a user ID , a PKG P_i provides $H_1(ID)^{\alpha_{i,1}}$ and $H_2(ID)^{\alpha_{i,2}} \cdot S_i$ (where $S_0 = 1_{\mathbb{G}}$ by convention). The signing and verification algorithms of our scheme can then be extended straightforwardly. Verifying an aggregate containing a signature by a user ID whose PKG is P_i requires the verifier to have obtained the public keys of all the PKGs on the path in the tree from P_0 to P_i . But as long as the hierarchy is not too large, a value $e_i = \mathbf{e}(H_4(P_1 \| \dots \| P_i), g^{\alpha_{i-1}})$ can be cached, and thus verification does not take much longer than in the basic scheme. (Recall that if e_i is cached, $\mathbf{e}(H_4(P_i)^\beta, g^{\alpha_{i-1}})$ for $\beta \in \mathbb{Z}_p$ can be computed as e_i^β ; one can also use “compressed pairings” [43] here.) Adapting Definition 4.1 to the hierarchical setting can be done following [25]; security of the extended scheme follows from [25, Theorem 2] and Theorem 4.3.

5 On the Hardness of IBSAS-CDH

While it would certainly be preferable to prove the security of our IBSAS scheme under the hardness of a more established computational problem, given the new functionality that the scheme provides, this is not always possible. To more carefully justify our use of the IBSAS-CDH problem,

we consider the generic group model [44], which models an inability of algorithms to use group representation (i.e. special properties of a group beyond the mere fact that it is a group) in solving a computational problem. Below, we establish a strong lower bound on the hardness of IBSAS-CDH in the extension of the generic group model to the bilinear group setting [8]. This has become a standard way of building confidence in the hardness of new computational problems in bilinear groups (see e.g. [8, 9]).

THE GENERIC BILINEAR GROUP MODEL. We briefly recall the model of [8], making only minor syntactic changes. For simplicity, we fix a bilinear-group generation algorithm \mathcal{G} that always outputs some fixed $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e})$. Let g, g_T be generators of \mathbb{G}, \mathbb{G}_T , respectively. An algorithm A being executed in the model means that it is run by a corresponding *generic bilinear group experiment* that encodes elements of these groups given to A as random strings of length $\lceil \log p \rceil$ via injective maps $\xi, \xi_T: \mathbb{Z}_p \rightarrow \{0, 1\}^{\lceil \log p \rceil}$, where $\xi(a)$ is the encoding of g^a and $\xi_T(a)$ the encoding of $(g_T)^a$ for all $a \in \mathbb{Z}_p$. That is, any group elements in A 's input (its input being that in its usual experiment, as well as $1_{\mathbb{G}}$), in A 's oracle queries and the responses it gets back, and in A 's output are so-encoded. At any time-step, A can in particular query one of two group operation oracles for \mathbb{G}, \mathbb{G}_T , respectively, with encodings $\gamma_1, \gamma_2 \in \{0, 1\}^{\lceil \log p \rceil}$ and a bit b to get back $\xi(\xi^{-1}(\gamma_1) \cdot (\xi^{-1}(\gamma_2))^{-b})$, where “ \cdot ” denotes the corresponding group operation. Likewise, it can query a bilinear map oracle with encodings $\gamma_3, \gamma_4 \in \{0, 1\}^{\lceil \log p \rceil}$ to get back $\xi_T(\mathbf{e}(\xi^{-1}(\gamma_3), \xi^{-1}(\gamma_4)))$. We use the same notation for the algorithm's advantage when executed in the model as for its advantage in its usual experiment.

OUR RESULT. Intuitively, the following establishes that the IBSAS-CDH is (unconditionally) hard in the generic bilinear group model.

Theorem 5.1 Let \mathcal{G} be as above. Suppose there is an algorithm A solving the IBSAS-CDH relative to \mathcal{G} in the generic bilinear group model that runs in time at most t and makes at most q_s queries to its oracle. Then we have that

$$\text{Adv}_{\mathcal{G}}^{\text{IBSAS-CDH}}(A) \leq \frac{4 \binom{t+3q_s+5}{2} + 4}{p} . \blacksquare$$

Proof: See Appendix C. \blacksquare

As

$$\binom{t+3q_s+5}{2} \leq (t+3q_s+5)^2 ,$$

the theorem shows that, asymptotically, an algorithm's advantage in solving the IBSAS-CDH in the generic bilinear group model can increase at most quadratically in the work it performs. This is fairly standard, e.g. for computing discrete logs. In practice, q_s corresponds roughly to the maximum number of signatures that an adversary may see, so could be set to about, say, 2^{30} .

6 Conclusions and Open Problems

This work presented two new cryptographic schemes for use in securing several network routing applications, which we believe to be more attractive in practice than existing alternatives. To conclude, we point out some interesting open problems in this area. Our results indicate that it would be useful to devise an identity-based OMS scheme that is more efficient than existing identity-based aggregate constructions. It also remains an excellent open problem to devise an identity-based aggregate signature scheme based on a more standard computational problem (e.g. CDH), but without the limitations of previous constructions. Secondly, it is important to devise such schemes secure in the standard model (without random oracles).

7 Acknowledgments

We thank Nick Feamster, Murtaza Motiwala, Krzysztof Pietrzak, Michel Abdalla, Younho Lee, Brent Waters, Gene Tsudik, Xavier Boyen, Farbod Shokrieh, Yael Tauman Kalai, and the anonymous reviewers of CCS 2007 for helpful comments, suggestions, and references. Alexandra Boldyreva was supported in part by NSF CAREER award 0545659. Craig Gentry was supported by the Herbert Kunzel Stanford Graduate Fellowship. Adam O’Neill was supported in part by the above-mentioned grant of the first author. Dae Hyun Yum was supported by the Korea Research Foundation Grant funded by the Korean Government (MOEHRD), KRF-2006-352-D00177.

References

- [1] W. Aiello, J. Ioannidis, and P. McDaniel. Origin authentication in interdomain routing. In *ACM CCS*, 2003. (Cited on page 3.)
- [2] M.-H. Au, W. Susilo, and Y. Mu. Practical compact e-cash. Cryptology ePrint Archive, Report 2007/148, 2007. <http://eprint.iacr.org/>. (Cited on page .)
- [3] M. Bellare and C. Namprempre and G. Neven. Unrestricted Aggregate Signatures. In *ICALP*, 2007. (Cited on page 3, 5.)
- [4] M. Bellare and G. Neven. Identity-based multi-signatures from RSA. In *CT-RSA*, 2007. (Cited on page 5, 6.)
- [5] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS*, 1993. (Cited on page 4, 11.)
- [6] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the Gap-Diffie-Hellman-Group signature scheme. In *Public Key Cryptography*, 2003. (Cited on page 3, 4, 8, 9, 10, 12.)
- [7] D. Boneh and X. Boyen. Efficient selective-ID secure identity-based encryption without random oracles. In *EUROCRYPT*, 2004. (Cited on page 12, 23.)
- [8] D. Boneh and X. Boyen. Short signatures without random oracles. In *EUROCRYPT*, 2004. (Cited on page 6, 8, 17, 18.)
- [9] D. Boneh, X. Boyen, and E.-J. Goh. Hierarchical identity based encryption with constant size ciphertext. In *EUROCRYPT*, 2005. (Cited on page 6, 18.)
- [10] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In *EUROCRYPT*, 2003. (Cited on page 3, 5, 12, 13.)
- [11] D. Boneh, B. Lynn, and H. Shacham. Short Signatures from the Weil Pairing. In *Journal of Cryptology*, Volume 17 , Issue 4, 2004. (Cited on page 7.)
- [12] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, 1984. (Cited on page 5, 13.)
- [13] M. Burmester, Y. Desmedt, H. Doi, M. Mambo, E. Okamoto, M. Tada, and Y. Yoshifuji. A structured ElGamal-type multisignature scheme. In *PKC*, 2000. (Cited on page 4.)

- [14] C.-Y. Lin, T.-C. Wu, and F. Zhang. A Structured Multisignature Scheme from the Gap Diffie-Hellman Group. Cryptology ePrint Archive, Report 2003/090, 2003. <http://eprint.iacr.org/>. (Cited on page 4.)
- [15] K. Butler and W. Aiello. Optimizing BGP security by exploiting path stability. In *ACM CCS*, 2006. (Cited on page 3.)
- [16] K. Butler, F. Farley, P. McDaniel, and J. Rexford. A survey of BGP security. Apr. 2005. <http://www.research.att.com/jrex/>. (Cited on page 3, 6, 13.)
- [17] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocols. In *ACM SIGMETRICS*, 2006. (Cited on page 3.)
- [18] X. Cheng, J. Liu, and X. Wang Identity-based aggregate and verifiably encrypted signatures from bilinear pairing. In *ICCSA*, 2005. (Cited on page 6.)
- [19] J.-S. Coron. On the exact security of full domain hash. In *CRYPTO*, 2000. (Cited on page 22, 25.)
- [20] H. Doi, E. Okamoto, and M. Mambo. Multisignature schemes for various group structures. In *Symposium on Cryptography and Information Security*, 1994. (Cited on page 4.)
- [21] H. Doi, E. Okamoto, M. Mambo, and T. Uyematsu. Multisignature scheme with specified order. In *Conference on Communication, Control, and Computing*, 1999. (Cited on page 4.)
- [22] N. Feamster, H. Balakrishnan, and J. Rexford. Some foundational problems in interdomain routing. In *HotNets*, 2004. (Cited on page 4, 8, 10.)
- [23] D. Galindo, J. Herranz, and E. Kiltz. On the generic construction of identity-based signatures with additional properties. In *ASIACRYPT*, 2006. (Cited on page 6.)
- [24] C. Gentry and Z. Ramzan. Identity-based aggregate signatures. In *Public Key Cryptography*, 2006. (Cited on page 5, 6, 15, 16, 17.)
- [25] C. Gentry and A. Silverberg. Hierarchical ID-based cryptography. In *ASIACRYPT*, 2002. (Cited on page 6, 17.)
- [26] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. Mcdaniel, and A. Rubin. Working around BGP: An incremental approach to improving security and accuracy in interdomain routing. In *NDSS*, 2003. (Cited on page 3.)
- [27] J. Herranz. Deterministic identity-based signatures for partial aggregation. *J. Comput.*, 49(3), 2006. (Cited on page 6.)
- [28] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. In *ACM SIGCOMM*, 2004. (Cited on page 3.)
- [29] J.Y. Hwang, D.H. Lee, and M. Yung. Universal forgery of the identity-based sequential aggregate signature scheme. In *ASIACCS*, 2009. (Cited on page 6.)
- [30] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure border gateway protocol (S-BGP) - Real world performance and deployment issues. In *NDSS*, 2000. (Cited on page 6, 17.)
- [31] E. Kiltz, A. Mityagin, S. Panjwani, and B. Raghavan. Append-only signatures. In *ICALP*, 2005. (Cited on page 6.)

- [32] S. Lu, R. Ostrovsky, A. Sahai, H. Shacham, and B. Waters. Sequential aggregate signatures and multisignatures without random oracles. In *EUROCRYPT*, 2006. (Cited on page 3, 5, 10, 12, 13, 14, 23.)
- [33] A. Lysyanskaya, S. Micali, L. Reyzin, and H. Shacham. Sequential aggregate signatures from trapdoor permutations. In *EUROCRYPT*, 2004. (Cited on page 3, 5, 12, 13, 14.)
- [34] A. Menezes, T. Okamoto, S. A. Vanstone, Reducing elliptic curve logarithms to logarithms in a finite field. In *IEEE Transactions on Information Theory*, Vol. 39, Num. 5, 1993. (Cited on page 7.)
- [35] S. Mitomi and A. Miyaji. A multisignature scheme with message flexibility, order flexibility and order verifiability. In *ACISP*, 2000. (Cited on page 4.)
- [36] M. Motiwala, A. Bavier, and N. Feamster. In-band network path diagnosis. *Georgia Tech Technical Report GT-CS-07-07*. (Cited on page 4, 8, 10.)
- [37] M. Motiwala and N. Feamster. Position paper: Network troubleshooting on data plane coat-tails. In *WIRED*, 2006. (Cited on page 4, 8, 10.)
- [38] S.M. Bellovin. Position paper: Workable Routing Security. In *WIRED*, 2006. (Cited on page 3, 10.)
- [39] M. Naehrig. Personal communication, June 2009. (Cited on page 7.)
- [40] A. Saxena and B. Soh. One-way signature chaining - a new paradigm for group cryptosystems. Cryptology ePrint Archive, Report 2005/335, 2005. <http://eprint.iacr.org/>. (Cited on page 5.)
- [41] J. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. In *Journal of the ACM*, Vol. 27, 1980. (Cited on page 28.)
- [42] R. Granger and N.P. Smart. On computing products of pairings. Cryptology ePrint Archive, Report 2006/172, 2006. <http://eprint.iacr.org/>. (Cited on page 12.)
- [43] M. Scott and P.S.L.M. Barreto. Compressed Pairings. In *CRYPTO*, 2004. (Cited on page 17.)
- [44] V. Shoup. Lower bounds for discrete logarithms and related problems. In *EUROCRYPT*, 1997. (Cited on page 18.)
- [45] M. Tada. An order-specified multisignature scheme secure against active insider attacks. In *Australian Conference on Information Security and Privacy*, 2002. (Cited on page 4.)
- [46] T. Wan, E. Kranakis, and P. van Oorschot. Pretty secure BGP, psBGP. In *NDSS*, 2005. (Cited on page 3.)
- [47] S. Xu and Y. and W. Susilo. Online/offline signatures and multisignatures for AODV and DSR routing security. In *Australasian Conference on Information Security and Privacy*, 2006. (Cited on page 5.)
- [48] M. Zhao, S. Smith, and D. Nicol. Aggregated path authentication for efficient BGP security. In *ACM CCS*, 2005. (Cited on page 3, 5, 6, 13.)

Simulator $B(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, g^a, g^b)$
 $k^* \xleftarrow{\$} \{1, \dots, n_{\max}\}; t, u \xleftarrow{\$} \mathbb{Z}_p$
Initialize arrays K, δ, H, E to everywhere undefined
Run F on inputs $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, H), (g^a, (g^a)^{-uk^*} g^t, (g^a)^u)$:
On key registration query (pk', sk', c) :
If $\text{OKg}(I; c) \Rightarrow (pk', sk')$ then
 $K[pk'] \leftarrow sk'$; Return 1
Else return 0
On hash query m :
 $E[m] \xleftarrow{\$} \mathbb{Z}_p; \delta[m] \xleftarrow{\delta} \{0, 1\}$
If $\delta[m] = 1$ then $H[m] \leftarrow g^{E[m]}$; Else $H[m] \leftarrow g^b g^{E[m]}$
Return $H[m]$
On signing query $(m_j, \sigma, L = (pk_1, \dots, pk_{i-1}))$:
Parse σ as (Q, R) and set it as $(1_{\mathbb{G}}, 1_{\mathbb{G}})$ if $i = 1$
If $i > 1$ and $\text{OVf}(m_j, \sigma, L) \Rightarrow 0$ or $\exists z \in \{1, \dots, i-1\}$ such that $K[pk_z]$ is undefined
then return \perp
If $\delta[m_j] = 0$ and $i = k^*$ then abort
 $r \xleftarrow{\$} \mathbb{Z}_p$; Let $K[pk_z] = (s_z, t_z, u_z)$ for all $z \in \{1, \dots, i-1\}$
If $\delta[m_j] = 1$ then {
 $Q' \leftarrow (g^a)^{E[m_j]} ((g^a)^{-uk^*} g^t (g^a)^{iu})^r$; $R' \leftarrow g^r$
 $Q'' \leftarrow Q' \cdot \prod_{k=1}^{j-1} (g^{s_k})^{E[m_j]} g^{(t_k + ku_k)r}$ }
Else {
 $Q' \leftarrow (g^a)^{E[m_j]} (g^b)^{-t/(u(i-k^*))} ((g^a)^{-uk^*} g^t (g^a)^{iu})^r$; $R' \leftarrow g^r (g^b)^{1/(u(k^*-i))}$
 $Q'' \leftarrow Q' \cdot \prod_{k=1}^{j-1} (g^b g^{E[m_j]})^{s_k} (g^r (g^b)^{-1/(u(i-k^*))})^{(t_k + ku_k)}$ }
Return (Q'', R')
Let $(L^* = (pk_1^*, \dots, pk_n^*), m^*, \sigma^* = (Q^*, R^*))$ be the output of F
If L^*, σ^* is not a forgery (relative to H -values) then return \perp
Let $i^* \in \{1, \dots, n\}$ satisfy condition (2) of a forgery
If $\delta[m^*] = 1$ or $i^* \neq k^*$ then abort
Let $K[pk_z^*] = (s_z, t_z, u_z)$ for all $z \in \{1, \dots, n\}$ such that $z \neq i^*$
 $Q \leftarrow Q^* / (\prod_{j \neq i^*} (g^b g^{E[m^*]})^{s_j} (R^*)^{t_j + ju_j})$; $Z \leftarrow Q / ((R^*)^t (g^a)^{E[m^*]})$
Return Z

Figure 1: Simulator B for the proof of Theorem 3.6.

A Proof of Theorem 3.6

We construct a “simulator” B that on inputs $p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, g^a, g^b$, runs F to solve the CDH.

THE SIMULATOR. For simplicity, we assume F always queries messages to its hash oracle prior to using them in its signing queries and its final output, and that F never repeats a hash query. The description of the simulator B for the proof is given in Figure 1. In responding to hash queries by F , using Coron’s technique [19] we have B assign query m a bit (aka. δ -value) $\delta[m]$ equal to 1 with probability δ , for some value of $0 \leq \delta \leq 1$ that we optimize later. Intuitively, B hopes that F never queries a message with δ -value 0 to its signing oracle where the honest signer is at the k^* -th

position in the OMS, but that F 's forgery contains such a message and that the position of the honest signer in the forgery is k^* .

ANALYSIS. We first need the following claim.

Claim A.1 On executions of B on which it does not abort, F 's view (consisting of its input and answers it receives to its oracle queries) in the simulation provided by B comes from an identical distribution to that in its real UF-OMS experiment.

Proof: We first note that, as in the proof of [32, Theorem 3.1], the ‘signature reconstruction’ technique used by B to answer F 's signing queries provides responses coming from the same distribution as in the UF-OMS experiment because the OMS-so-far is verified and re-randomized in the signing algorithm, and the distribution of the resulting OMS is independent of the order in which the individual components of each signers were actually combined (which is why we require verifying the OMS-so-far in the signing algorithm anyway in Definition 3.1). Moreover, in the case that F makes a signing query m such that $\delta[m] = 0$ but the position of the honest signer in the OMS is some $i \neq k^*$, our code for B first uses a trick of Boneh and Boyen [7] to first create the honest signer's individual component in the OMS with the correct distribution from F 's perspective. To see this, we can write the honest signer's component (Q', R') in this case as

$$\begin{aligned}
Q' &= (g^a)^{\text{E}[m_j]} (g^b)^{-t/(u(i-k^*))} ((g^a)^{-uk^*} g^t (g^a)^{iu})^r \\
&= (g^a)^{\text{E}[m_j]} g^{-bt/(u(i-k^*))} (g^{(i-k^*)au} g^t)^r \\
&= (g^a)^{\text{E}[m_j]} g^{ab} (g^{(i-k^*)au} g^t)^{r-b/(u(i-k^*))} \\
&= (g^b g^{\text{E}[m_j]a})^a (g^{(i-k^*)au} g^t)^{r-b/(u(i-k^*))} ; \\
R' &= g^r (g^b)^{-1/(u(i-k^*))} \\
&= g^{r-b/(u(i-k^*))} ,
\end{aligned}$$

which, given that $r \in \mathbb{Z}_p$ is chosen randomly by B , indeed comes from the same distribution from the perspective of F as the response given to it in its real experiment. Now, it is straightforward to verify that when B does not abort, it provides F with a view coming from an identical distribution to that in its real experiment during the rest of the simulation as well. ■

Now based on the code for B , conditions (1) and (3) of a forgery in Definition 3.2, and by using the bilinearity and non-degeneracy properties of a pairing, we have that on run of B on which it does not abort and on which F produces a forgery with signature (Q^*, R^*) , where $R^* = g^{r^*}$ for some $r^* \in \mathbb{Z}_p$ unknown to B , B 's output can be written as

$$\begin{aligned}
g^{ab} g^{\text{E}[m_j]a} ((g^a)^{-uk^*} g^t (g^a)^{k^*u})^{r^*} / ((g^a)^{\text{E}[m_j]} (g^{r^*})^t) &= g^{ab} g^{\text{E}[m_j]a} ((g^a)^{k^*-k^*} g^t)^{r^*} / ((g^a)^{\text{E}[m_j]} (g^{r^*})^t) \\
&= g^{ab} g^{\text{E}[m_j]a} (g^t)^{r^*} / ((g^a)^{\text{E}[m_j]} (g^{r^*})^t) \\
&= g^{ab} ,
\end{aligned}$$

which is a solution to its input CDH problem instance. In light of the above claim, then, we can wlog consider runs of the CDH game played by B and of the UF-OMS experiment with F using randomly-chosen coin sequences drawn from a common finite space of coins, where if F produces a forgery on a run of its experiment using some chosen sequence of coins from this space then, on a run of B using the same coins, the latter provides input and oracle replies to F identical to that its real UF-OMS experiment and hence outputs a solution to its CDH instance if it does not abort. Let forge be the event that F produces a forgery when run in its UF-OMS experiment. Let abort

be the probability that B aborts. Then we have the probability $\mathbf{Adv}_{\mathcal{G}}^{\text{CDH}}(B)$ that B succeeds in solving the CDH relative to \mathcal{G} is bounded as follows:

$$\begin{aligned} \mathbf{Adv}_{\mathcal{G}}^{\text{CDH}}(B) &= \Pr [\text{forge} \wedge \overline{\text{abort}}] \\ &= \Pr [\overline{\text{abort}} \mid \text{forge}] \cdot \Pr [\text{forge}] \\ &= \Pr [\overline{\text{abort}} \mid \text{forge}] \cdot \mathbf{Adv}_{\text{OMS}}^{\text{UF-OMS}}(F). \end{aligned}$$

Note that the probabilities above are taken over the choice of the coin sequence drawn from the common finite space, and that the last equality is by definition. To continue the analysis, we make the following claim.

Claim A.2 We claim that

$$\Pr [\overline{\text{abort}} \mid \text{forge}] \geq (1/n_{\max}) \cdot (1 - \delta) \cdot \delta^{q_s}.$$

Proof: Note that the probability that B aborts seems difficult to analyze directly, because F may repeat messages in its queries to its signing oracle, which have the same δ -values. Instead, we analyze it by looking at a run of B and of the UF-OMS experiment with F using the same coin sequence drawn from the common finite space of coins, on which F outputs a forgery on a message m^* in the latter. Let abort_s be the event that B aborts when responding to a signing query by F and abort_f be the event that B aborts after F outputs a forgery. Let goodf be the event that B sets $\delta[m^*] = 0$ and $k^* = i^*$, where i^* is the position of the honest signer in forgery that F outputs in its experiment when run on the same coin sequence. Then we have

$$\begin{aligned} \Pr [\overline{\text{abort}} \mid \text{forge}] &= \Pr [\overline{\text{abort}}_s \wedge \overline{\text{abort}}_f \mid \text{forge}] \\ &= \Pr [\overline{\text{abort}}_s \wedge \overline{\text{abort}}_f \wedge \text{goodf} \mid \text{forge}] \\ &= \Pr [\overline{\text{abort}}_f \mid \overline{\text{abort}}_s \wedge \text{goodf} \wedge \text{forge}] \cdot \Pr [\overline{\text{abort}}_s \wedge \text{goodf} \mid \text{forge}] \\ &= \Pr [\overline{\text{abort}}_s \wedge \text{goodf} \mid \text{forge}]. \end{aligned} \tag{2}$$

For the second equality above, we use the fact that $\overline{\text{abort}}_f$ occurs only if goodf does, by definition of B . For the last, we use that $\Pr [\overline{\text{abort}}_f \mid \overline{\text{abort}}_s \wedge \text{goodf} \wedge \text{forge}] = 1$, because abort_f cannot occur if goodf does. Now, consider the set of distinct messages $S = \{m_1, \dots, m_k\}$ that F queries to its signing oracle when executed in its experiment. We assume wlog that $m^* = m_k$, where m^* is the message on which F forges. Let badq_j be the event that F when executed by B makes a signing query of the form $m_j, \sigma, L = (pk_1, \dots, pk_{k^*-1})$ for which the reply is not \perp . We next claim the following sequence of inequalities:

$$\begin{aligned} \Pr [\overline{\text{abort}}_s \wedge \text{goodf} \mid \text{forge}] &= \Pr \left[\text{goodf} \wedge \bigwedge_{j=1}^k \delta[m_j] = 1 \vee (\delta[m_j] = 0 \wedge \overline{\text{badq}}_j) \mid \text{forge} \right] \\ &\geq \Pr \left[\text{goodf} \wedge (\delta[m_k] = 0 \wedge \overline{\text{badq}}_k) \bigwedge_{j=1}^{k-1} \delta[m_j] = 1 \mid \text{forge} \right] \\ &\geq \Pr \left[\text{goodf} \wedge \bigwedge_{j=1}^{k-1} \delta[m_j] = 1 \mid \text{forge} \right] \\ &\geq \frac{1}{n_{\max}} \cdot (1 - \delta) \cdot \delta^{k-1} \\ &\geq \frac{1}{n_{\max}} \cdot (1 - \delta) \cdot \delta^{q_s}. \end{aligned}$$

Above, the first equation is just by the definition of B . The next follows by dropping events in disjunctions. For the third, we use the fact that if `goodf` occurs and B sets all δ -values except $\delta[m^*]$ to 1, then `badk` (where $m^* = m_k$) cannot occur. This is because, by condition (4) in the definition of a forgery, F does not make a signing query of the form $m^*, \sigma, L = (pk_1, \dots, pk_{i^*-1})$ for any $\sigma \in \{0, 1\}^*$ on the run of its experiment and hence on the run of B , because the latter provides the same input and oracle replies to F as the former in this case when run on the same coins. The fourth follows from the fact that all B always sets k^* and any δ -values independently of each other and of F . Finally, the last uses $1 \leq k \leq q_s$ and $0 \leq \delta \leq 1$. Substituting the last inequality into equation (2) above proves the claim. \blacksquare

Using the above claim, we now have

$$\mathbf{Adv}_{\mathcal{G}}^{\text{CDH}}(B) \geq (1/n_{\max}) \cdot (1 - \delta) \cdot \delta^{q_s} \cdot \mathbf{Adv}_{\text{OMS}}^{\text{UF}}(F). \quad (3)$$

To finish the analysis, let us define for $0 \leq \delta \leq 1$ the function

$$f(\delta) \stackrel{\text{def}}{=} \delta^{q_s} \cdot (1 - \delta).$$

It is not hard to see that f is maximized at $\delta_{\text{OPT}} = q_s / (q_s + 1)$, at which $f(\delta_{\text{OPT}}) \geq 1 / (e(q_s + 1))$. Setting δ to δ_{OPT} in the description of B and substituting the above for $f(\delta)$ in equation (3), then re-arranging terms, gives equation (1) in Theorem 3.6.

Finally, to justify the running-time analysis of B , we take into account our convention to include in the running-time of F that of its overlying experiment. B 's extra work is one additional exponentiation on each hash query F makes, as well as an additional number of exponentiations linear in the number signers in the OMS (of which there are at most n_{\max}) on each signing query and on a forgery. This takes time at most $\tau(\mathcal{G}) \cdot O(q_h + n_{\max}(q_s + 1))$, as asserted. \blacksquare

B Proof of Theorem 4.3

We construct a algorithm B for solving the IBSAS-CDH problem that simulates the UF-IBSAS experiment for F .

THE SIMULATOR. The description of simulator B is given in Figure 2, under the following simplifying assumptions. We assume wlog that F never repeats a query to its hash oracles. For convenience we also assume that F makes hash queries to its H_1 and H_2 oracles at the same time; that is, it submits a query ID to both oracles and receives back both the values of $H_1(ID)$ and $H_2(ID)$. In responding to these hash queries, using Coron's technique [19] we have B assign query ID a bit (aka. δ -value) $\delta[ID]$ equal to 1 with some probability $0 \leq \delta \leq 1$ that we optimize later. Moreover, in B 's code, we assume for simplicity that when F makes a signing query $ID_i, m_i, ((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1})), \sigma$, it has previously queried identity ID_k to its H_1 and H_2 oracles for all $k \in \{1, \dots, i\}$; similarly, on F 's final output $L^* = ((ID_1^*, m_1^*), \dots, (ID_n^*, m_n^*)), \sigma^*$, we assume wlog it has queried identity ID_k to its H_1 and H_2 oracles and $ID_k || m_k$ to its H_3 oracle for all $k \in \{1, \dots, n\}$. Intuitively, B hopes that F does not ask for the secret key of any ID with $\delta[ID] = 0$, but that at most one such identity occurs in each signing query and exactly one such identity occurs in its forgery.

ANALYSIS. For the analysis, let `collide` be the event that B outputs (m', X, Y, Z) such that it has previously queried m' to its oracle, let `forge` be the event that F produces a forgery according to Definition 4.1 when executed by B , and let `abort` be the event that B aborts. (We emphasize that `forge` is defined with respect to an execution of the simulator B here, unlike in the proof

Simulator $B^{\mathcal{O}_{g,a_1,b_1,a_2,b_2}^{\text{IBSAS-CDH}}}$ ($p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g, g^{a_1}, g^{b_1}, g^{a_2}, g^{b_2}$)
Initialize arrays $E_1, E_2, \delta, H_1, H_2, H_3$ to everywhere undefined
Run F on input $p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, u, v, g, g^a, H_1, H_2, H_3$:

On H_1, H_2 query ID :
 $E_1[ID], E_2[ID] \xleftarrow{\$} \mathbb{Z}_p$; $\delta[ID] \xleftarrow{\delta} \{0, 1\}$
If $\delta[ID] = 1$ then $H_1[ID] \leftarrow g^{E_1[ID]}$; $H_2[ID] \leftarrow g^{E_2[ID]}$
Else $H_1[ID] \leftarrow g^{b_1} g^{E_1[ID]}$; $H_2[ID] \leftarrow g^{b_2} g^{E_2[ID]}$
Return $H_1[ID], H_2[ID]$

On H_3 -query x :
 $H_3[x] \xleftarrow{\$} \mathbb{Z}_p^*$; Return $H_3[x]$

On signing query $(ID_i, m_i, ((ID_1, m_1), \dots, (ID_{i-1}, m_{i-1})), \sigma)$:
If the input is not a valid signature then return \perp
Parse σ as (X, Y, Z) and set it as $(1_{\mathbb{G}}, 1_{\mathbb{G}}, 1_{\mathbb{G}})$ if $i = 1$
If there exists $\delta[ID_j] = 0$ for some $1 \leq j \leq i$ then
 If there exists $\delta[ID_{j'} = 0]$ for some $j' \neq j$ then abort
 Else $(X', Y', Z') \xleftarrow{\$} \mathcal{O}_{g,g^{a_1},g^{b_1},g^{a_2},g^{b_2}}^{\text{IBSAS-CDH}}(H_3[ID_j || m_j])$
For all $1 \leq \ell \leq i$ do
 $X' \leftarrow X' \cdot (g^{a_1})^{E_1[ID_\ell]} (g^{a_2})^{E_2[ID_\ell]} H_3[ID_\ell || m_\ell]$
Return (X', Y', Z')

On key-derivation query ID :
If $\delta[ID] = 0$ then abort ; Else return $((g^{a_1})^{E_1[ID]}, (g^{a_2})^{E_2[ID]})$

Let $(L^* = ((ID_1^*, m_1^*), \dots, (ID_n^*, m_n^*)), \sigma^*)$ be the output of F
If L^*, σ^* is not a forgery then return \perp
Let $i^* \in \{1, \dots, n\}$ satisfy condition (2) of a forgery
If $\delta[ID_{i^*}] = 1$ or there exists $\delta[ID_{i'}] = 0$ for some $1 \leq i' \neq i^* \leq n$ then abort
Parse σ^* as (X, Y, Z) and for all $1 \leq j \leq n$ do
 $X \leftarrow X / ((g^{a_1})^{E_1[ID_j]} (g^{a_2})^{E_2[ID_j]} H_3[ID_j || m_j])$
Return $(H_3[ID_{i^*} || m_{i^*}], X, Y, Z)$

Figure 2: Simulator B for the proof of Theorem 4.3.

of Theorem 3.6 in Appendix A.) We claim that the probability $\mathbf{Adv}_{\mathcal{G}}^{\text{IBSAS-CDH}}(B)$ that B succeeds in solving the IBSAS-CDH problem relative to \mathcal{G} is bounded as follows:

$$\begin{aligned}
\mathbf{Adv}_{\mathcal{G}}^{\text{IBSAS-CDH}}(B) &\geq \Pr [\text{forge} \wedge \overline{\text{collide}} \wedge \overline{\text{abort}}] \\
&= \Pr [\text{forge} \wedge \overline{\text{collide}} \mid \overline{\text{abort}}] \cdot \Pr [\overline{\text{abort}}] \\
&= \Pr [\text{forge} \setminus \text{collide} \mid \overline{\text{abort}}] \cdot \Pr [\overline{\text{abort}}] \\
&= (\Pr [\text{forge} \mid \overline{\text{abort}}] - \Pr [\text{collide} \mid \overline{\text{abort}}]) \cdot \Pr [\overline{\text{abort}}] . \quad (4)
\end{aligned}$$

To see the first inequality above, consider a run of B in which it does not abort, and suppose F 's output $L^* = ((ID_1^*, m_1^*), \dots, (ID_n^*, m_n^*)), \sigma^*$ when executed by B is a forgery. Let $i^* \in \{1, \dots, n\}$ satisfy condition (2) of a forgery. Based on the description of B , condition (1) of a forgery, and by using the bilinearity and non-degeneracy properties of a pairing, we have that B 's output can be

written as

$$\left(m', g^{r'x'} g^{a_1 b_1} g^{m' a_2 b_2}, g^{r'}, g^{x'} \right)$$

for $r', x' \in \mathbb{Z}_p$ unknown to B , where $m' = H_3[ID_{i^*} || m_{i^*}]$. So, the output is a solution to the IBSAS-CDH problem (meaning causes the IBSAS-CDH game to output 1) if B did not previously query $H_3[ID_{i^*} || m_{i^*}]$ to its oracle, i.e. collide did not occur. This gives us the first inequality. To lower-bound the last line, we use the following claims.

Claim B.1 We claim that

$$\Pr [\overline{\text{abort}}] \geq \delta^{n_{\max} q_s} \delta^{q_k} (1 - \delta) \delta^{n_{\max} - 1}.$$

Proof: To see this let us consider runs of the IBSAS-CDH game played by B and the UF-IBSAS experiment with F using coin sequences drawn from a common finite set of coins. Such a coin sequence defines an identity and message ID^*, m^* for which F forges, if any, and a set $\{ID_1, \dots, ID_k\}$ comprised of distinct identities query by F to its key-derivation oracle or occurring in the lists used in its signing queries and in its forgery. (If F does not successfully forge in its experiment on this coin sequence, just choose ID^* arbitrarily.) Wlog suppose $ID_1 = ID^*$. Then for B not to abort it suffices that $\delta[ID_1] = 0$ and $\delta[ID_i] = 1$ for all $1 < i \leq k$. Since these δ -values are all chosen independently at random this occurs with probability at least $(1 - \delta) \delta^{k-1}$. Using $k \leq n_{\max} q_s + (n_{\max} - 1)$ yields the claim. ■

Claim B.2 We claim that

$$\Pr [\text{forge} | \overline{\text{abort}}] = \mathbf{Adv}_{\text{AS}}^{\text{UF-IBSAS}}(F).$$

Proof: We again consider runs of the IBSAS-CDH game played by B and of the UF-IBSAS experiment with F using randomly-chosen coin sequences drawn from a common finite space of coins. Imagine a new game where we first run B using a randomly-chosen coin sequence from this space and then run the UF-IBSAS experiment with F using the same coins. Note that on executions of B where it does not abort, F 's view in the simulation comes from a distribution identical to that in its real experiment. So, the probability of forge given that abort does not occur is the same as the probability that F outputs a forgery when run in the new game given that B did not abort when run in this game, which is clearly just $\mathbf{Adv}_{\text{AS}}^{\text{UF-IBSAS}}(F)$ as desired. ■

Claim B.3 We claim that

$$\Pr [\text{collide} | \overline{\text{abort}}] \leq \frac{q_{h_2}(q_{h_2} - 1)}{2p}.$$

Proof: As before, consider runs of the IBSAS-CDH game played by B in which B does not abort and of the UF-IBSAS experiment with F using the same sequence of coins drawn from the common finite space. Let q_1, \dots, q_k be the queries made by F in the latter to its H_3 oracle. If $H_3[q_1], \dots, H_3[q_k]$ are all distinct then this implies collide does not occur. Since the former are chosen independently at random we can apply a standard birthday bound for the claim. ■

Plugging the previous claims into equation (4), we now have

$$\mathbf{Adv}_{\mathcal{G}}^{\text{CDH-IBSAS}}(B) \geq \left(\mathbf{Adv}_{\text{AS}}^{\text{UF-IBSAS}}(F) - \frac{q_{h_2}(q_{h_2} - 1)}{2p} \right) \cdot \delta^{n_{\max} q_s + q_k} (1 - \delta) \delta^{n_{\max} - 1}. \quad (5)$$

To complete the analysis, let us define the function

$$\begin{aligned} f(\delta) &\stackrel{\text{def}}{=} \delta^{n_{\max}q_s+q_k}(1-\delta)\delta^{n_{\max}-1} \\ &= \delta^{n_{\max}(q_s+1)+q_k-1}(1-\delta). \end{aligned}$$

Denoting the exponent $n_{\max}(q_s+1)+q_k-1$ by z , it is not hard to see that f is maximized at $\delta_{\text{OPT}} = z/(z+1)$, for which we have $f(\delta_{\text{OPT}}) \geq 1/(e(z+1)) = 1/(e(n_{\max}(q_s+1)+q_k))$. Setting δ to δ_{OPT} in the code for B and substituting the above for $f(\delta)$ in (5) then re-arranging the inequality yields equation (2) in Theorem 4.3.

Finally, to justify our running-time analysis of B , recall our convention to include in the running-time of F that of its overlying experiment. On answering H_1, H_2 and key-derivation queries by F , B 's overhead is at most one exponentiation in \mathbb{G} . On each signing query, B 's overhead is a linear (in n_{\max}) number of exponentiations in \mathbb{G} . After a forgery, B 's overhead is again at most a linear number of exponentiations in \mathbb{G} , giving total time $\tau(\mathcal{G}) \cdot O(q_h + q_k + n_{\max}(q_s+1))$. Moreover, the number of queries made by B to its oracle is at most q_s , as asserted. ■

C Proof of Theorem 5.1

Here we provide the crux of the proof. Based on it the theorem can be obtained in standard ways, using the Schwartz-Zippel Lemma [41].

C.1 A Useful Lemma

Let us first establish the following useful lemma. To state it let us call a degree-2 monomial in $\mathbb{Z}_p[x_1, \dots, x_n]$ *odd consecutive* if it has the form $x_i x_{i+1}$ for some *odd* i , and otherwise *non-odd-consecutive*.

Lemma C.1 For p prime and $n \geq 2$, there are no linear polynomials $r = r(x_1, \dots, x_n), s = s(x_1, \dots, x_n) \in \mathbb{Z}_p[x_1, \dots, x_n]$ such that in the product rs at least two of the odd consecutive degree-2 monomials have non-zero coefficients, but all the coefficients of non-odd-consecutive degree-2 monomials are zero.

Proof: Assume that there are such polynomials. Let r_i be the coefficient of x_i in polynomial r ; define s_i analogously. By assumption there exist odd $1 \leq a \neq b \leq n-1$ such that the monomials $x_a x_{a+1}$ and $x_b x_{b+1}$ in rs have non-zero coefficients. Since $r_a s_{a+1} + r_{a+1} s_a$ (the coefficient of $x_a x_{a+1}$ in rs) is non-zero, either $r_a s_{a+1}$ or $r_{a+1} s_a$ is non-zero, meaning that at least one of the pairs $(r_a, s_{a+1}), (r_{a+1}, s_a)$ contains two non-zero coefficients. Similarly for (r_b, s_{b+1}) and (r_{b+1}, s_b) . On the other hand, since $r_i s_i = 0$ for all i (since the coefficient of x_i^2 in rs is 0 for all i), we have that, for each i , at most one of r_i or s_i can be non-zero. Therefore, *exactly* one pair (r_a, s_{a+1}) or (r_{a+1}, s_a) has two non-zero coefficients, and the other pair has both equal to 0; similarly for (r_b, s_{b+1}) and (r_{b+1}, s_b) . Suppose (r_a, s_{a+1}) and (r_b, s_{b+1}) are all non-zero; then, $r_a s_{b+1}$ (the coefficient of $x_a x_{b+1}$ in rs , since $s_a r_{b+1}$ is assumed zero) is non-zero, a contradiction since $x_a x_{b+1}$ is non-odd-consecutive. A similar analysis pertains to the other three possibilities. ■

C.2 Main Argument

Essentially, in the generic bilinear group model we can represent (the discrete logs of) group elements given to the adversary as multivariate polynomials over \mathbb{Z}_p , where each randomly-chosen

exponent unknown to the adversary is a formal indeterminate. The group operations available to the adversary correspond to computing linear combinations of these polynomials. (The adversary can also use the bilinear map operation, but since the result lies in a different group this operation can be effectively ignored here.) We need to argue that the adversary cannot win the IBSAS-CDH game in this model, or equivalently that a solution to the IBSAS-CDH problem is linearly independent from the polynomials the adversary is given. In the remainder we denote indeterminates using capital letters and scalar variables (i.e. coefficients) using Greek letters.

THE GAME. In the IBSAS-CDH game the adversary is initially given polynomials $1, A_1, B_1, A_2, B_2$, where the indeterminates A_1, B_1, A_2, B_2 represent the discrete logs a_1, b_1, a_2, b_2 chosen in the IBSAS-CDH game. Observe that a solution to the IBSAS-CDH problem can be respresented by non-zero μ and polynomials $poly_1, poly_2, poly_3$ over \mathbb{Z}_p such that

$$poly_1 + poly_2 \cdot poly_3 = A_1 B_1 + \mu A_2 B_2, \quad (6)$$

where the adversary did not query μ to its oracle. The overall strategy for the proof is to assume the adversary outputs such a solution and derive a contradiction.

THE NO-QUERY CASE. As a warm-up (that will be useful to us later), let us first consider an adversary that makes no oracle queries. In this case we may write

$$poly_j = \alpha_{j,1} A_1 + \beta_{j,1} B_1 + \alpha_{j,2} A_2 + \beta_{j,2} B_2 + \gamma_j$$

for each j . As $poly_1$ contains no degree-2 monomial with non-zero coefficient, the adversary must have produced $poly_2, poly_3$ such that in $poly_2 \cdot poly_3$ the monomial $A_1 B_1$ has coefficient 1 and $A_2 B_2$ has coefficient μ . Applying Lemma C.1 with $n = 4, x_1 = A_1, x_2 = B_1, x_3 = A_2, x_4 = B_2$ yields the desired contradiction.

THE GENERAL CASE. Now let us consider an adversary that makes up to q oracle queries. On the i -th query it submits a value $\mu_i \in \mathbb{Z}_p$ and receives

$$(X_i R_i + A_1 B_1 + \mu_i A_2 B_2, X_i, R_i).$$

We may write

$$\begin{aligned} poly_j &= \alpha_{j,1} A_1 + \beta_{j,1} B_1 + \alpha_{j,2} A_2 + \beta_{j,2} B_2 + \gamma_j \\ &\quad + \sum_{k=1}^q \delta_{j,k} X_k + \zeta_{j,k} R_k + \eta_{j,k} (X_k R_k + A_1 B_1 + \mu_k A_2 B_2) \end{aligned} \quad (7)$$

for each j . In fact, we have the following claim.

Claim C.2 We may assume wlog that $\eta_{j,k}$ is zero for $j = 2, 3$ and all k (i.e., that $poly_2, poly_3$ are *linear*).

The proof is below. Observe (by equating coefficients) that for (6) to hold the projection⁴ of $poly_2 \cdot poly_3$ onto monomials $X_i R_i$ for all i must be $\sum_{i=1}^q -\eta_{1,i} X_i R_i$. Furthermore, its only possible degree-2 monomials in A_1, B_1, A_2, B_2 with non-zero coefficients are $A_1 B_1$ and $A_2 B_2$. Thus, if the above sum has at least two non-zero terms, then (using the above claim) $poly_2, poly_3$ contradict Lemma C.1 with $n = 2q + 4, x_i = X_i$ and $x_{i+1} = R_i$ for all odd $1 \leq i \leq 2q$ and $x_{2q+1} = A_1, x_{2q+2} = B_1, x_{2q+3} = A_2, x_{2q+4} = B_2$. So, there is at most *one* value of k for which $\eta_{1,k}$ is non-zero. If there is no such value of k , then we can conclude via the no-query analysis above. Suppose there is exactly one such value, say k^* . We use the following claim.

⁴A *projection* of a polynomial onto specified monomials means we zero-out the coefficients of all other monomials.

Claim C.3 If (6) holds and not all $\eta_{1,k}$ are zero, then the coefficient of A_1B_1 in $poly_2 \cdot poly_3$ is zero.

Again, the proof is below; let us first see that the result follows. Indeed, it follows by equating coefficients of A_1B_1 in (6) and the above claim that $\sum_{k=1}^q \eta_{1,k} = 1$, so $\eta_{1,k^*} = 1$ since *exactly* η_{1,k^*} , and no other $\eta_{1,k}$ for $k \neq k^*$, is assumed non-zero. Thus $\mu = \mu_{k^*}$. This contradicts the fact that the adversary did not query μ to its oracle (since it did so in its k^* -th query). So it remains to prove the two claims.

Proof: (of Claim C.2) Given an IBSAS-CDH solution $(poly_1, poly_2, poly_3, \mu)$, we give another one $(poly'_1, poly'_2, poly'_3, \mu)$ in which $poly'_2, poly'_3$ are linear; moreover, the new solution can be output by a generic bilinear group adversary if the old one can. First note that there are two disjoint cases: (1) $poly_2$ is non-linear, or (2) $poly_3$ is non-linear. They are disjoint because, say, in case (1) $poly_3$ must be a *constant* (i.e., $poly_3 = \gamma_3$) as otherwise the LHS of (6) would have total degree at least 3, contradicting the RHS; similarly in case (2). In case (1) define $poly'_1, poly'_2$ to be the same as $poly_1, poly_2$ respectively, except that for every non-zero $\eta_{2,i}$ add $\eta_{2,i} \cdot \gamma_3$ to the coefficient of the term $X_iR_i + A_1B_1 + \mu_1A_2B_2$ in $poly'_1$ (written as in (7)) and change the coefficient of this term in $poly'_2$ to zero. Also define $poly'_3 = poly_3$. It is easy to verify that this gives an IBSAS-CDH solution with the desired properties. A similar transformation applies in case (2). ■

Proof: (of Claim C.3) Suppose (6) holds and not all $\eta_{1,k}$ are zero, but the coefficient of A_1B_1 in $poly_2 \cdot poly_3$ is non-zero. Recall from above that for (6) to hold the projection of $poly_2 \cdot poly_3$ onto monomials X_iR_i for all i must have the form $\sum_{i=1}^q -\eta_{1,i}X_iR_i$. Furthermore, its only possible degree-2 monomials in A_1, B_1, A_2, B_2 with non-zero coefficients are A_1B_1 and A_2B_2 . Here the above sum has at least one non-zero term and A_1B_1 is non-zero by assumption. Thus, similarly to before $poly_2, poly_3$ contradict Lemma C.1 with $n = 2q + 4$, $x_i = X_i$ and $x_{i+1} = R_i$ for all odd $1 \leq i \leq 2q$ and $x_{2q+1} = A_1, x_{2q+2} = B_1, x_{2q+3} = A_2, x_{2q+4} = B_2$. ■ ■